IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

DYNAMIC MESH NETWORKS, INC. DBA MESHDYNAMICS, Plaintiff, v. CISCO SYSTEMS, INC., Defendant.	Case No. 2:25-cv-00781-JRG (Lead Case) Jury Trial Demanded

MESHDYNAMICS'S LOCAL PATENT RULES 3-1 AND 3-2 INITIAL DISCLOSURES

Pursuant to Local Patent Rules (L.P.R.) 3-1 and 3-2 of the United States District Court for the Eastern District of Texas, Plaintiff Dynamic Mesh Networks, Inc. ("MeshDynamics" or "Plaintiff") submits the following initial patent disclosures with respect to U.S. Patent Nos. 7,420,952 ("the '952 Patent"), 7,885,243 ("the '243 Patent"), 7,894,385 ("the '385 Patent"), 8,520,691 ("the '691 Patent"), and 11,368,537 ("the '537 Patent") (collectively, the "MeshDynamics Asserted Patents").

I. INTRODUCTION

This case is in the early stages and discovery and investigation in connection with this lawsuit has not begun. These contentions are thus based on information reasonably available to MeshDynamics to-date. These contentions apply the plain and ordinary meaning to the claim terms and will be supplemented should the Court issue a *Markman* order construing any term of the MeshDynamics Asserted Patents. MeshDynamics has made reasonable assumptions, to the extent necessary and appropriate, with respect to the meaning of patent claim terms for the purpose of preparing these contentions. MeshDynamics reserves the right to rely upon different meanings in the conduct of this litigation, to assert different meanings as appropriate, and revise the factual content of these contentions in light of any such different meanings.

MeshDynamics's contentions are subject to revision, modification and/or supplementation as MeshDynamics's investigation and discovery in the case progresses. Many claim elements are software limitations (as noted in the attached claim charts), and Cisco Systems, Inc. ("Cisco" or "Defendant") has yet to provide access to the relevant source code in this case. Pursuant to Judge Gilstrap's discovery order paragraph 3(a), MeshDynamics will supplement these contentions following the production and review of the relevant source code. MeshDynamics further reserves the right to revise, modify and/or supplement these contentions in response to any invalidity contentions. MeshDynamics further reserves the right to revise, modify, amend and/or supplement

these contentions in accordance with the Local Patent Rules, any other applicable Rules, or by order of this Court.

In addition, MeshDynamics reserves its right to amend, alter, or supplement these contentions at any time and/or as otherwise authorized or permitted by the aforesaid Rules, based on the following:

- Any further investigation;
- Any fact or expert discovery materials, including, but not limited to, documents, source code, reports, guides, communications and deposition testimony, which have not yet been obtained, produced or provided to MeshDynamics;
- Evaluation of the scope and content of the prior art;
- Clarification of Cisco's non-infringement positions;
- Any claim construction order;
- Any claim contentions that Cisco may assert that are different from any assumptions MeshDynamics has made in preparing these contentions;
- Any claim constructions or additional or different allegations that Cisco may assert at a later date; or
- Any other grounds otherwise provided under the Local Patent Rules, any other applicable Rules, upon leave of the Court, and/or any other reason requiring same.

These contentions may be asserted in the alternative and do not constitute any concession by MeshDynamics for purposes of claim construction, non-infringement, or invalidity. *See* Fed. R. Civ. P. 8(d).

Furthermore, these contentions are provided without prejudice to the rights of MeshDynamics to introduce at trial any subsequently-discovered evidence or expert opinions and to produce and introduce at trial all evidence, whenever discovered, relating to the proof of subsequently-discovered facts. Moreover, facts, documents and things now known may be imperfectly understood and, accordingly, such facts, documents and things may not be included in these contentions. In particular, Cisco's publicly available information is unclear and ambiguous as to the operation of its products. MeshDynamics reserves its right to refer to, conduct discovery with reference to, or offer into evidence at the time of trial, any and all facts, expert opinion

testimony, documents and things notwithstanding the written statements herein and subject to the Rules of this Court. MeshDynamics further reserves its right to refer to, conduct discovery with reference to, or offer into evidence at the time of trial, any and all facts, documents and things that are not currently recalled but might be recalled at some time in the future.

MeshDynamics provides these contentions without waiving any claim of attorney-client privilege, attorney work-product immunity, or any other applicable privilege or immunity. To the extent that MeshDynamics inadvertently discloses information that may be protected from discovery under the attorney-client privilege, the attorney work-product immunity, or any other applicable privilege or immunity, such inadvertent disclosure does not constitute a waiver of any such privilege or immunity.

The information set forth below is provided without waiver of (1) the right to object to the use of any statement for any purpose, in this action or in any other action, on the grounds of privilege, relevance, materiality or any other appropriate grounds; (2) the right to object to any request involving or relating to the subject matter of the statements herein; or (3) the right to revise, correct, supplement or clarify any of the statements provided below at any time.

Subject to the foregoing, MeshDynamics hereby submits the following initial disclosures in accordance with L.P.R. 3-1 and L.P.R. 3-2.

II. INFRINGEMENT CONTENTIONS

Pursuant to L.P.R. 3-1, and subject to the reservation of rights outlined above, MeshDynamics provides below, and in the charts attached as Exhibits 1-10, MeshDynamics's infringement contentions.

A. Identification of the Asserted Claims

Pursuant to L.P.R. 3-1(a), the asserted claims of the MeshDynamics Asserted Patents include those claims identified in the attached claim charts. MeshDynamics identified these

asserted claims based on its current and preliminary understanding and reserves the right to supplement its identification of infringed claims as discovery proceeds, including identifying additional claims following review of Cisco's proprietary information.

MeshDynamics accuses Cisco of infringing under 35 U.S.C. § 271(a), 35 U.S.C. § 271(b), and 35 U.S.C. § 271(c) each of the following claims of the MeshDynamics Asserted Patents (hereinafter referred to as the "Asserted Claims"):

- Claims 1-21 of the '952 Patent
- Claims 1-13 of the '243 Patent
- Claims 1-3 of the '385 Patent
- Claims 1, 9-11, 15-16, 18-19, 35 of the '691 Patent
- Claims 1-19 of the '537 Patent

Cisco has infringed and continues to infringe (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more of the Asserted Claims of the MeshDynamics Asserted Patents by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Accused Instrumentalities ("Cisco Accused Products") as detailed in the attached claim charts. Cisco has indirectly infringed and continues to indirectly infringe (literally and/or under the doctrine of equivalents) one or more of the Asserted Claims (which were all known to it well before this litigation was commenced) by making, using, offering for sale, selling, and/or importing in or into the United States the Cisco Accused Products, and by actively and intentionally inducing its customers and end users to infringe by making, using, offering for sale, selling, and/or importing the Cisco Accused Products as detailed in the attached claim charts. Cisco has also indirectly infringed and continues to indirectly infringe by offering for sale, selling, and/or

importing in or into the United States components that are material parts of the Accused Instrumentalities and are not staple articles of commerce, with knowledge of their infringing use, and without substantial non-infringing uses as detailed in the attached claim charts. *See* Exhibits 1-10. MeshDynamics incorporates fully herein the factual and infringement allegations set forth in its Complaint (*see* Dkt. 1).

B. Identification of the Accused Instrumentalities

Pursuant to L.P.R. 3-1(b), the Accused Instrumentalities ("Cisco Accused Products") include certain of Cisco's controllers, access points, cloud architecture, and software made, used, offered for sale, sold and/or imported in or into the United States for operation in or for Cisco's Wireless Mesh Networking. Attached as Exhibit 11 is the list of Accused Instrumentalities of which MeshDynamics is currently aware. MeshDynamics asserts that any other Cisco product(s) that operates in the same or similar manner as those specifically identified in Exhibit 11 are also Accused Instrumentalities. MeshDynamics reserves all rights to revise, amend, and/or supplement this list as discovery and its investigation in this case proceeds. MeshDynamics asserts that all Accused Instrumentalities identified on Exhibit 11, and any other Cisco product(s) that operates in the same or similar matter, infringe the Asserted Claims of each Asserted Patent as identified in Section II.A *supra.*[1] *See* Exhibits 1-10.

C. Infringement Charts

Pursuant to L.P.R. 3-1(c), a chart identifying specifically where each element of each Asserted Claim is found within representative Accused Instrumentalities, including for each

[1] MeshDynamics does not assert Claim 1 of the '385 Patent against Cisco's Meraki Wireless Mesh Network (and related products). MeshDynamics only asserts Claim 1 of the '385 Patent against Cisco's Ultra Reliable Wireless Backhaul (and related products). MeshDynamics does not assert Claims 15-16, 18-19 of the '691 Patent against Cisco's Unified Wireless Mesh Network (and related products). MeshDynamics only asserts Claims 15-16, 18-19 of the '691 Patent against Cisco's Meraki Wireless Mesh Network (and related products).

element that such party contends is governed by 35 U.S.C. § 112 (¶ 6), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function, is attached for the MeshDynamics Asserted Patents. *See* Exhibits 1-10.

The claim charts are based on public information reasonably available to MeshDynamics, and do not include any confidential or source code information in Cisco's possession, custody or control. MeshDynamics reserves all rights to revise, amend, and/or supplement these charts as discovery proceeds, including but not limited to adding citations to evidence and source code for software limitations, whose operation and structure is not publicly available and cannot be known to MeshDynamics at this time. Accordingly, pursuant to the Court's standing Discovery Order, MeshDynamics reserves the right to supplement these charts once such information is made available to MeshDynamics.

On information and belief, Cisco's Unified Wireless Mesh Network mesh access points employ Cisco's proprietary Adaptive Wireless Path Protocol (AWPP).[2] To the extent that any of the technology present in the AWPP has evolved, been modified, incorporated in, or replaced by another same or similar operating Cisco proprietary protocol or enhancement, such protocol or enhancement infringed or infringes for the same reasons as AWPP and as detailed in the attached claim charts. Details regarding any other protocol used in mesh networking for Cisco Unified Wireless Networking other than AWPP are not publicly available and not accessible to

[2]Recent guides discuss Catalyst and Aironet mesh access points using AWPP in Cisco's Unified Wireless Mesh Networks. *See, e.g.*,
https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html;
https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-11/config-guide/b_wl_17_eleven_cg/m_mesh_ewlc.html;
https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html;
https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf.

MeshDynamics to fully chart at this time.

Furthermore, MeshDynamics reserves the right to revise these charts, as appropriate, upon issuance of any Claim Construction Order, including the identification of whether any arguable differences between the Accused Instrumentalities and the claims are insubstantial and/or whether the Accused Instrumentalities read on the claim elements under the doctrine of equivalents.

The claim charts are not intended to be an expert report on infringement or provide detailed analysis of the claim terms or infringement. MeshDynamics will disclose and produce an expert report regarding infringement with appropriate analysis pursuant to Fed. R. Civ. P. 26(a)(2), applicable Federal and Local Rules, and this Court's Orders. These charts include publicly available documentary evidence by way of example and not limitation.

D. Literal Infringement and Infringement Under the Doctrine of Equivalents

Pursuant to L.P.R. 3-1(d), MeshDynamics contends that the Accused Instrumentalities literally infringe or infringed the Asserted Claims of the MeshDynamics Asserted Patents as more specifically explained in the attached claim charts. *See* Exhibits 1-10. With respect to any claim element(s) that may be found not to be literally infringed, MeshDynamics contends in the alternative that the Cisco Accused Products infringed or infringe such claim element(s) under the doctrine of equivalents and that any claim element(s) not found to be literally met is equivalently met because any difference between the claim element(s) and the Accused Instrumentalities is not a substantial difference. In other words, to the extent literal infringement is purportedly absent, the Accused Instrumentalities identified in the attached claim charts perform substantially the same function of the invention recited in the Asserted Claims, in substantially the same way, to achieve substantially the same result. Accordingly, MeshDynamics contends that if any of the Accused Instrumentalities is found not to have infringed any Asserted Claims literally it is nevertheless embodied by the Accused Instrumentalities under at least one operative doctrine of equivalents

test (*e.g.*, function-way-result or insubstantial difference). *See* Exhibits 1-10.

MeshDynamics reserves the right to add allegations of infringement under the doctrine of equivalents or amend its doctrine of equivalents allegations in response to discovery, Cisco's claim construction and non-infringement positions, claim construction proceedings, and this Court's claim construction rulings or other rulings by the Court.

E. Priority Dates

Pursuant to L.P.R. 3-1(e), MeshDynamics identifies the priority date to which each Asserted Claim is entitled:[3]

- Claims 1-21 of the '952 Patent: The asserted claims of the '952 Patent are entitled to a priority date at least as early as October 28, 2002.
- Claims 1-3 of the '385 Patent: The asserted claims of the '385 Patent are entitled to a priority date at least as early as June 19, 2006.
- Claims 1-13 of the '243 Patent: The asserted claims of the '243 Patent are entitled to a priority date at least as early as October 28, 2002 (*i.e.*, the priority date of the '952 Patent).
- Claims 1, 9-11, 15-16, 18-19, 35 of the '691 Patent: The asserted claims of the '691 Patent are entitled to a priority date at least as early as October 28, 2002 (*i.e.*, the priority date of the '952 Patent), or in the alternative at least as early as January 30, 2009.

[3] MeshDynamics asserts that each of the MeshDynamics Asserted Patents is entitled to a priority date *as least as early as* the earliest priority date identified in its priority claims to the Patent Office (*i.e.*, shown on the front of the patent). In the event of any invalidity challenge that raises the issue, either in this litigation or before the Patent Office, MeshDynamics reserves the right to rely on the documentation identified below and/or corroborating testimony supporting earlier dates of conception and/or reduction to practice to establish priority dates earlier than those listed here.

- Claims 1-19 of the ’537 Patent: The asserted claims of the ’537 Patent are entitled to a priority date at least as early as October 28, 2002 (*i.e.*, the priority date of the ’952 Patent).

F. MeshDynamics Embodiments

Pursuant to L.P.R. 3-1(f), MeshDynamics reserves its right to rely upon its products as embodying all or some of the Asserted Claims and identifies at least the following instrumentalities, which practice claim(s) of the MeshDynamics Asserted Patents. MeshDynamics reserves the right to revise, amend, or supplement this identification following any claim construction proceedings. MeshDynamics’s identification of embodiments does not constitute any concession regarding the applicability of 35 U.S.C. § 287.

- MeshDynamics’s legacy MD-4000 product line, specifically the MD-4350 configuration, was an embodiment of claim 12 of the ’952 Patent.

III. DOCUMENT PRODUCTION UNDER L.P.R. 3-2

Pursuant to L.P.R. 3-2 and subject to the reservation of rights outlined above, MeshDynamics is producing concurrently with service of these infringement contentions the following documents:

a) Documents sufficient to evidence each discussion with, disclosure to, or other manner of providing to a third party, or sale of or offer to sell, the claimed invention prior to the date of application for the patent in suit. A party’s production of a document as required herein shall not constitute an admission that such document evidences or is prior art under 35 U.S.C. §102. *See* MeshD00052012-52040.

b) All documents evidencing the conception, reduction to practice, design, and development of each claimed invention, which were created on or before the date of application for the patent in suit or the priority date identified pursuant to L.P.R. 3-1(e),

whichever is earlier: MeshD00036313-341, MeshD00032412-32824 (including the identified source code), MeshD00037659-676, MeshD00052011, MeshD00052012-040, MeshD00052044-063, and MeshD00052075. Source code and simulation available on a standalone computer, which will be made available as specified under the Protective Order to be entered in the case, may also be relied on for evidencing the conception, reduction to practice, design, and development of each claimed invention.

c) A copy of the file history for each patent in suit. *See* MeshD00000001-1848; *see also* MeshD00032412-32824.

Dated: November 10, 2025

Respectfully submitted,

/s/ Elizabeth Bernard

Jennifer Parker Ainsworth
TX State Bar No. 00784720
jainsworth@wilsonlawfirm.com
WILSON ROBERTSON & VANDEVENTER, P.C.
909 ESE Loop 323, Suite 400
Tyler, TX 75701
Telephone: (903) 509-5000

Scott R. Samay* (*pro hac vice*)
Jason S. Charkow* (*pro hac vice*)
Elizabeth Bernard* (admitted DC Bar No. 986675)
Oded Burger* (admitted NYS Bar No. 4910808)
Lisa Phillips* (*pro hac vice*)
Taylor Lepore* (*pro hac vice*)
Erin Hadi* (*pro hac vice*)
Louay Meroueh* (*pro hac vice*)
ssamay@daignaultiyer.com
jcharkow@daignaultiyer.com
ebernard@daignaultiyer.com
oburger@daignaultiyer.com
lphillips@daignaultiyer.com
tlepore@daignaultiyer.com
ehadi@daignaultiyer.com
lmeroueh@daignaultiyer.com

DAIGNAULT IYER LLP
8229 Boone Blvd - Suite 450
Vienna, VA 22182
*Not admitted to practice in Virginia

Attorneys for Plaintiff Dynamic Mesh Networks, Inc. dba MeshDynamics

CERTIFICATE OF SERVICE

I certify that a copy of **MESHDYNAMICS'S LOCAL PATENT RULES 3-1 AND 3-2 INITIAL DISCLOSURES** was served on November 10, 2025 via email at the email addresses indicated below:

Tamir Packin
tpackin@desmaraisllp.com
Lindsey Miller
lmiller@desmaraisllp.com
Michael R Rhodes
mrhodes@desmaraisllp.com
Emily Weber
eweber@desmaraisllp.com
Ryan Thorne
rthorne@dllp.com
Ha Vu Thu Huynh
hhuynh@desmaraisllp.com

Melissa Richards Smith
melissa@gillamsmithlaw.com

Attorneys for Defendant Cisco Systems, Inc.

/s/ Elizabeth Bernard
Elizabeth Bernard
Counsel for MeshDynamics

Exhibit 1 - U.S. Patent No. 7,420,952 (Cisco Wireless Mesh Networks with Catalyst and Aironet Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Wireless Mesh Network including but not limited to Cisco's Catalyst and Aironet Access Points, with related controllers and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1-21 (the "Asserted Claims") of U.S. Patent No. 7,420,952 (the "'952 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '952 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco's access points directly infringed at least Claim 12 of the '952 Patent when made, used, sold, offered for sale and imported into the United States. Cisco further directly infringed by testing, setting up, and using the Cisco Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Wireless Mesh Network in the United States.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the '952 Patent and make and use the infringing Cisco Wireless Mesh Network. Specifically, Cisco's Wireless Mesh Access Points, Design and Deployment Guides, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco's website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-21 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Wireless Mesh Network, and the infringing configuration and operation of the controller, access points, and software that comprise the infringing Cisco Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Wireless Mesh Network (and

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g.*, https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the ’952 Patent since 2009 when Cisco executives met with MeshDynamics to discuss partnering opportunities and discussed the ’952 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’952 Patent, or in the alternative, was willfully blind to the ’952 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’952 Patent. Cisco Accused Products come enabled as mesh access points and automatically connect to the controller to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particularly suited for Cisco’s Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringed the ’952 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the ’952 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

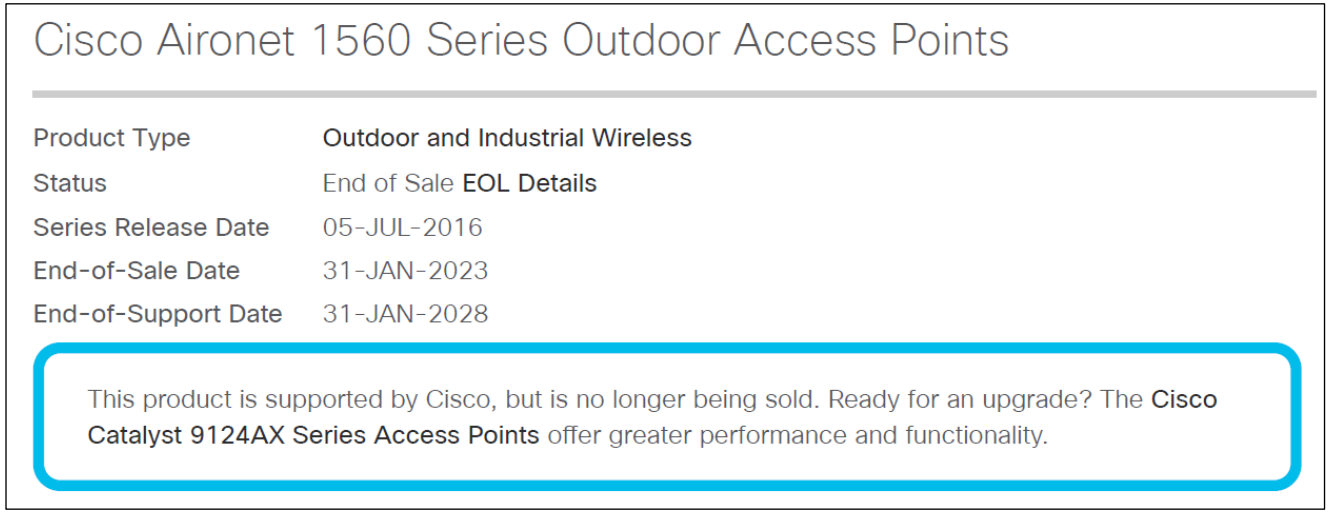
The below chart presents MeshDynamics’s preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts

presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

#	Claim 1	Cisco's Wireless Mesh Network with Catalyst and Aironet Access Points
1.pre	1. A method of controlling a wireless mesh network, comprising:	Cisco's Wireless Mesh Network (a component of the Cisco Unified Wireless Network solution) comprises a plurality of Wi-Fi relay nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Access Points), which employ and provide a method of controlling a wireless mesh network. Cisco Catalyst 9100 and Aironet Access Points support Wi-Fi capabilities. The Catalyst series is a follow-up from the Cisco Aironet series access points, and the Aironet and Catalyst Access Points contain the same mesh networking capabilities, as shown below in Cisco's documentation and guides that apply to both product lines: Q. Which Cisco access points support Wi-Fi 6? A. Cisco Catalyst 9100 enterprise-class access points support Wi-Fi 6 capabilities extending the power of intent-based networking with hardware and software innovations. The Cisco Catalyst 9100 APs offer resiliency and security while enabling intelligence at the network edge. Q. What is the portfolio transition from the Cisco Aironet Series to the Cisco Catalyst 9100 platform? A. The Cisco Catalyst 9100 Access Points consist of the Cisco Catalyst 9105, 9115, 9120, 9130, 9136, 9162, 9164, 9166 and 9166D1, and outdoor 9163E Series. These access points are the follow-ups to the Cisco Aironet 1815, 1850, 2800, and 3800 Series Access Points, respectively. The Cisco Catalyst 9100 Access Points come equipped with Wi-Fi 6/6E capabilities, a better industrial design, and improved RF performance, and deliver reliability, security, and intelligence at scale.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html



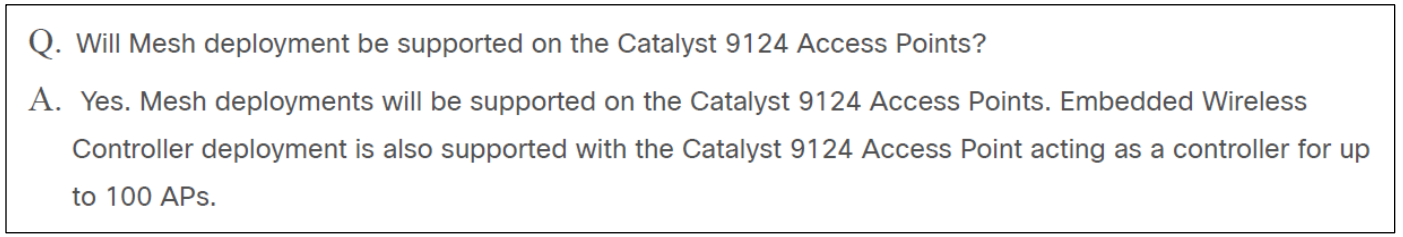
Cisco Aironet 1560 Series Outdoor Access Points

Product Type	Outdoor and Industrial Wireless
Status	End of Sale EOL Details
Series Release Date	05-JUL-2016
End-of-Sale Date	31-JAN-2023
End-of-Support Date	31-JAN-2028

This product is supported by Cisco, but is no longer being sold. Ready for an upgrade? The **Cisco Catalyst 9124AX Series Access Points** offer greater performance and functionality.

Source: https://www.cisco.com/c/en/us/support/wireless/aironet-1560-series/series.html

Mesh deployments are supported on the Cisco Catalyst 9100 Access Points:



Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points?

A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

In addition to the Wave 1 and Wave 2 access points, mesh deployments are supported on the Cisco Catalyst 9100 Access Points, including, *e.g.*, Catalyst 9124 Access Points, Catalyst 9130 Access Points and "AX" access points:

Families of mesh APs

The two families of mesh APs available are the Cisco Wave 1 APs and the Cisco Wave 2 APs + the AX APs. These two families of mesh APs have different capabilities and features, and they are not fully compatible with each other. The Cisco Wave 1 APs do not support the same feature set as the Cisco Wave 2 + AX APs.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg/mesh-access-points.html

From this release, mesh support is included in the Cisco Catalyst 9130AX Series Access Points. All traditional capabilities of mesh are included in the Cisco Catalyst 9130AX Series APs operating in Cisco IOS XE Dublin 17.12.1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html

A plurality of Cisco Catalyst 9100 and Aironet Access Points can be deployed in a Cisco Wireless Mesh Network with a controller and software, such as the Catalyst 9800 Controller and Cisco DNA and Prime Infrastructure:

Mesh Network Components

This section describes the mesh network components.

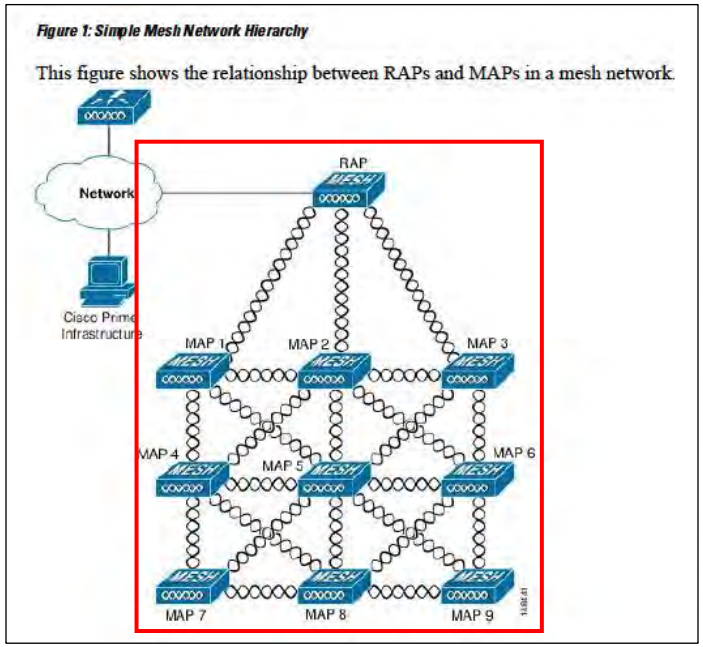
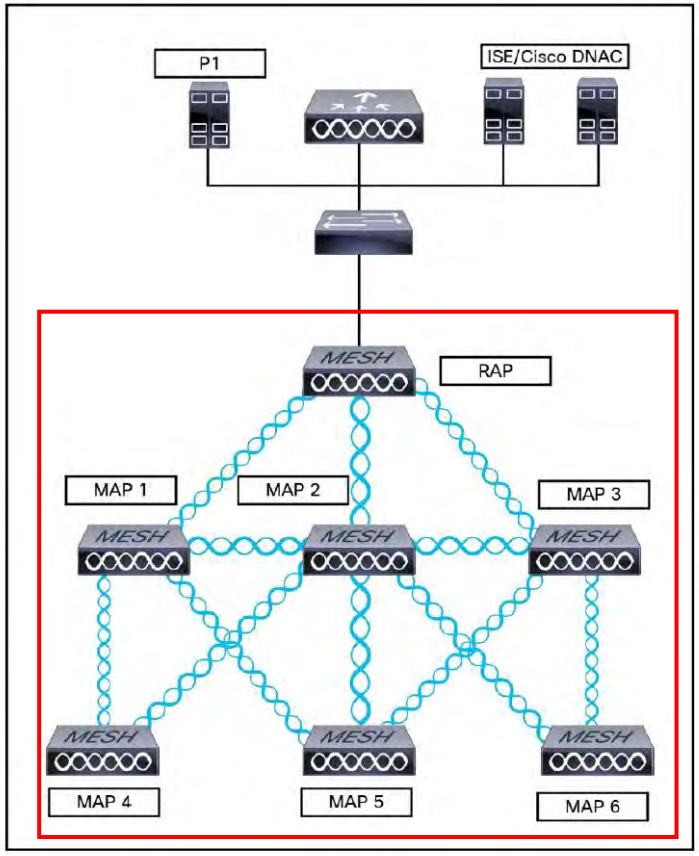
The Cisco wireless mesh network has four core components:

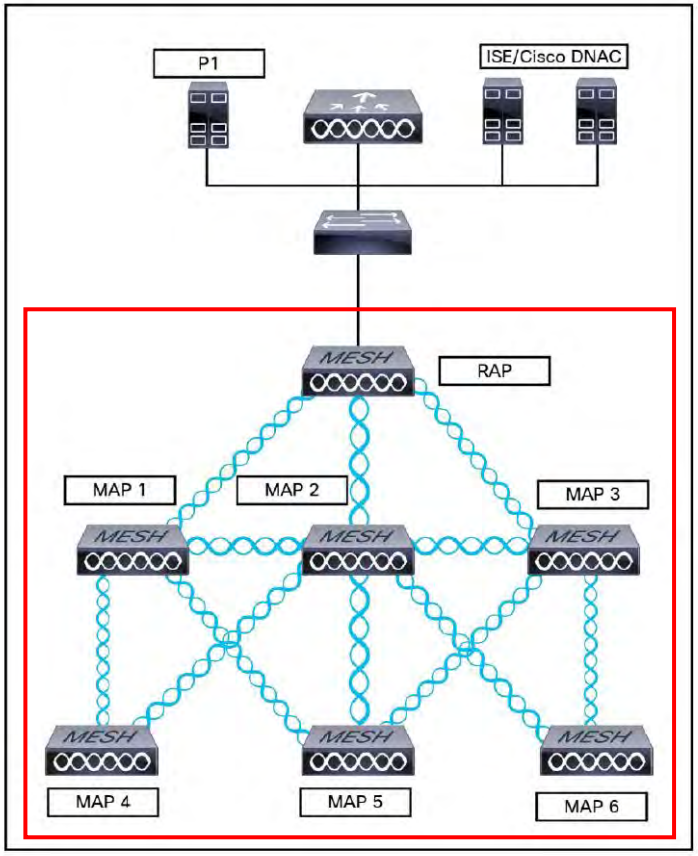
- Cisco C9800 and IOS-XE 17.1
- Cisco Aironet and Catalyst series access points
- Cisco PI and Cisco DNA Center
- Mesh software architecture

Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network.

- Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series.
- Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

		  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
1.a	a) configuring the wireless mesh network with a first configuration, the wireless mesh network comprising a plurality of relay nodes, each of the relay nodes being	Cisco's Wireless Mesh Network (as shown in the diagrams above and in Cisco's documentation) comprises a plurality of relay nodes (*e.g.*, Cisco Catalyst 9100 or Aironet Access Points in the configuration shown above and below in red), which employ and provide a method of controlling a wireless mesh network comprising the step of configuring the wireless mesh network with a first configuration:

	capable of relaying traffic from a child node associated with the relay node, each of the relay nodes including a first radio interface configurable to relay traffic to a parent relay node and a second dual purpose radio interface configurable to both relay traffic from the child relay node and to communicate with a client node that is not a relay node, and at least one of the relay nodes is configured in the first configuration to relay traffic from a child relay node to a parent relay node;	 Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf Cisco Catalyst 9100 and Aironet Access Points within the Cisco Wireless Mesh Network operate as either a Root AP (RAP) or Mesh AP (MAP), and are capable of relaying traffic from a child node (*e.g.*, MAP 5) associated with the parent relay node (*e.g.*, MAP 2) and include a first radio interface configurable to relay traffic to the parent relay node (*e.g.*, MAP 2 or RAP). Cisco Catalyst 9100 and Aironet Access Points perform the step of configuring the Cisco Wireless Mesh Network in a first configuration, as the access points are configured in the “first configuration” shown

above to relay traffic from a child relay node to a parent relay node. All Cisco access points are configured and shipped as MAPs.

For Mesh mode, an AP should be configured with Bridge mode. Access points within a mesh network operate in one of the following two ways:

- Root AP (RAP)
- Mesh AP (MAP)

Note: All APs are configured and shipped as mesh APs. To use an APs as a RAP, you must reconfigure the mesh AP as a RAP. In all mesh networks, ensure that there is at least one RAP.

Network access

Wireless mesh networks can simultaneously carry two different traffic types:

- WLAN client traffic
- MAP Ethernet port traffic

WLAN client traffic terminates on the controller, and the Ethernet traffic terminates on the Ethernet ports of the mesh APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

MAPs communicate traffic among themselves and back to the RAP using wireless connections over the 802.11 a/n/ac/ax radio backhaul ("a first radio interface configurable to relay traffic").

While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

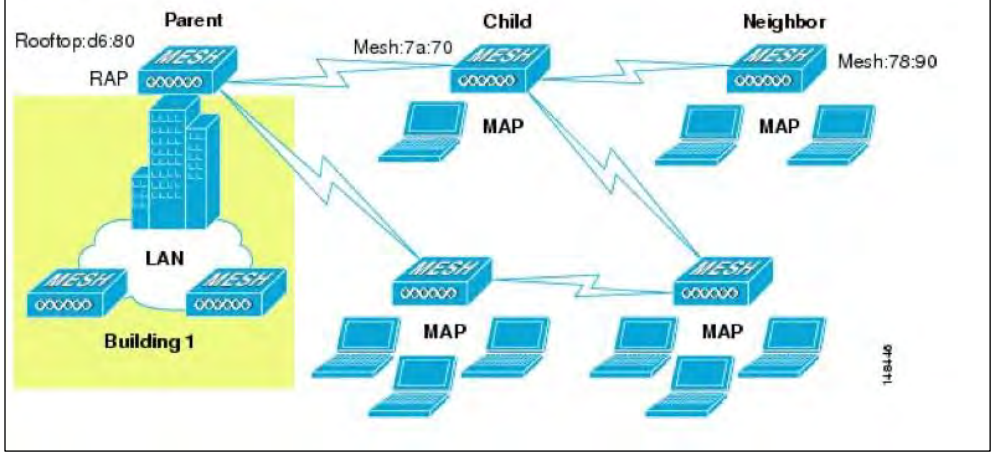
As described and shown in Cisco's wireless mesh networking documentation, relationships among mesh access points are a parent, child or neighbor. A parent access point offers the best route back to the RAP based on ease values and signal strength, and a parent can be either the RAP or another MAP.

Relationships among mesh access points are as a parent, child, or neighbor (see Figure 6: Parent, Child, and Neighbor Access Points).

- A parent access point offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
    - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an access point with a higher ease value is selected.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

Figure 6: Parent, Child, and Neighbor Access Points



Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's Adaptive Wireless Path Protocol (AWPP) is designed specifically for mesh networking and enables each relay node (*e.g.*, Cisco Catalyst 9100 or Aironet Access Point) to dynamically find the best path back to a RAP and relay traffic to a parent node.

AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Cisco Catalyst 9100 and Aironet Access Points also include a second dual purpose radio interface (*e.g.*, 2.4 GHz or 5-GHz radio) configurable to both relay traffic from the child relay node (*e.g.*, MAP 5) and to communicate with a client node (*e.g.*, laptop) that is not a relay node:

Frequency bands

Both the 2.4-GHz and 5-GHz frequency bands are supported on the indoor and outdoor access points.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's Catalyst 9100 and Aironet Access Points support multiple radios, including one or more radios for client access and one or more radios for wireless backhaul and flexible radio assignment. For instance, Cisco Catalyst 9100 and Aironet Access Points can operate in tri-radio mode (known as split 5 GHz mode) or XOR mode (the ability to operate a single radio interface in many modes), which allows one 5 GHz radio to be split into two separate 5 GHz radios. The access point carries out the radio reassignment, and this enables the access point to utilize three radios simultaneously: one 2.4 GHz radio and two 5 GHz radios. The radios can serve different user groups on different channels.

Table 3. AP models and types of hardware managed by FRA

AP model	FRA radios	Functions
Cisco Aironet 2800 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Aironet 3800 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Aironet 4800 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Catalyst 9120 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Catalyst 9130 Series Access Points	5-GHz Tri-Radio	2.4-GHz 4x4 and single 5-GHz 8x8, or 2.4-GHz 4x4 and dual 5-GHz 4x4
Cisco Catalyst Wireless 9166 Access Points	5/6-GHz XOR	2.4-GHz 4x4 and dual 5-GHz 4x4, or 5-GHz 4x4 and 6-GHz 4x4

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html

For example, Cisco Catalyst 9130AX Wi-Fi 6 access point is tri-radio (2.4 GHz + 5 GHz with dual-5GHz using Flexible Radio Assignment) that supports mesh. One radio can be used for client access, the others used flexibly for backhaul (uplink and downlink) at different channels and client access.

Flexible Radio Assignment (FRA) with tri-radio mode	FRA allows the access points to intelligently determine the operating mode of serving radios based on the RF environment and traffic demands. The access points can operate in the following modes: • Dual-radio mode: One 8x8 5 GHz and one 4x4 2.4 GHz. One radio will serve clients on the 5-GHz band, while the other serves clients on the 2.4-GHz band. • Tri-radio mode: * Dual 4x4 5 GHz and one 4x4 2.4 GHz. With two 4x4 5-GHz and one 4x4 2.4-GHz radios (tri-radio) inside the access point, client device capacity can be increased on demand. The access point's default mode is dual radio with 8x8 5 GHz and 4x4 2.4 GHz. It has the ability to split the 8x8 radio into two separate 4x4 5-GHz radios through software, thereby enabling the benefits of FRA while allowing the 2.4-GHz radio to remain active.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html

Cisco Access Points with Tri-Radio

This topic describes the Tri-Radio feature for Cisco Access Points (APs).

Access Points with three radios are designed for high density environments. The APs by default run one dedicated 2.4-GHz 4x4 mode radio and one 5-GHz 8x8 mode radio. In the default mode, the radios are managed by the Flexible Radio Assignment (FRA), and the Dual Radio Mode is in the disabled state indicating that the radios have either been assigned as client serving 8x8 radio or have not yet been evaluated by FRA.

When you enable the dual radio mode setting, the 8x8 radio is split to two independent 5-GHz 4x4 radios. In this mode, slot 1 and slot 2 are active independent 4x4 radio interfaces. They can serve different user groups with different assigned channels.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/cisco_access_points_with_tri_radio.pdf

In addition, while 5 GHz is often used for the primary backhaul (communication between access points), both the 2.4 GHz and 5GHz radios are configurable to be used for backhaul and client access (*e.g.*, "second dual purpose radio interface"), for example, when Universal Backhaul Access (UBA) is enabled.

.The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios:

- From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable
- 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled

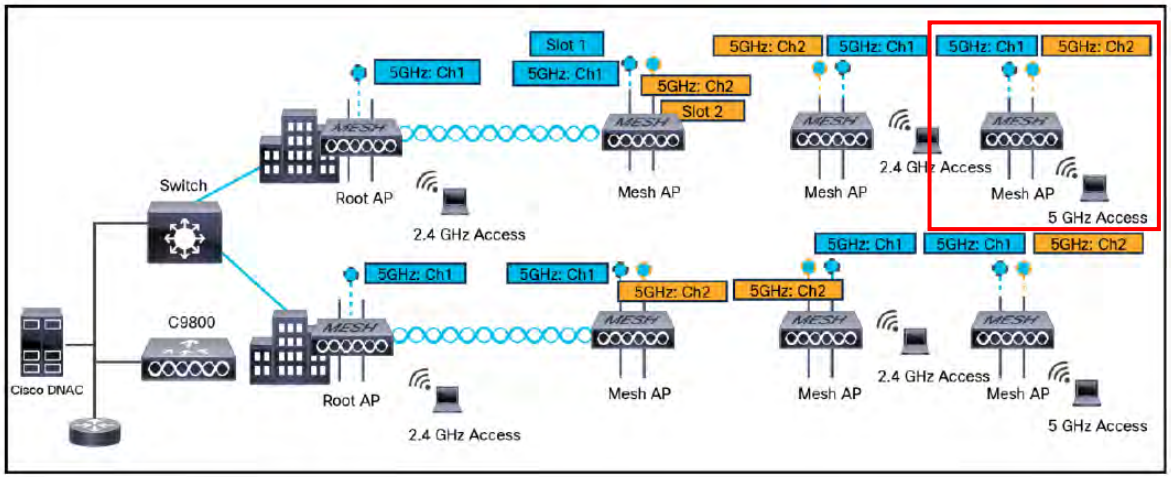
The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-6/b_mesh_86/m_mesh-network-components.html; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/technotes/8-5/cisco-enterprise-mobility-design-guide-8-5.pdf

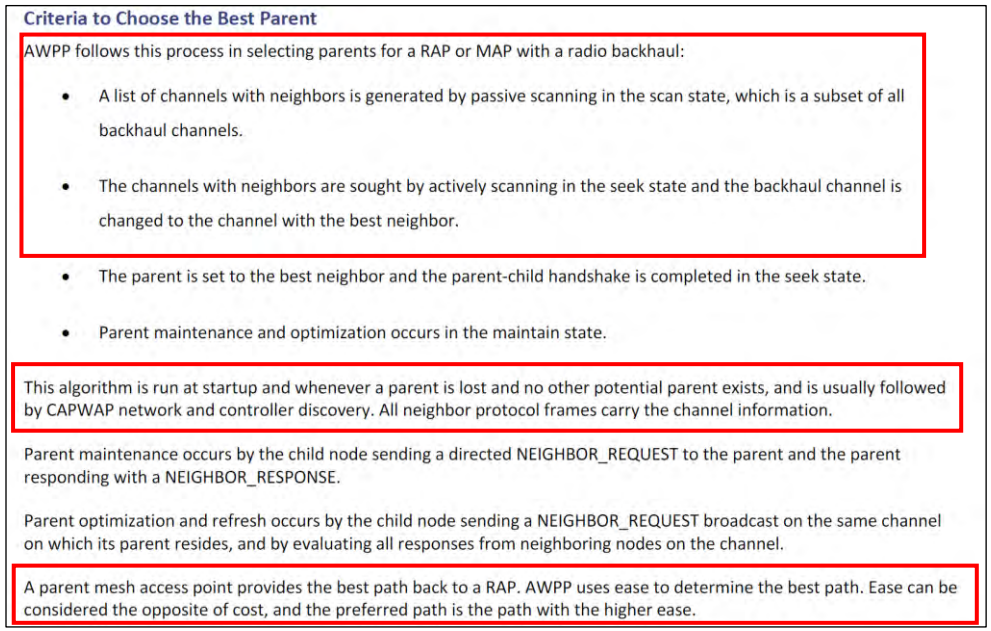
With Mesh Serial Backhaul, two of the radios are reserved for backhaul. The Mesh Serial Backhaul feature in a mesh access point (MAP) allows different channels for uplink (*e.g.*, "first radio interface") and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. Basic client access functionality is offered on the 2.4-GHz radio and the 5-GHz radio, which are not used in serial backhaul. Universal access is made available on the downlink radio (*e.g.*, "second dual purpose radio interface").

Mesh serial backhaul

Serial backhaul enables us to use different channels on different backhauls, which helps in reducing interference and maximizing throughput. The intention is that in units with third and fourth radios, two of the radios are reserved for backhaul.



		Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-7/config-guide/b_wl_17_7_cg/m_mesh_ewlc.html#mesh-backhaul-24ghz-5ghz If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of establishing a first wireless mesh configuration of multiple relay nodes (each with dual radio interfaces for relaying traffic and serving client nodes) in the same way by using AWPP to deploy multiple nodes, each equipped with a backhaul radio and a dual-purpose radio (or tri-radio mode) and configuring at least one node (*e.g.*, an access point) to relay a child node's traffic to a designated parent node (*e.g.*, another access point), and achieves the same result by creating a wireless mesh network in which a relay node forwards traffic from its child node to its parent node, while the child serves clients. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's AWPP is used to deploy multiple dual-radio (or tri-radio) mesh access points to relay traffic via backhaul links, achieving the same multi-node mesh configuration of relay nodes serving clients. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.b	b) determining that a first one of the plurality of relay nodes in the wireless mesh network requires a channel selection change;	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network employ and provide a method of controlling a wireless mesh network, comprising the step of determining that that a first one of the plurality of relay nodes in the wireless mesh network requires a channel selection change. Cisco Catalyst 9100 and Aironet Access Points ("plurality of relay nodes in the wireless mesh network") can determine (*e.g.*, for parent optimization and upon the loss of a parent connection) that it should change channels and connect with a better potential parent across any channel to maintain its uplink with the best parent (*i.e.*, the best path back to the RAP based on signal strength and number of hops/ease criteria). This required "channel selection change" will choose the best parent.



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

On information and belief, when the backhaul "channel is changed to the channel with the best neighbor" (as shown in the preceding image), AWPP may work in conjunction with, or coordinate with, the "advance radio management features in the CAPWAP WLAN solution." This may include working with Radio Resource Management ("RRM") and/or Dynamic Channel Allocation ("DCA"). For instance, mesh path changes may be triggered by DCA-driven channel changes. AWPP may react by recalculating mesh paths, possibly selecting a new parent or RAP if the link quality changes.

> AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

Dynamic Channel Assignment

Two adjacent access points on the same channel can cause either signal contention or signal collision. In a collision, data is not received by the access point. This functionality can become a problem, for example, when someone reading an e-mail in a café affects the performance of the access point in a neighboring business. Even though these are separate networks, someone sending traffic to the café on channel 1 can disrupt communication in an enterprise using the same channel. Controllers can dynamically allocate access point channel assignments to avoid conflict and increase capacity and performance. Channels are *reused* to avoid wasting scarce RF resources. In other words, channel 1 is allocated to a different access point far from the café, which is more effective than not using channel 1 altogether.

The controller's Dynamic Channel Assignment (DCA) capabilities are also useful in minimizing adjacent channel interference between access points. For example, two overlapping channels in the 802.11b/g band, such as 1 and 2, cannot simultaneously use 11 or 54 Mbps. By effectively reassigning channels, the controller keeps adjacent channels that are separated.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-9/config-guide/b_cg89/radio_resource_management.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining when a first relay node (*e.g.*, a access point) in the mesh requires a channel selection change (*i.e.*, detecting that the node's current channel must be changed) in the same way by employing a dynamic channel selection mechanism that monitors wireless conditions and identifies when the relay node's channel needs to be switched, and achieves the same result by recognizing that a channel selection change is necessary to maintain proper network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's software continuously monitors wireless conditions and identifies when an access point's current channel must be changed, accomplishing the same function of determining that a channel selection change is required. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1.c	c) dynamically selecting a second one of the plurality of relay nodes in the wireless mesh network to form a parent-child relationship with the first relay node; and	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network employ and provide a method of controlling a wireless mesh network, comprising the step of dynamically selecting a second one of the plurality of relay nodes in the wireless mesh network to form a parent-child relationship with the first relay node. Cisco's Adaptive Wireless Path Protocol (AWPP) is designed specifically for mesh networking and enables each relay node (*e.g.*, Cisco Catalyst 9100 and Aironet Access Points) to dynamically select the best parent to form a parent-child relationship and the best path to the controller. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

With Mesh Background Scanning and Auto parent selection, a MAP can better find and dynamically connect with a potential parent across any channels and maintain its uplink with the best parent all the time. After the detection of parent loss, a MAP can directly select a best parent from the neighbor list and establish the AWPP link.

Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

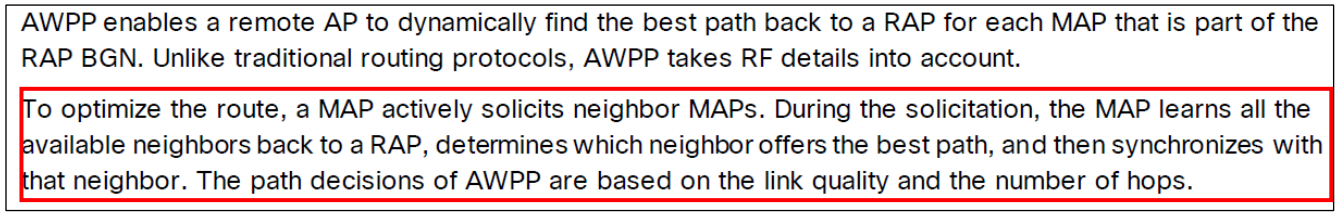
Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Mesh background scanning

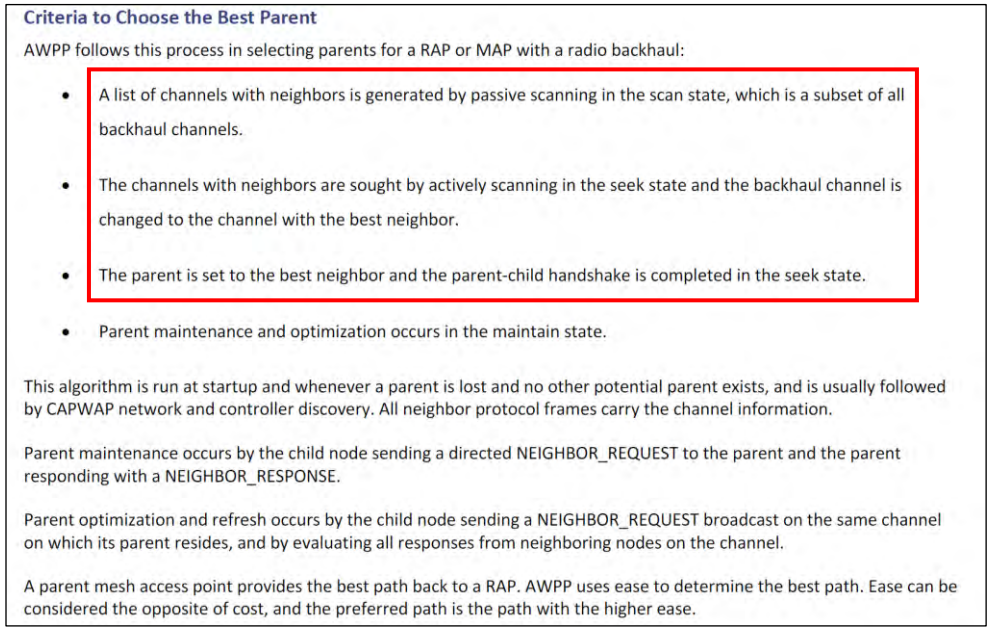
Cisco MAPs are interconnected over wireless links in a spanning tree-like topology. A MAP connected to the network via Ethernet uplink is designated as the root AP. AWPP is used to maintain and form the tree. When a MAP comes up, it tries to look for another MAP (parent) to join to reach the gateway eventually via a RAP. The same happens when a MAP lose connectivity with its existing parent. This procedure is known as mesh tree convergence. This feature aims to make the convergence procedure faster and more robust.

This procedure is time consuming. To simplify the process, we have introduced mesh background scanning and automatic parent selection. This feature is available in Cisco IOS XE starting with Release 17.11.1 in the Aironet 1562 and Catalyst 9124AX APs. This mechanism allows the MAP to find and connect faster to a better potential parent across channels and always maintain its uplink with the best parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; *see also* https://www.cisco.com/c/en/us/support/docs/wireless/embedded-wireless-controller-on-catalyst-9124ax-access-points/221557-configure-point-to-point-mesh-link-with.pdf

		If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of dynamically selecting a second relay node (*e.g.*, an access point) to form a parent-child relationship with the first relay node (*e.g.*, an access point) in the same way by using AWPP or background scanning to evaluate available neighboring nodes and choose the best candidate as the new parent for the first node, and achieves the same result by establishing a parent-child link between the first node and the selected second node. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's AWPP protocol dynamically selects the optimal neighboring access point to serve as a new parent, thereby forming the same parent-child relationship between the first node and the selected second node as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.d	d) making a channel selection change wherein prior to being dynamically selected, the second relay node was a sibling node of the first relay node; and making a channel selection includes making a channel selection change to meet a performance requirement.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network employ and provide a method of controlling a wireless mesh network, comprising the step of making a channel selection change wherein prior to being dynamically selected the second relay node (*e.g.*, a MAP) was a sibling node of the first relay node (*e.g.*, a neighboring MAP); and making a channel selection includes making a channel selection change to meet a performance requirement (*e.g.*, reducing latency in backhaul data traffic). AWPP follows a dynamic process for selecting a parent to make any necessary channel selection change, including making a channel selection change to meet a performance requirement (*i.e.*, by choosing a parent based on the best path based on signal strength and number of hops). The AWPP sets the parent (*e.g.*, changes or selects the backhaul channel) to the best neighbor (*e.g.*, "sibling") and completes the parent-child handshake.  AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's documentation demonstrates that choosing a parent involves making a channel selection change wherein prior to being dynamically selected, the second relay node (*e.g.*, a neighboring MAP) was a sibling node of the first relay node. As illustrated below, AWPP scans available backhaul channels, changes the backhaul channel to that of the best neighbor, and then sets that neighbor as the new parent.



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

On information and belief, when the backhaul "channel is changed to the channel with the best neighbor" (as shown in the preceding image), AWPP may work in conjunction with, or coordinate with, the "advance radio management features in the CAPWAP WLAN solution." This may include working with Radio Resource Management ("RRM") and/or Dynamic Channel Allocation ("DCA"). For instance, mesh path changes may be triggered by DCA-driven channel changes. AWPP may react by recalculating mesh paths, possibly selecting a new parent or RAP if the link quality changes.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

Dynamic Channel Assignment

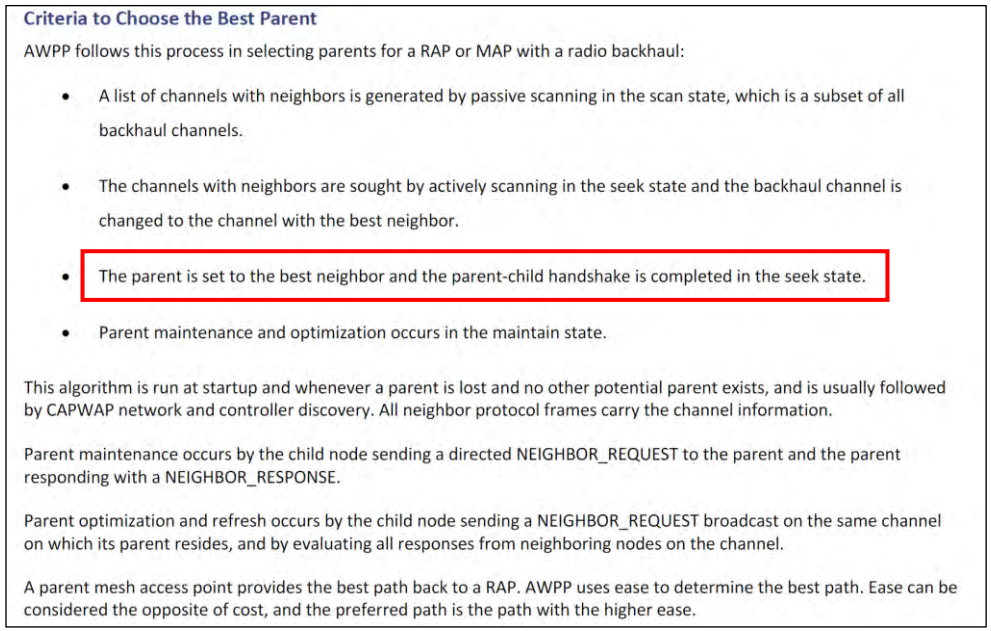
Two adjacent access points on the same channel can cause either signal contention or signal collision. In a collision, data is not received by the access point. This functionality can become a problem, for example, when someone reading an e-mail in a café affects the performance of the access point in a neighboring business. Even though these are separate networks, someone sending traffic to the café on channel 1 can disrupt communication in an enterprise using the same channel. Controllers can dynamically allocate access point channel assignments to avoid conflict and increase capacity and performance. Channels are *reused* to avoid wasting scarce RF resources. In other words, channel 1 is allocated to a different access point far from the café, which is more effective than not using channel 1 altogether.

The controller's Dynamic Channel Assignment (DCA) capabilities are also useful in minimizing adjacent channel interference between access points. For example, two overlapping channels in the 802.11b/g band, such as 1 and 2, cannot simultaneously use 11 or 54 Mbps. By effectively reassigning channels, the controller keeps adjacent channels that are separated.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-9/config-guide/b_cg89/radio_resource_management.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of making a channel selection change that promotes a neighboring node to become the new parent of the first node in order to meet a performance requirement in the same way by dynamically reassigning the first node's backhaul connection to an optimal neighboring relay node based on metrics such as signal quality, latency, throughput, etc. and achieves the same result by improving the mesh network's performance (for example, reducing latency or hop count) through the channel selection change while establishing the second node as the first node's parent. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's AWPP protocol automatically switches a first node's backhaul connection to the "best" neighboring

		node (previously a sibling node) as a new parent whenever needed to meet network performance requirements (such as better signal or fewer hops), thereby making the same performance-driven channel selection change to a new parent. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 2	
2	2. A method as recited in claim 1, wherein the second node is selected to be a parent node of the first node.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the second node is selected to be a parent node of the first node. Cisco Catalyst 9100 and Aironet Access Points in a mesh network implement this limitation by dynamically selecting a neighboring node to serve as a new parent for the first node. The Adaptive Wireless Path Protocol (AWPP) running on each mesh AP will set the parent to the "best" neighboring AP and establish a parent-child link, thereby selecting the second node as the parent node of the first node. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

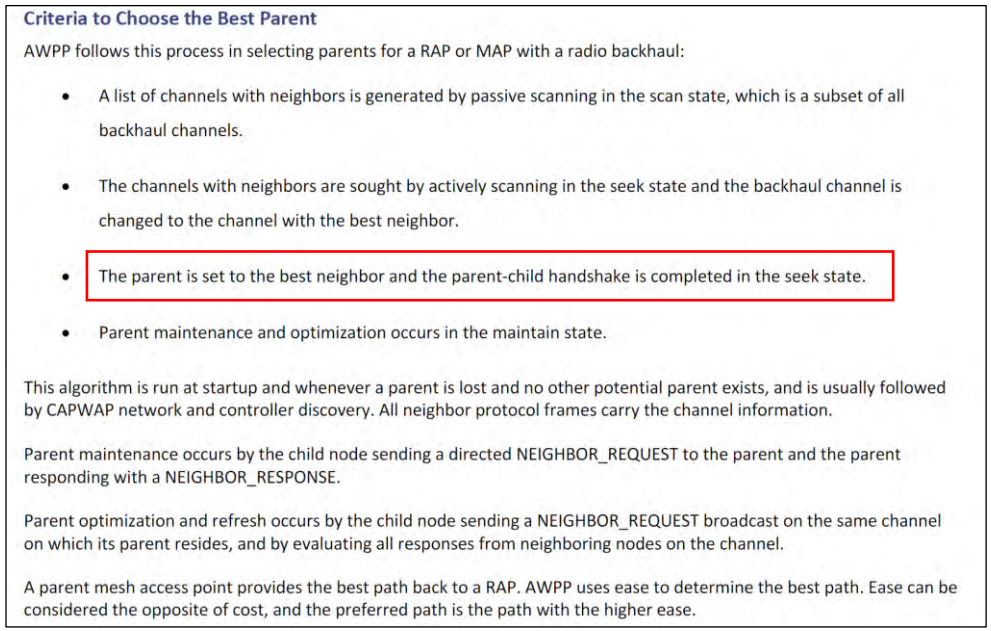
A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of selecting the second node to serve as the parent node of the first node in the same way by utilizing AWPP to dynamically designate the best neighboring access point as the new parent for the first access point, and achieves the same result by establishing the second node as the first node's parent within the wireless mesh topology. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh nodes dynamically choose a neighboring access point to act as the new parent node for the first node (via AWPP's parent selection algorithm), fulfilling the same function of selecting the second node as a parent for the first node. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

	Claim 3	
3	3. A method as recited in claim 1, wherein the second node is selected to be a child node of the first node.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the second node is selected to be a child node of the first node. Cisco Catalyst 9100 and Aironet Access Points in a mesh network implement this limitation by dynamically selecting the first node to serve as a new parent for the second node. The Adaptive Wireless Path Protocol (AWPP) running on each mesh AP will set the parent to the "best" neighboring AP and establish a parent-child link, thereby selecting the second node as the child node of the first node. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

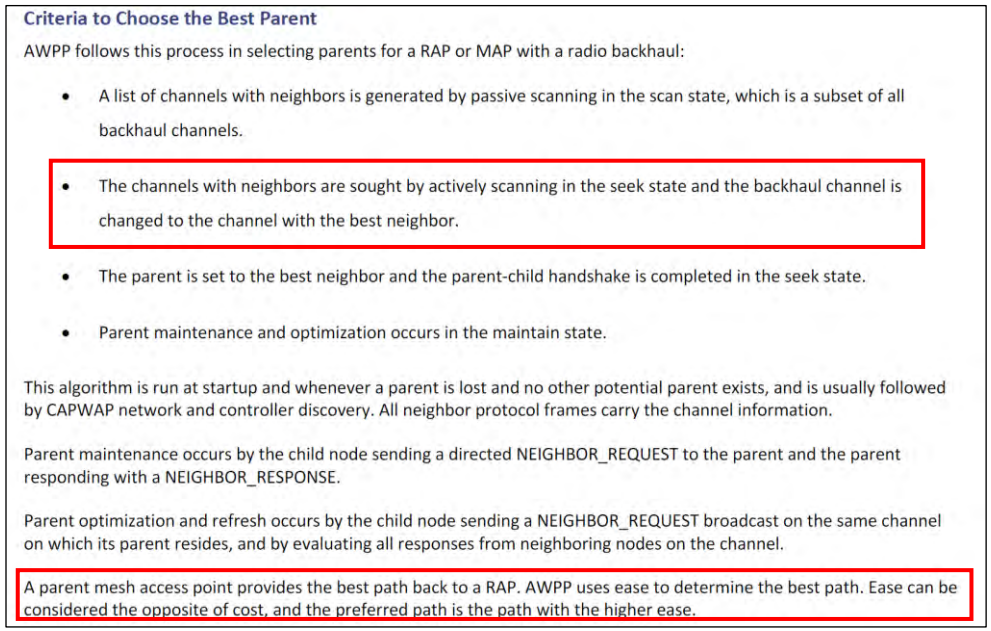
Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of selecting the second node (*e.g.*, an access point) to be a child node of the first node (*e.g.*, another access point) in the same way by dynamically reconfiguring the mesh so that the first access point becomes the parent and the second access point is assigned as its child by using AWPP, and achieves the same result by establishing a new parent-child relationship in which the second node is connected as a child under the first node. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's algorithm can equivalently reconfigure node roles such that the first access point becomes a parent and the second access point is assigned as its child (using AWPP's dynamic reconfiguration), achieving the same result of establishing a new parent-child relationship with the second node as the child under the first node. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

		MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 4	
4	4. A method as recited in claim 1, wherein making a channel selection includes making a channel selection change to meet a latency requirement.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein making a channel selection includes making a channel selection change to meet a latency requirement, including Cisco's "ease calculation" functionality, which considers the latency related to the "hop count" or "hop value" back to the RAP. **Ease Calculation** Ease is calculated using the SNR and hop value of each neighbor, and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf **Parent Decision** A parent mesh access point is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: adjusted ease = min (ease at each hop) Hop count Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf Cisco Catalyst 9100 and Aironet Access Points perform channel selection changes to meet specific performance requirements such as latency. For example, when using AWPP, Cisco's Catalyst 9100 and Aironet Access may switch the backhaul channel (following the "parent decision" shown above) in order to reduce latency or hop value thereby making a channel selection change to satisfy a latency requirement as recited.



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

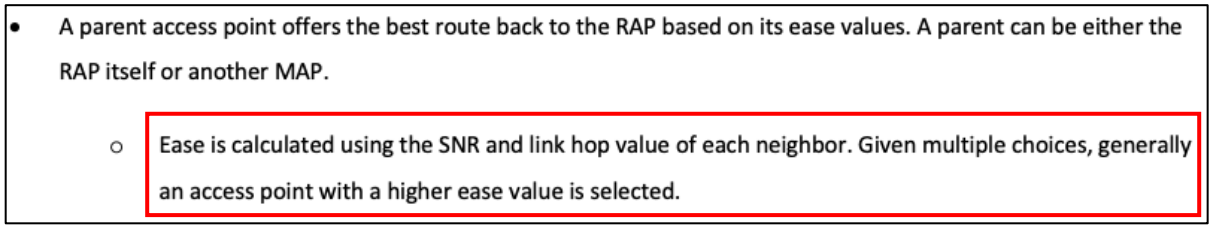
Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

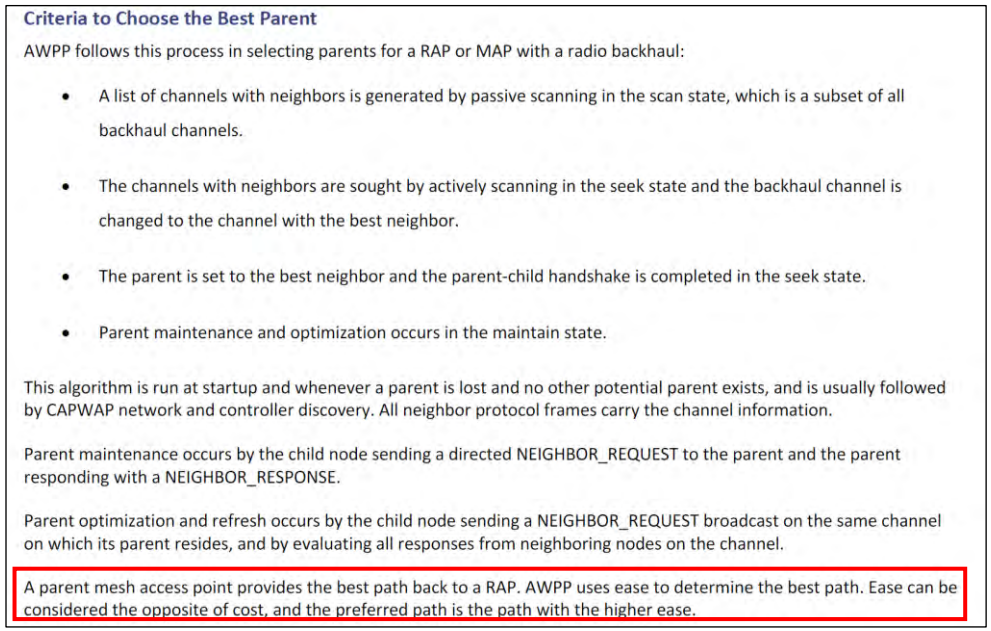
Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of making a channel selection change to meet a latency requirement in the same way by automatically switching the first node's wireless backhaul to a better channel or alternate parent node whenever network latency exceeds a defined threshold and achieves the same result by reducing communication latency so that the mesh network maintains the required performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco access points automatically trigger a channel selection change (or parent switch) whenever network latency rises above a defined threshold – for example, the mesh software will select and switch the first node's backhaul channel or parent to a better option if latency becomes too high – thereby meeting the same latency requirement as the claimed method. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

		MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 5	
5	5. A method as recited in claim 1, wherein making a channel selection includes making a channel selection change to meet a throughput requirement.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein making a channel selection includes making a channel selection change to meet a throughput requirement. Cisco Catalyst 9100 and Aironet Access Points perform channel selection changes to meet specific performance requirements such as throughput based on link hop values and/or signal strength. For example, when using AWPP the access point may switch the backhaul channel in order to meet throughput requirements associated with high link hop values and lower signal strength in backhaul data traffic, thereby making a channel selection change to satisfy a throughput requirement as recited.  • A parent access point offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP. ○ Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an access point with a higher ease value is selected. Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf.



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

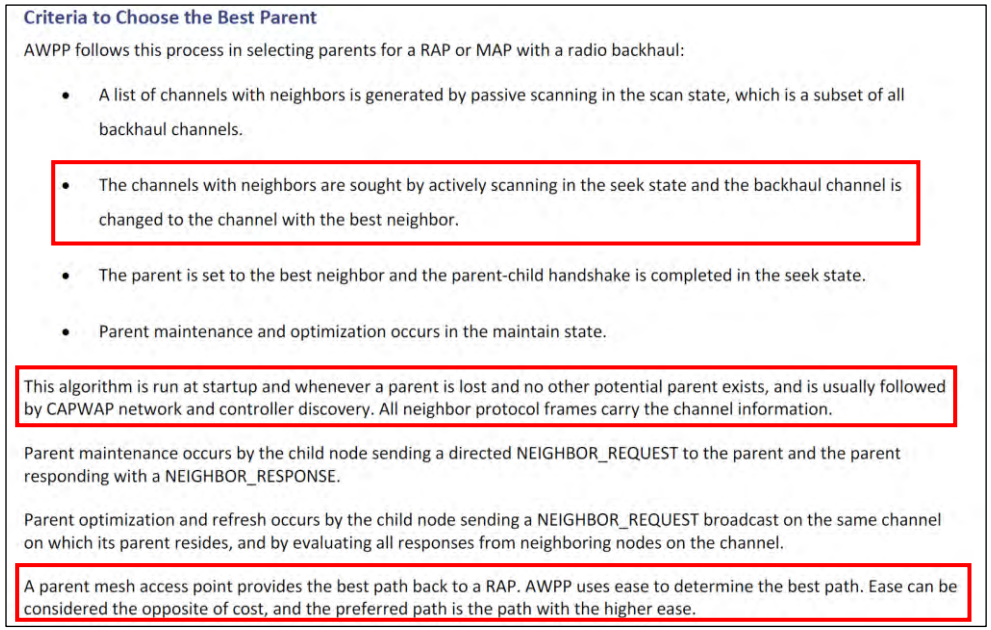
Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of making a channel selection change to meet a throughput requirement in the same way by dynamically adjusting the channel or route based on throughput-related metrics (for example, detecting excessive link hops or low signal strength and switching to a channel that provides higher capacity), and achieves the same result by increasing the data throughput so that the mesh network satisfies the required throughput level. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's channel selection algorithms dynamically adjust the wireless backhaul (by changing the channel and routing path) based on throughput-related metrics – for example, detecting excessive link hops or load or low signal strength and switching to a higher-capacity channel – thereby making the same kind of channel selection change to satisfy a throughput requirement. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

		MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 6	
6	6. A method as recited in claim 1, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to a performance requirement change.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, further employ and provide a method of controlling a wireless mesh network, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to a performance requirement change (*e.g.*, a performance requirement change associated with an updated link hop value). As shown above for claims 4 and 5 (incorporated herein by reference), Cisco Catalyst 9100 and Aironet Access Points perform channel selection changes to meet specific performance requirements such as throughput and/or latency. For example, when using AWPP, the access points may switch the backhaul channel in order to meet performance requirements associated with high link hop values in backhaul data traffic, thereby making a channel selection change to satisfy a performance requirement change (*e.g.*, a high hop value) as recited.  **Criteria to Choose the Best Parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: • A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels. • The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor. • The parent is set to the best neighbor and the parent-child handshake is completed in the seek state. • Parent maintenance and optimization occurs in the maintain state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease. Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-

network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

As another example, when using AWPP, Cisco Catalyst 9100 and Aironet Access Points may switch the backhaul channel in order to reduce latency in backhaul data traffic, thereby making a channel selection change due to a performance requirement change as recited.

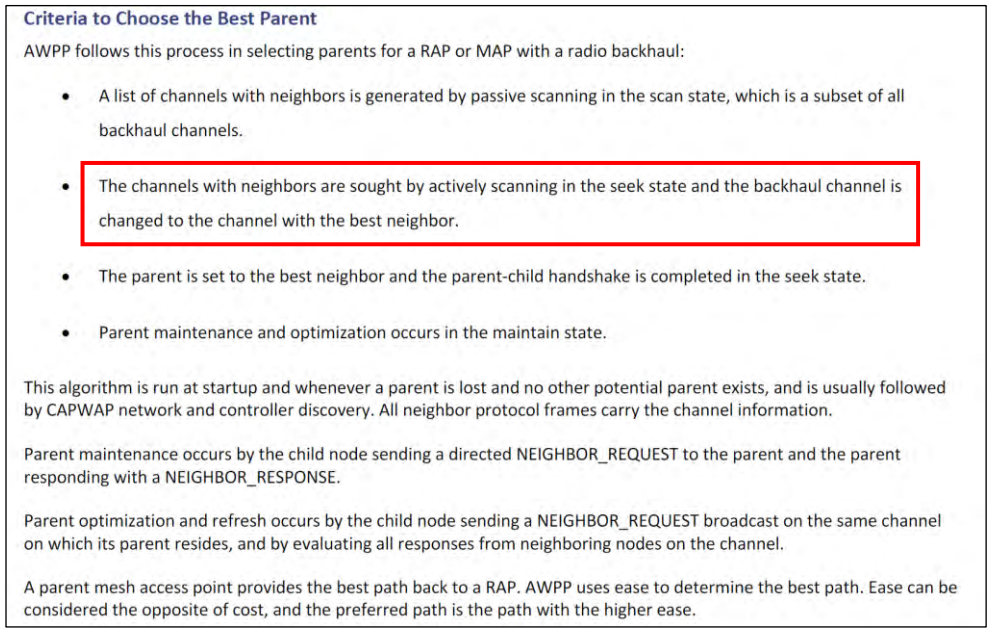
- A parent access point offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
    - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an access point with a higher ease value is selected.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining that a channel selection change is required due to a change in a performance requirement in the same way by monitoring network performance and identifying when an updated performance requirement indicates the current channel is inadequate, and achieves the same result by detecting that the first node must switch channels in response to the changed performance requirement. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products monitor network performance criteria and can tell when an updated performance requirement makes the current channel inadequate, thus detecting the need for a channel change due to the changed performance requirement in essentially the same way as the claimed invention. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

	Claim 7	
7	7. A method as recited in claim 1, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to congestion.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to congestion. Cisco Catalyst 9100 and Aironet Access Points channel selection process accounts for congestion and will change channels to alleviate noise and interference problems associated with congestion when needed. Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network determine that the first node requires a channel selection change by determining that the channel selection change is required due to congestion by considering “SNR” and/or “load values” based on their parent selection algorithm. This feature tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink. This 'switching' parent procedure need not be a triggered event like parent loss detection, but also on identifying a better parent using 'Auto parent selection algorithm' when the child MAP still has its current parent uplink active. The " Auto Parent selection algorithm" is based on the new "ease" values. For better convergence calculation in rel 8.3 a new "Ease" value introduced for smoother and faster parent or neighbor finding and auto parent connection algorithm. The Ease value is based on SNR, number of Hops, Timers and load values. For the off-channel neighbors the AdjustedEase value will be used and a best neighbor per off-channel shall be identified based on its highest AdjustedEase value. StickyEase shall be applicable only for on-channel parent. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf AWPP’s and the AP’s focus on factors like link quality/signal strength, reducing the number of hops and load value helps to minimize and avoid congestion in the mesh network. For example, lower throughput is an indication of congestion. By choosing the “best” path to the wired network based on these criteria, and switching parents, AWPP’s addresses the level of network congestion experienced by an available parent node. Following the parent selection algorithm shown in the previous image, Cisco Catalyst 9100 and Aironet Access Points determine that an AP requires a channel selection as shown in the image below:



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

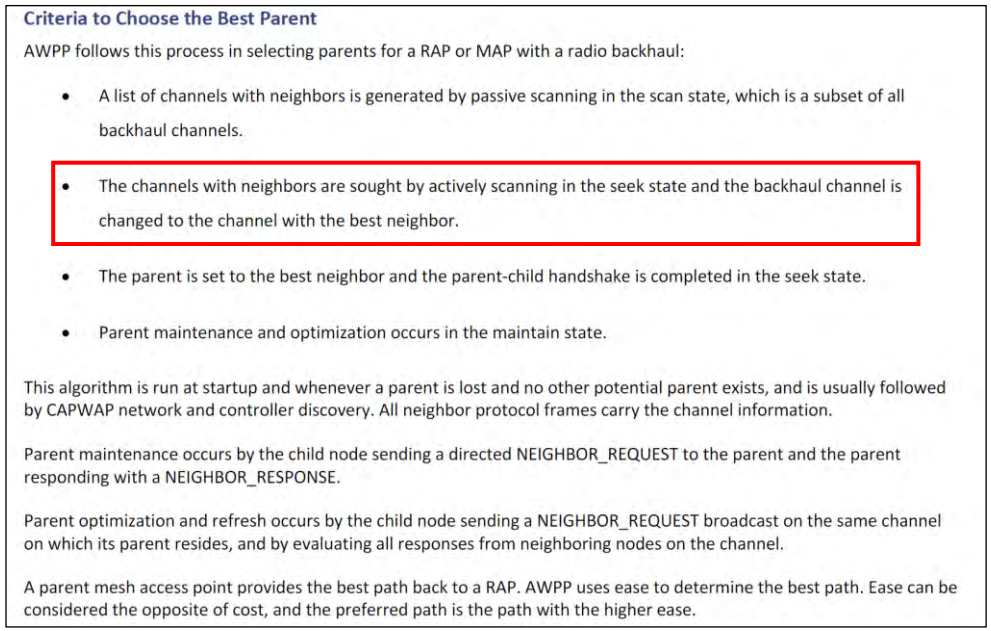
A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

On information and belief, when the backhaul "channel is changed to the channel with the best neighbor" (as shown in the preceding image), AWPP may work in conjunction with, or coordinate with, the "advance radio management features in the CAPWAP WLAN solution." This may include working with Radio Resource Management ("RRM") and/or Dynamic Channel Allocation ("DCA"). For instance, mesh path changes may be triggered by DCA-driven channel changes due to channel interference or congestion. AWPP may react by recalculating mesh paths, possibly selecting a new parent or RAP if the link quality changes.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining that a channel selection change is required due to congestion in the same way by analyzing channel utilization (via Cisco's parent selection algorithms that consider metrics like traffic load) to decide when the first node's channel is overly congested and must be changed, and achieves the same result by triggering a channel switch whenever the current channel becomes too congested, thereby alleviating congestion in the mesh network. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial

		because the Cisco Accused Products' parent selection algorithm takes congestion into account by measuring load values, and it automatically switches the first node's backhaul to a less-loaded channel when the current channel becomes overly congested, thereby performing the same congestion-triggered channel change as in the claim. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 8	
8	8. A method as recited in claim 1, wherein the channel selection change is made to avoid channel interference.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the channel selection change is made to avoid channel interference. Cisco Catalyst 9100 and Aironet Access Points channel selection process accounts for channel interference and will change channels to alleviate channel interference problems when needed. Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network determine that the first node requires a channel selection change by determining that the channel selection change is required due to interference by considering "SNR" and/or "load values" based on their parent selection algorithm. This feature tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink. This 'switching' parent procedure need not be a triggered event like parent loss detection, but also on identifying a better parent using 'Auto parent selection algorithm' when the child MAP still has its current parent uplink active. The " Auto Parent selection algorithm" is based on the new "ease" values. For better convergence calculation in rel 8.3 a new "Ease" value introduced for smoother and faster parent or neighbor finding and auto parent connection algorithm. The Ease value is based on SNR, number of Hops, Timers and load values. For the off-channel neighbors the AdjustedEase value will be used and a best neighbor per off-channel shall be identified based on its highest AdjustedEase value. StickyEase shall be applicable only for on-channel parent. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf Following the parent selection algorithm shown in the previous image, Cisco Catalyst 9100 and Aironet Access Points determine that an AP requires a channel selection as shown in the image below:



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

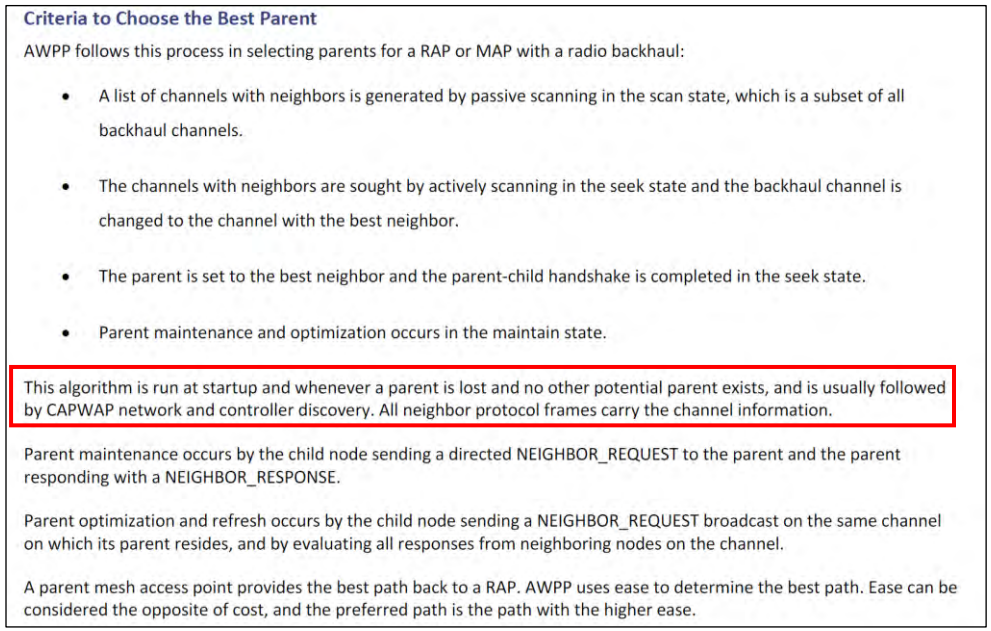
A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

On information and belief, when the backhaul "channel is changed to the channel with the best neighbor" (as shown in the preceding image), AWPP may work in conjunction with, or coordinate with, the "advance radio management features in the CAPWAP WLAN solution." This may include working with Radio Resource Management ("RRM") and/or Dynamic Channel Allocation ("DCA"). For instance, mesh path changes may be triggered by DCA-driven channel changes due to channel interference or congestion. AWPP may react by recalculating mesh paths, possibly selecting a new parent or RAP if the link quality changes.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of making a channel selection change to avoid channel interference in the same way by continually evaluating the interference on the current channel (using Cisco's AWPP algorithm or DCA to monitor interference) and switching the first node to a cleaner channel when interference is excessive, and achieves the same result by maintaining communications on a channel with substantially reduced interference. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products

		continually evaluate interference on the current channel and switches the first node to a cleaner channel whenever the existing channel suffers excessive interference, achieving the same result of avoiding channel interference as the claimed method. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 9	
9	9. A method as recited in claim 1, wherein the channel selection change is made based on a change to network topology.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the channel selection change is made based on a change to network topology. Cisco Catalyst 9100 and Aironet Access Points implement this step by dynamically selecting a neighboring node to serve as a new parent for the first node. The Adaptive Wireless Path Protocol (AWPP) running on each node will set the parent to the "best" neighboring AP and establish a parent-child link, thereby selecting the second node as the parent node of the first node. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If a MAP's current parent "is lost" (*e.g.*, a topology change), the AP detects the loss and scans other channels to find and connect to a new parent on a different channel. In this way, a channel selection change is performed when the network topology changes. A new and better neighbor is another example of a topology change.



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

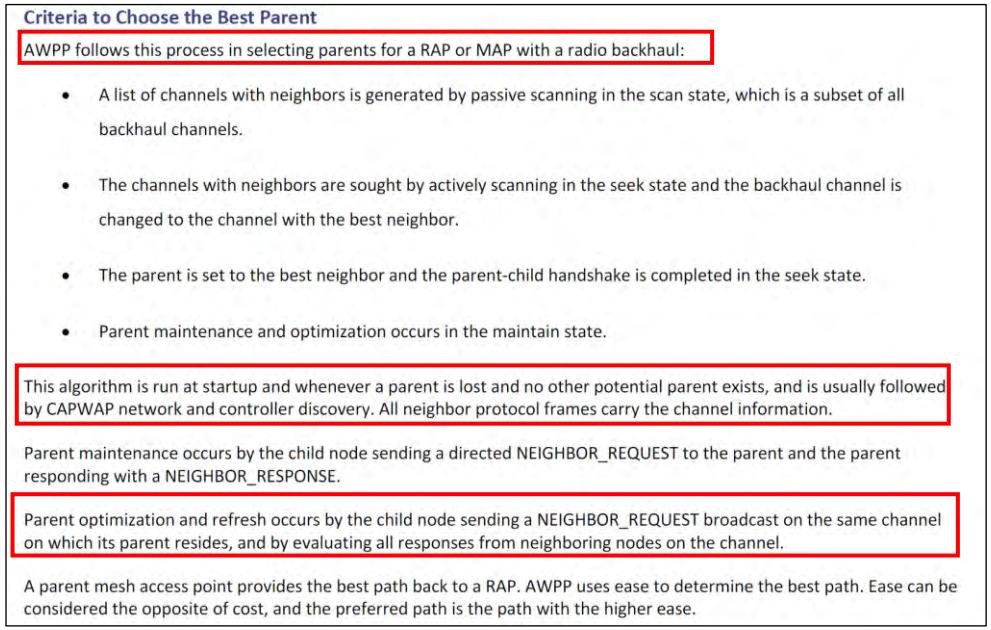
Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

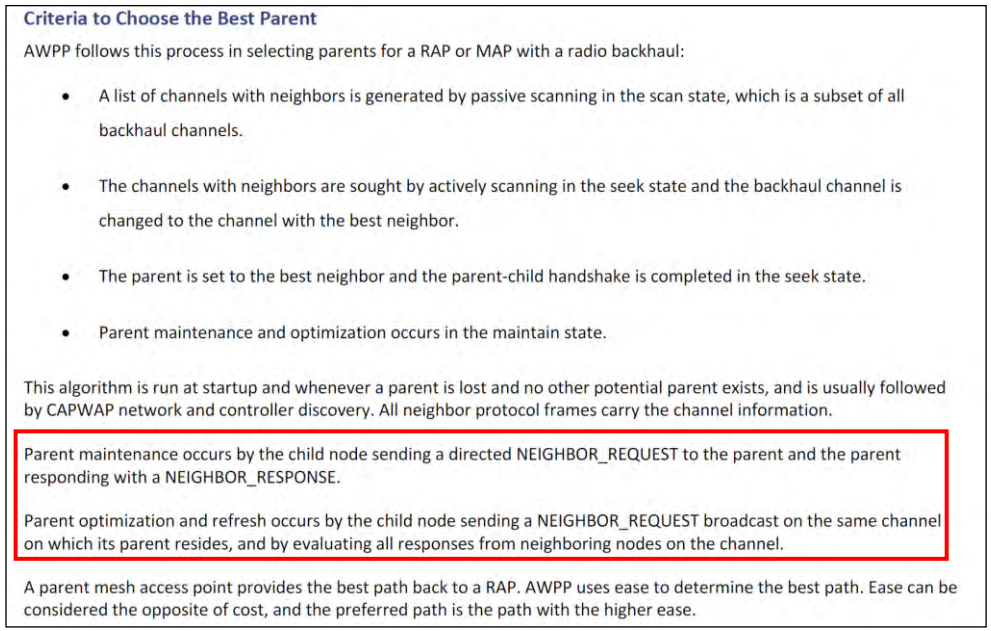
A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of making a channel selection change based on a change in network topology in the same way by automatically detecting topology changes (for example, the loss of the first node's parent or the appearance of a new neighboring node) and scanning for an alternative parent on a different channel to reroute the first node's traffic, and achieves the same result by reestablishing the mesh link via a new parent-child connection after the topology change, thereby preserving network connectivity and performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products automatically detect changes in network topology (for example, if the first node's parent is lost or a new neighbor becomes available) and responds by selecting an alternate parent on a different channel to reroute the first node's traffic, thus making the same kind of channel selection change based on a topology change as described in the claim. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

		MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 10	
10	10. A method as recited in claim 1, further comprising repeating b)-d).	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network further comprising repeating steps b)-d). Cisco Catalyst 9100 and Aironet Access Points repeatedly perform steps (b)–(d) as an ongoing adaptive process. In particular, the AWPP-based parent selection algorithm runs initially at network startup and then again whenever conditions require it (for example, when a parent is lost or during periodic optimization of the mesh). Thus, the sequence of steps (b)–(d) is continually repeated by the access points in the Cisco Wireless Mesh Network to maintain optimal routing.  **Criteria to Choose the Best Parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: • A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels. • The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor. • The parent is set to the best neighbor and the parent-child handshake is completed in the seek state. • Parent maintenance and optimization occurs in the maintain state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease. Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-

		network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of repeatedly performing steps (b)–(d) of the channel changing process as needed in the same way by using AWPP to continuously run Cisco's mesh routing algorithm (initially at network startup and again whenever conditions require, such as upon parent node failure or during periodic optimization) to re-execute the detection and channel-selection steps on an ongoing basis, and achieves the same result by continually reapplying steps (b)–(d) to adapt the mesh network and maintain optimal routing of traffic over time. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's AWPP protocol continuously runs the mesh routing algorithm such that it repeatedly performs the channel-change process (steps of detecting need, selecting new node, and selecting channel) whenever conditions require – for example, at startup, upon parent failures, or during periodic optimizations – which is effectively the same repetitive execution of steps (b)–(d) as needed to maintain optimal routing. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 11	
11	11. A method as recited in claim 1, further comprising receiving a monitoring signal that includes information about available nodes from which the second node is selected.	Cisco Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, further comprising receiving a monitoring signal that includes information about available nodes from which the second node is selected. Cisco Catalyst 9100 and Aironet Access Points receive and utilize "neighbor" signals carrying information about other available nodes to facilitate the dynamic selection of a new parent node. In particular, a MAP or child sends out Neighbor Discovery requests and receives Neighbor Response messages from nearby APs, which provide information (*e.g.*, identity and signal strength of neighboring nodes) that is used to identify and select the next parent node.



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of receiving a monitoring signal that includes information about available nodes from which the second node is selected in the same way by utilizing AWPP's neighbor discovery mechanism in which mesh access points exchange signals (*e.g.*, the first node receives neighbor response messages containing identifiers, etc.) to learn which nodes are available, and achieves the same result by obtaining information about candidate neighboring nodes so that the second node can be identified and selected as the new parent node based on that information. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products exchange neighbor discovery signals carrying information about other available access points (such as identifiers and signal strengths of nearby nodes), so that each node receives data about which neighboring nodes are available for selection – an equivalent mechanism to the claimed monitoring signal with information about available nodes from which the second node is selected. A person skilled in the relevant

		art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 12	
12.pre	12. A computer program product for controlling a wireless mesh network, the computer program product being embodied in a computer readable medium and comprising computer instructions for:	*See* 1.Pre-1.d above, the contents of which are incorporated herein by reference. Cisco Catalyst 9100 and Aironet Access Points comprise a computer program product (Cisco's proprietary AWPP algorithms and related software, *e.g.*, IOS XE and/or a universal IOS XE image or other firmware) for controlling a wireless mesh network, the computer program product being embodied in a computer readable medium (*e.g.*, the storage device or memory on the Catalyst 9100 and Aironet Access Points) with all of the claimed computer instructions, as shown in claim 1 *supra*. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
12.a	a) configuring the wireless mesh network with a first configuration, the wireless mesh network comprising a plurality of relay nodes, each of the relay nodes being capable of relaying traffic from a child node associated with the relay node, each of the relay nodes including a first radio interface configurable to relay traffic to a parent relay node and a second dual	*See* 1.a above, the contents of which are incorporated herein by reference.

	purpose radio interface configurable to both relay traffic from the child relay node and to communicate with a client node that is not a relay node, and at least one of the relay nodes is configured in the first configuration to relay traffic from a child relay node to a parent relay node;	
12.b	b) determining that a first one of the plurality of relay nodes in the wireless mesh network requires a channel selection change;	*See* 1.b above, the contents of which are incorporated herein by reference.
12.c	c) dynamically selecting a second one of the plurality of relay nodes in the wireless mesh network to form a parent-child relationship with the first relay node; and	*See* 1.c above, the contents of which are incorporated herein by reference.
12.d	d) making a channel selection change; wherein prior to being dynamically selected, the second relay node was a sibling node of the first relay node; and making a channel	*See* 1.d above, the contents of which are incorporated herein by reference.

	selection includes making a channel selection change to meet a performance requirement.	
	Claim 13	
13	13. A computer program product as recited in claim 12, wherein the second node is selected to be a parent node of the first node.	*See* claims 2 and 12 above, the contents of which are incorporated herein by reference.
	Claim 14	
14	14. A computer program product as recited in claim 12, wherein the second node is selected to be a child node of the first node.	*See* claims 3 and 12 above, the contents of which are incorporated herein by reference.
	Claim 15	
15	15. A computer program product as recited in claim 12, wherein making a channel selection includes making a channel selection change to meet a latency requirement.	*See* claims 4 and 12 above, the contents of which are incorporated herein by reference.
	Claim 16	
16	16. A computer program product as recited in claim 12, wherein making a channel	*See* claims 5 and 12 above, the contents of which are incorporated herein by reference.

	selection includes making a channel selection change to meet a throughput requirement.	
	Claim 17	
17	17. A computer program product as recited in claim 12, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to a performance requirement change.	*See* claims 6 and 12 above, the contents of which are incorporated herein by reference.
	Claim 18	
18	18. A computer program product as recited in claim 12, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to congestion.	*See* claims 7 and 12 above, the contents of which are incorporated herein by reference.
	Claim 19	
19	19. A computer program product as recited in claim 12, wherein the	*See* claims 9 and 12 above, the contents of which are incorporated herein by reference.

	channel selection change is made based on network topology.	
	Claim 20	
20	20. A computer program product as recited in claim 12, the computer program product further comprising computer instructions for repeating b)-d).	*See* claims 10 and 12 above, the contents of which are incorporated herein by reference.
	Claim 21	
21	21. A computer program product as recited in claim 12, the computer program product further comprising computer instructions for receiving a monitoring signal that includes information about available nodes from which the second node is selected.	*See* claims 11 and 12 above, the contents of which are incorporated herein by reference.

Exhibit 2 - U.S. Patent No. 7,420,952 (Cisco Wireless Mesh Networks with Meraki and Catalyst CW Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Meraki Wireless Mesh Network including but not limited to Cisco's Meraki and Catalyst CW Access Points, with related software and cloud architecture (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1-21 (the "Asserted Claims") of U.S. Patent No. 7,420,952 (the "'952 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '952 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco's access points directly infringed at least Claim 12 of the '952 Patent when made, used, sold, offered for sale and imported into the United States. Cisco further directly infringed by testing, setting up, and using the Cisco Meraki Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Meraki Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Meraki Wireless Mesh Network in the United States. As indicated in the Meraki Mesh Networking Guide, each Meraki access point has unique mesh characteristics independent of general technology specifications listed on the datasheet. Cisco recommends a Meraki Systems Engineer as a resource to ensure that an access point's capabilities match the desired design targets.[3] The Meraki Trial Program allows customers to run a real-world proof of concept to test the design, which is arranged by a Cisco Meraki account representative. Meraki is also cloud-based architecture, and the infringing mesh networking methods and systems were controlled and put into service by Cisco.[4] Cisco knew how its customers were configuring the specific access points sold, as customers must register to the Meraki cloud and use the invoice or serial number of each purchased access point.[5]

---

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g*., https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

[3] *See, e.g*., https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

[4] *See Centillion Data Systems, LLC v. Qwest Communications International, Inc.*, 631 F.3d 1279 (Fed. Cir. 2011); *NTP, Inc. v. Research in Motion, Ltd*., 418 F. 3d 1282, 1317-1318 (Fed. Cir. 2005); *see also SiRF Tech. Inc. v. ITC*, No. 2009-1262 (Fed. Cir. Apr. 12, 2010) (direct infringement where the claims do not specifically require action by its customers or end users).

[5] *See, e.g*., https://documentation.meraki.com/MR/MR_Installation_Guides/CW9163E-MR_Installation_Guide.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the ’952 Patent and make and use the infringing Cisco Meraki Wireless Mesh Network. Specifically, Cisco’s Meraki Mesh Deployment Guides, Meraki Mesh Networking Guide, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco’s website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Meraki Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-21 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Meraki Wireless Mesh Network, and the infringing configuration and operation of the cloud architecture, access points and software that comprise the infringing Cisco Meraki Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Meraki Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the ’952 Patent since 2009 when Cisco executives met with MeshDynamics to discuss partnering opportunities and discussed the ’952 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Meraki Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’952 Patent, or in the alternative, was willfully blind to the ’952 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Meraki Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’952 Patent. Cisco Accused Products come enabled as mesh access points and automatically connect to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particularly suited for Cisco’s Meraki Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Meraki Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale the Cisco Accused Products to other entities, the other entities used these products for their intended purpose and according to their instructions with the result that such entities directly infringed the ’952 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the ’952 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such

claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics's preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

#	Claim 1	Cisco's Meraki Wireless Mesh Network with Meraki and Catalyst CW Access Points
1.pre	1. A method of controlling a wireless mesh network, comprising:	Cisco's Meraki Wireless Mesh Network comprises a plurality of Wi-Fi relay nodes (*e.g.*, Cisco Meraki and Catalyst CW Access Points), which employ and provide a method of controlling a wireless mesh network. The Meraki cloud-managed wireless series is a line of enterprise-grade WLAN access points ("APs") that leverage the Cisco Meraki cloud-managed architecture. Both Cisco Meraki and Catalyst CW Access Points support Wi-Fi capabilities and can be paired with the Meraki cloud platform. Meraki AP model numbers start with MR, *e.g.* MR33, MR36, MR46E, and the Catalyst CW APs with CW, *e.g.* CW9162, CW9164, CW9166, are dual-mode APs which can be used with the Catalyst 9800 controller or the Meraki cloud architecture.

Source: Meraki Family Datasheet; Meraki Wi-Fi Models

Cisco Meraki Cloud Management

Pairing the Cisco Catalyst 9166 Series Access Points with the Meraki cloud platform gives organizations a unified IT experience for network monitoring and management. The Meraki dashboard provides an intuitive and interactive web interface connecting your network to the industry's leading cloud IT platform.

Through the dashboard, Meraki provides sophisticated and scalable tools to automate network optimization, deploy policy and segmentation configurations across thousands of sites and devices, and manage a full-stack network from SD-WAN to Access to IoT technologies. The platform supports over 3.5 million active networks around the world.

Working together, the Catalyst 9166 Series and Cisco Meraki offer such features as:

- Cisco Spaces
- Cisco Identity Services Engine
- Meraki Health intelligent optimization and assurance
- Meraki Vision, smart cameras, and sensors for network closet monitoring

Source: CW9166 Datasheet; CW9166

The Meraki and Catalyst CW Access Points contain the same mesh networking capabilities, as shown in Cisco's documentation and the Meraki Family Datasheet, and the Meraki Wireless Mesh Networking and Mesh Deployment Guide apply to both product lines for setting up a Cisco Meraki Wireless Mesh Network. Multiple access points communicate over wireless interfaces to form a single network. Meraki's mesh networking functionality is automatic, self-healing and available on all Meraki (including Catalyst CW) Access Points.

cisco Meraki

Wireless Mesh Networking

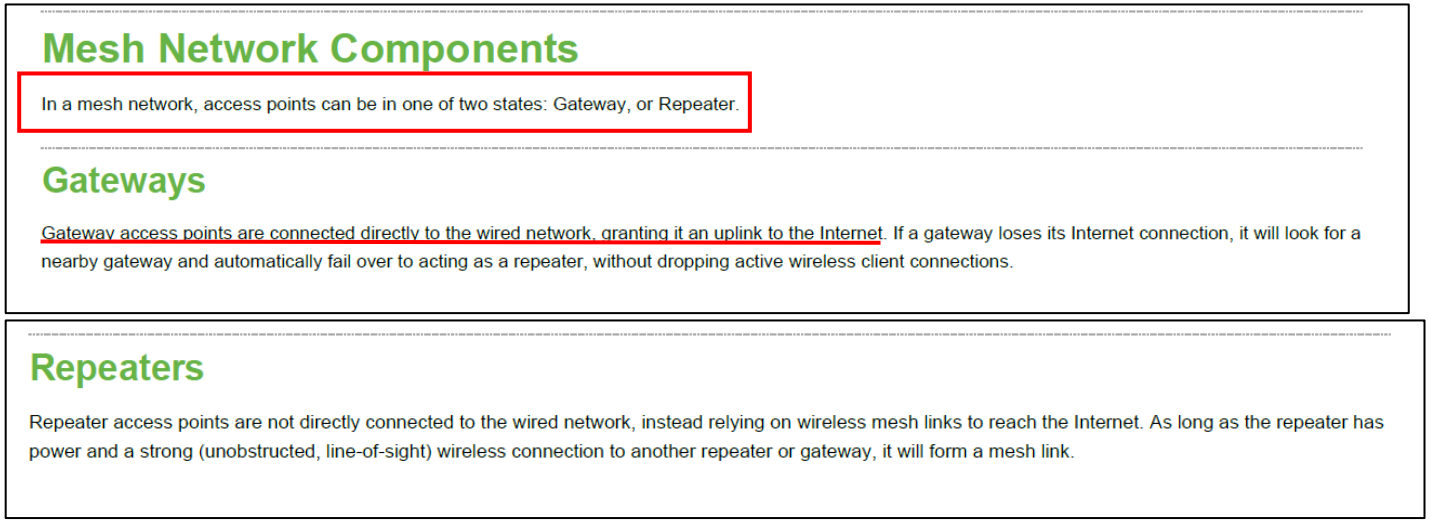
In a wireless mesh deployment, multiple Access points (with or without Ethernet connections) communicate over wireless interfaces to form a single network. This wireless communication between Access points is called Mesh Networking. Meraki's mesh networking functionality is automatic, self-healing and available on all access points.

A detailed article on designing wireless Mesh networks can be found on Mesh Deployment Guide.

Source: Wireless Mesh Networking; *see also* Meraki Best Practices - Wireless Mesh Networking

Mesh functionality is enabled by default in every Meraki AP. All Meraki APs support wireless meshing without any additional configuration ("zero-configuration mesh") or licensing.

		Access Point Capabilities All Meraki access points support wireless meshing without any additional configuration or licensing. The most successful mesh deployments undergo both pre-deployment design and post-design testing phases. Source: Mesh Deployment Guide **Class-leading enterprise features** Meraki MR wireless access points come equipped with industry-leading features that make them ideal for demanding enterprise deployments: • Self-configuring plug-and-play deployment • 802.11be MU-MIMO with up to sixteen spatial streams built for voice and video • Integrated Bluetooth® IoT radio • IoT ready (ESL integration) • Self-healing, zero-configuration mesh Source: Meraki Family Datasheet Mesh networking, included in every Meraki AP, extends coverage to hard-to-wire areas and creates a self-healing network that is resilient to cable and switch failures, continuing to operate despite failures or configuration changes in the rest of the network without the need for manual configuration or optimization. Source: Meraki Family Datasheet
1.a	a) configuring the wireless mesh network	Cisco's Meraki Wireless Mesh Network (as shown in Cisco's documentation) comprises a plurality of relay nodes (*e.g.*, Cisco Meraki or Catalyst CW Access Points), which employ and provide a method of

	with a first configuration, the wireless mesh network comprising a plurality of relay nodes, each of the relay nodes being capable of relaying traffic from a child node associated with the relay node, each of the relay nodes including a first radio interface configurable to relay traffic to a parent relay node and a second dual purpose radio interface configurable to both relay traffic from the child relay node and to communicate with a client node that is not a relay node, and at least one of the relay nodes is configured in the first configuration to relay traffic from a child relay node to a parent relay node;	controlling a wireless mesh network comprising the step of configuring the wireless mesh network with a first configuration: APs within the Cisco Meraki Wireless Mesh Network operate as either a gateway (*i.e.*, connected to the wired network) or repeater (*i.e.*, relying on wireless mesh links to communicate and reach the network), and are capable of relaying traffic from a child node (*i.e.*, repeater) associated with a parent relay node (*i.e.*, gateway or other repeater) through a mesh link and include a first radio interface configurable to relay traffic to the parent relay node.  Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide Cisco Meraki and Catalyst CW Access Points perform the step of configuring the Cisco Meraki Mesh Network in a first configuration, as the access points are configured in the "first configuration" to relay traffic from a child relay node (*i.e.*, repeater) to a parent relay node (*i.e.*, repeater or gateway). All Cisco Meraki access points are configured and shipped to support mesh networking. Meraki uses a proprietary routing protocol called "Meraki Mesh Algorithm," which on information and belief, operates the same or similar to Cisco's Adaptive Wireless Path Protocol ("AWPP") and forms parent-child relationships in the mesh, with parents selected based on criteria or characteristics of the wireless network including link quality and performance through multiple hops. Like AWPP (*see* Exhibit 1 for AWPP details, incorporated herein by reference), and the self-healing nature of meshing, the APs will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki Access

Points (and the Catalyst CW Access Points that can be used in a Meraki Wireless Mesh Network) are available for connections from repeaters (*i.e.*, are available to form a parent-child relationship and capable of relaying traffic from a child or repeater).

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode.

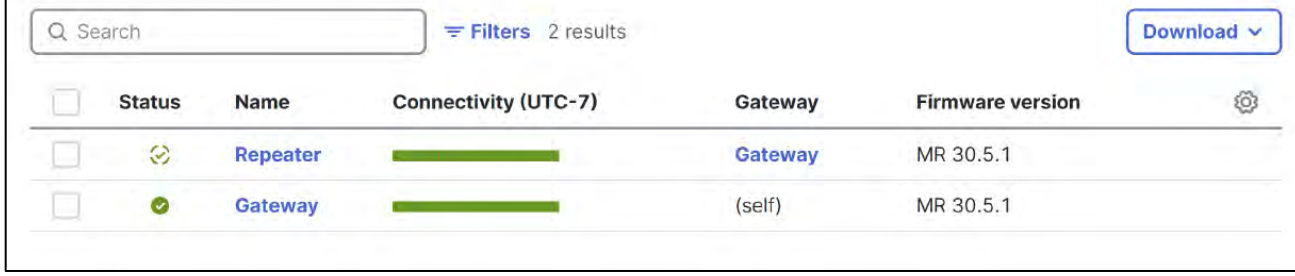
Source: Wireless Mesh Networking

Repeaters will identify the other AP in the network that it is using as its gateway, and its route back to the gateway, as well as identify and track neighbors. Each repeater knows the distance to its neighbors as well as the link quality. This is shown at the access server using identifier values to identify the repeaters, *e.g.*, “Name.”

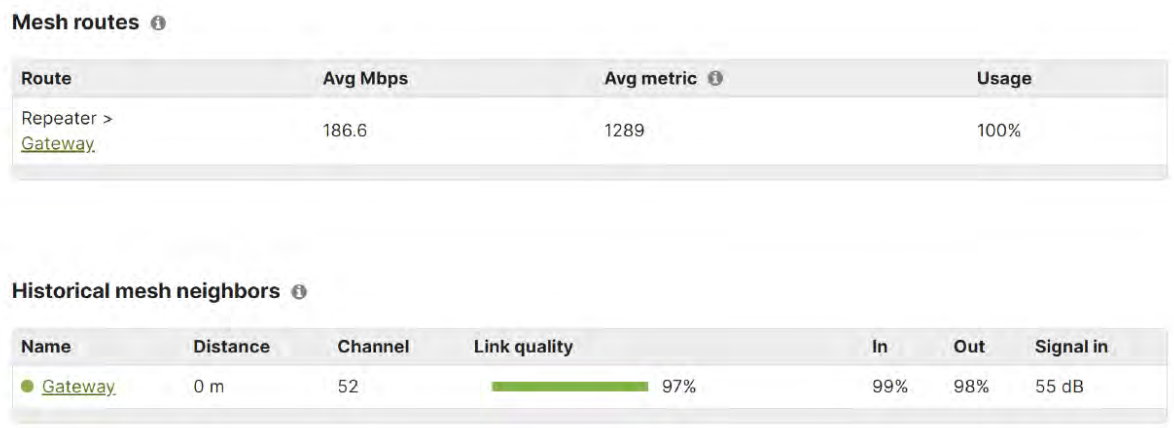
Identifying a Repeater access point vs Gateway access point

1. Navigate to **Wireless > Monitor > Access points**
2. Click the Gear icon on the right-hand side and make sure the **Gateway** option is selected:

3. A gateway access point will be listed as "(self)" under the **Gateway** column, while a repeater access point will list some other access point in the network (denoting that it is using that access point as its gateway):

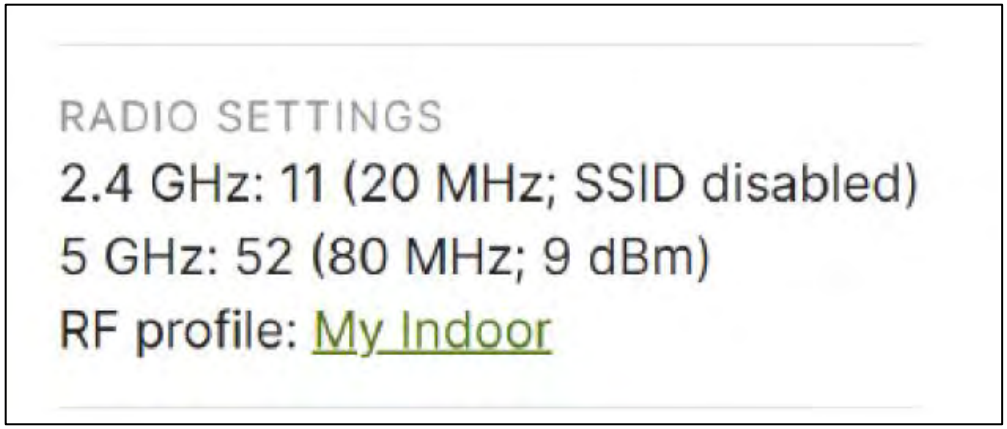


7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:

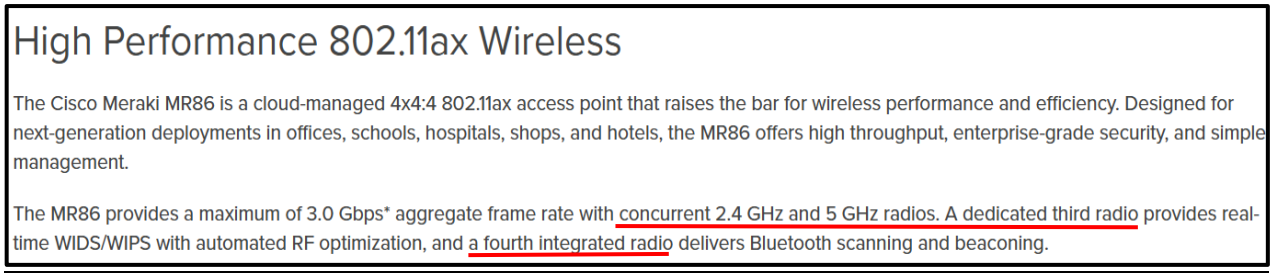


Source: Wireless Mesh Networking

APs in the Cisco Meraki Wireless Mesh Network have multiple configurable radio interfaces, including 2.4 and 5 GHz radios.



Source: Wireless Mesh Networking



Source: MR86 Datasheet

Product options: Wi-Fi 6E indoor models

	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: Meraki Family Datasheet

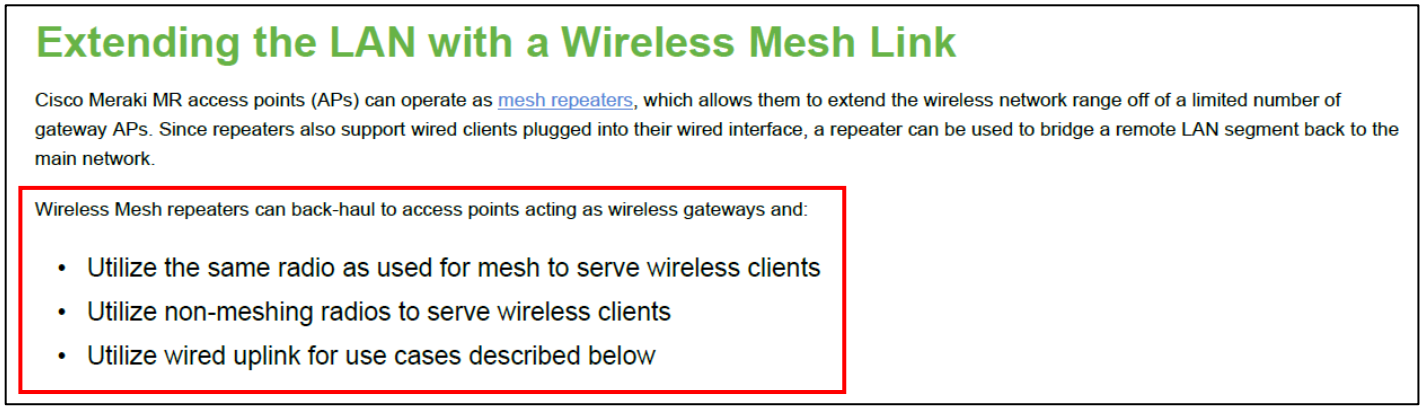
Radios may be used concurrently for meshing while also serving wireless clients (*e.g.*, the claimed "a first radio interface configurable to relay traffic to a parent relay node and a second dual purpose radio interface configurable to both relay traffic from the child relay node and to communicate with a client node").

Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

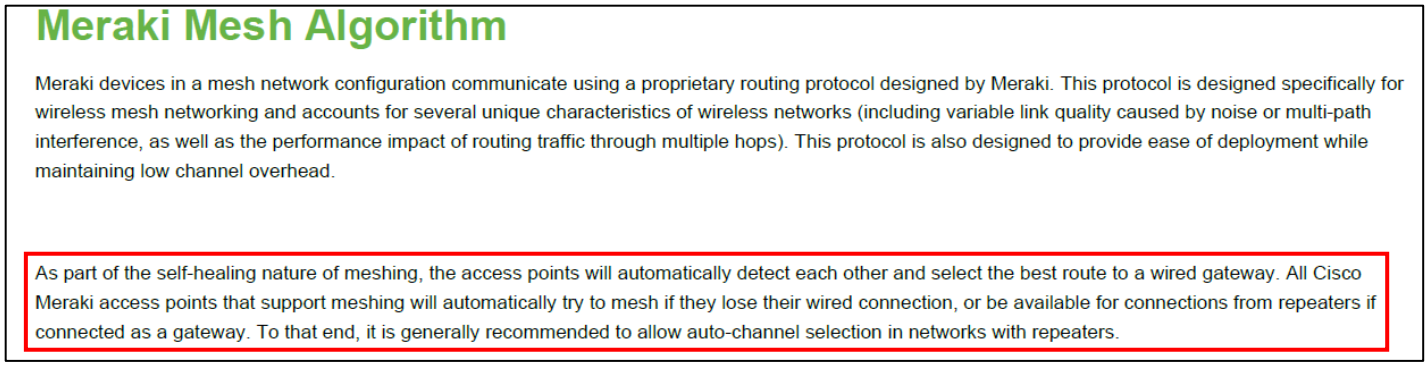
Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: Mesh Deployment Guide

		 Source: Extending the LAN with a Wireless Mesh Link If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of establishing a first wireless mesh configuration of multiple relay nodes (each with dual radio interfaces for relaying traffic and serving client nodes) in the same way by using the Meraki Mesh Algorithm to deploy multiple nodes, each equipped with a backhaul radio and a dual-purpose radio (or tri-radio mode) and configuring at least one node (*e.g.*, an access point) to relay a child node's traffic to a designated parent node (*e.g.*, another access point), and achieves the same result by creating a wireless mesh network in which a relay node forwards traffic from its child node to its parent node, while the child serves clients. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's Meraki Mesh Algorithm is used to deploy multiple dual-radio (or tri-radio) mesh access points to relay traffic via dedicated backhaul links, achieving the same multi-node mesh configuration of relay nodes serving clients as is claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.b	b) determining that a first one of the plurality of relay nodes in the wireless mesh network requires a channel selection change;	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network employ and provide a method of controlling a wireless mesh network, comprising the step of determining that a first one of the plurality of relay nodes in the wireless mesh network requires a channel selection change. Cisco Meraki and Catalyst CW Access Points ("plurality of relay nodes in the wireless mesh network") can determine (*e.g.*, for gateway or parent optimization and upon the loss of a gateway or parent connection) that it

should change channels and connect with a better potential parent across any channel to maintain its uplink with the best parent (*i.e.*, the best path back to the gateway based on link quality and number of hops). This required "channel selection change" will choose the best parent or gateway.

Cisco Meraki and Catalyst CW Access Points will automatically detect each other and select the best route to a wired gateway. Cisco Meraki Mesh Algorithm has an "auto channel selection" feature in networks with repeaters or child nodes.



Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

Repeaters in a Cisco Meraki Wireless Mesh Network look for an AP (*e.g.*, determining that it requires a channel selection change) when the gateway is down or it finds a better gateway. AP will scan the entire spectrum in the first case. In the second, a repeater is constantly evaluating the current channel for better gateways. If a repeater finds an AP with a better link quality metric, it determines whether to change the channel and, if so, will move to the new gateway.

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: Wireless Mesh Networking

Each AP in Cisco's Meraki Wireless Mesh Network also scans all channels to collect information from its neighbors. If a valid neighbor is found, the AP switches to the channel. Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. The number of hops is considered as well, with preference give to gateways with the lower number of hops. An AP finalizes its route based on link quality and number of hops and will mesh on the channels with the least interference using channel selection algorithms.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in.

Source: Manually Changing Channels

		If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of determining when a first relay node (*e.g.*, a access point) in the mesh requires a channel selection change (*i.e.*, detecting that the node's current channel must be changed) in the same way by employing a dynamic channel selection mechanism that monitors wireless conditions and identifies when the relay node's channel needs to be switched, and achieves the same result by recognizing that a channel selection change is necessary to maintain proper network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's software continuously monitors wireless conditions and identifies when an access point's current channel must be changed, accomplishing the same function of determining that a channel selection change is required as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.c	c) dynamically selecting a second one of the plurality of relay nodes in the wireless mesh network to form a parent-child relationship with the first relay node; and	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network employ and provide a method of controlling a wireless mesh network, comprising the step of dynamically selecting a second one of the plurality of relay nodes in the wireless mesh network to form a parent-child relationship with the first relay node. As discussed above, Cisco's Meraki Mesh Algorithm was designed specifically for mesh networking and enables (like AWPP) each relay node (*e.g.*, Cisco Meraki and Catalyst CW Access Points) to dynamically select the best AP to form a parent-child relationship and the best path to the gateway. In the event of a mesh gateway failure or a gateway with a better routing metric, all new traffic flows will be routed to the new mesh gateway (*i.e.*, the repeater will dynamically select a gateway or another repeater in the best path to form a parent-child relationship).

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

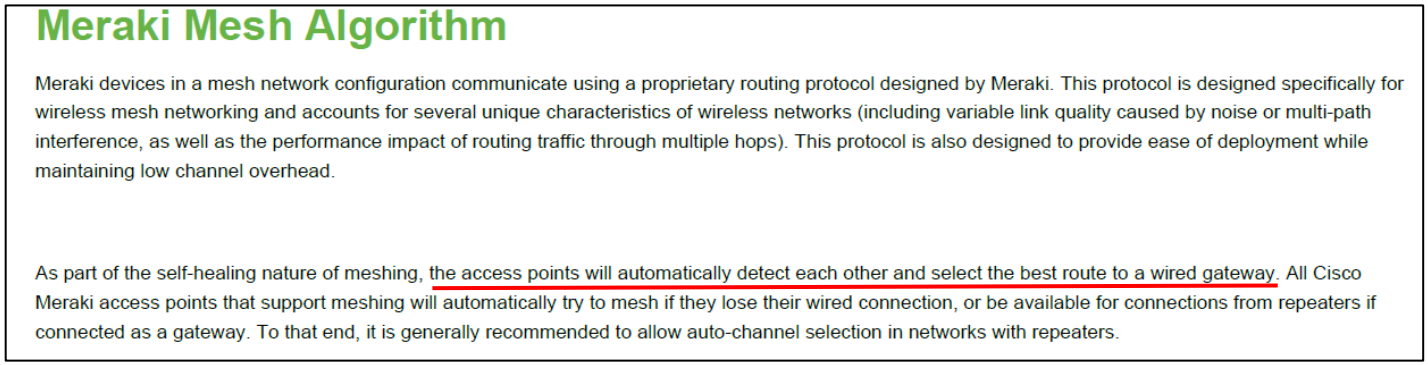
Source: Wireless Mesh Networking

Repeaters will automatically choose the gateway that has the strongest connection.

Note: Both gateways and repeaters can serve wireless clients. It is possible to have multiple gateways in a mesh network, and repeaters will automatically choose the gateway to which it has the strongest connection.

Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of dynamically selecting a second relay node (*e.g.*, an access point) to form a parent-child relationship with the first relay node (*e.g.*, an access point) in the same way by using the Meraki Mesh Algorithm to evaluate available neighboring and gateway nodes and choose the best candidate as the new parent for the first node, and achieves the same result by establishing a parent-child link between the first node and the selected second node. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's Meraki Mesh

		Algorithm dynamically selects the optimal neighboring access point or gateway to serve as a new parent, thereby forming the same parent-child relationship between the first node and the selected second node as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.d	d) making a channel selection change wherein prior to being dynamically selected, the second relay node was a sibling node of the first relay node; and making a channel selection includes making a channel selection change to meet a performance requirement.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network employ and provide a method of controlling a wireless mesh network, comprising the step of making a channel selection change wherein prior to being dynamically selected the second relay node (*e.g.*, a gateway or repeater) was a sibling node of the first relay node (*e.g.*, a neighboring AP); and making a channel selection includes making a channel selection change to meet a performance requirement (*e.g.*, better routing metric). Meraki's Mesh Algorithm follows a dynamic process for selecting a parent (*i.e.*, gateway or repeater) to make any necessary channel selection change, including making a channel selection change to meet a performance requirement (*i.e.*, by choosing a parent or gateway (based on the best path link quality, number of hops, and/or least interference)). Meraki's Mesh Algorithm sets the parent (*e.g.*, changes or selects the backhaul channel) to the best neighbor (*e.g.*, "sibling"), and the APs will automatically mesh on the best channel using the channel selection algorithms.  Source: Wireless Mesh Networking

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in.

Source: Manually Changing Channels

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of making a channel selection change that promotes a neighboring node to become the new parent of the first node in order to meet a performance requirement in the same way by dynamically reassigning the first node's backhaul connection to an optimal neighboring relay node based on metrics such as signal quality, latency, throughput, etc. and achieves the same result by improving the mesh network's performance (for example, reducing latency or hop count) through the channel selection change while establishing the second node as the first node's parent. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's Meraki Mesh Algorithm automatically switches a first node's backhaul connection to the "best" neighboring node (previously a sibling node) as a new parent whenever needed to meet network performance requirements (such as better signal or fewer hops), thereby making the same performance-driven channel

		selection change to a new parent as is claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 2	
2	2. A method as recited in claim 1, wherein the second node is selected to be a parent node of the first node.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the second node is selected to be a parent node of the first node. Cisco Meraki and Catalyst CW Access Points in a mesh network implement this limitation by dynamically selecting a neighboring node to serve as a new parent for the first node. The Meraki Mesh Algorithm running on each mesh AP will set the parent to the "best" neighboring AP and establish a parent-child link, thereby selecting the second node as the parent node of the first node. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: Wireless Mesh Networking Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. Source: Wireless Mesh Networking

		If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of selecting the second node to serve as the parent node of the first node in the same way by utilizing the Meraki Mesh Algorithm to dynamically designate the best neighboring access point as the new parent for the first access point, and achieves the same result by establishing the second node as the first node's parent within the wireless mesh topology. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh nodes dynamically choose a neighboring access point to act as the new parent node for the first node (via Meraki Mesh Algorithm's parent selection algorithm), fulfilling the same function of selecting the second node as a parent for the first node as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 3	
3	3. A method as recited in claim 1, wherein the second node is selected to be a child node of the first node.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the second node is selected to be a child node of the first node. Cisco Meraki and Catalyst CW Access Points in a mesh network implement this limitation by dynamically selecting the first node to serve as a new parent for the second node. The Meraki Mesh Algorithm running on each mesh AP will set the parent to the "best" neighboring AP and establish a parent-child link, thereby selecting the second node as the child node of the first node. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: Wireless Mesh Networking

		Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of selecting the second node (*e.g.*, an access point) to be a child node of the first node (*e.g.*, another access point) in the same way by dynamically reconfiguring the mesh so that the first access point becomes the parent and the second access point is assigned as its child by using the Meraki Mesh Algorithm, and achieves the same result by establishing a new parent-child relationship in which the second node is connected as a child under the first node. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's system can equivalently reconfigure node roles such that the first access point becomes a parent and the second access point is assigned as its child (using Meraki Mesh Algorithm's dynamic reconfiguration), achieving the same result of establishing a new parent-child relationship with the second node as the child under the first node as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 4	
4	4. A method as recited in claim 1, wherein making a channel selection includes making a channel selection change to meet a latency requirement.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein making a channel selection includes making a channel selection change to meet a latency requirement, including Cisco's routing metric requirement, which considers the latency related to link quality and the "hops" or "lower hops" back to the gateway. Cisco Meraki and Catalyst CW Access Points perform channel selection changes to meet specific performance requirements such as latency. For example, when using the Meraki Mesh Algorithm, Cisco's Meraki and Catalyst CW Access may switch the backhaul channel (following, *e.g.*, establishing the parent-child relationship

		as shown above) in order to reduce latency and number of hops thereby making a channel selection change to satisfy a latency requirement as recited. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of making a channel selection change to meet a latency requirement in the same way by automatically switching the first node's wireless backhaul to a better channel or alternate parent node whenever network latency exceeds a defined threshold and achieves the same result by reducing communication latency so that the mesh network maintains the required performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco access points automatically trigger a channel selection change (or parent switch) whenever network latency rises above a defined threshold – for example, the mesh software will select and switch the first node's backhaul channel or parent to a better option if latency becomes too high – thereby making a channel selection change to meet the same latency requirement as in claimed method. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 5	
5	5. A method as recited in claim 1, wherein making a channel selection includes making a channel selection change to meet	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein making a channel selection includes making a channel selection change to meet a throughput requirement, including Cisco's routing metric requirement, which considers, along with the number of hops, the throughput related to link quality and signal strength.

	a throughput requirement.	Cisco Meraki and Catalyst CW Access Points perform channel selection changes to meet specific performance requirements such as throughput based on link hop values and/or signal strength. For example, when using the Meraki Mesh Algorithm, the access point may switch the backhaul channel in order to meet throughput requirements associated with high link hop values and lower signal strength in backhaul data traffic, thereby making a channel selection change to meet a throughput requirement as recited. **Mesh Neighbors** The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link. Source: Wireless Mesh Networking **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking **Impact of Meshing on Throughput** Due to the half-duplex nature of wireless communication and that signals being passed through a repeater access point must be retransmitted to the next hop, throughput is greatly reduced when using a repeater. While many factors impact wireless throughput, it is safe to assume that the addition of meshing can reduce throughput by approximately 50%, with that reduction being applied for each subsequent repeater that must be traversed to reach a gateway. Therefore, it is advised to minimize the number of hops between a client and gateway. Source: Wireless Mesh Networking

		If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of making a channel selection change to meet a throughput requirement in the same way by dynamically adjusting the channel or route based on throughput-related metrics (for example, detecting excessive link hops or low signal strength and switching to a channel that provides higher capacity), and achieves the same result by increasing the data throughput so that the mesh network meets the required throughput level. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's channel selection algorithms dynamically adjust the wireless backhaul (by changing the channel and routing path) based on throughput-related metrics – for example, detecting excessive link hops or low signal strength and switching to a higher-capacity channel – thereby making the same kind of channel selection change to meet a throughput requirement as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 6	
6	6. A method as recited in claim 1, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to a performance requirement change.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, further employ and provide a method of controlling a wireless mesh network, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to a performance requirement change (*e.g.*, a performance requirement change associated with a link hop value increase). As shown *supra* for claims 4 and 5 (the contents of which are incorporated herein by reference), Cisco Meraki and Catalyst CW Access Points perform channel selection changes to meet specific performance requirements such as throughput and/or latency. For example, when using the Meraki Mesh Algorithm, the access points may switch the backhaul channel in order to meet performance requirements associated with high or updated link hop values in backhaul data traffic, thereby making a channel selection change to satisfy a performance requirement change as recited. As another example, when using the Meraki Mesh Algorithm, Cisco Meraki and Catalyst CW Access Points may switch the backhaul channel in order to increase throughput in backhaul data traffic.

Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

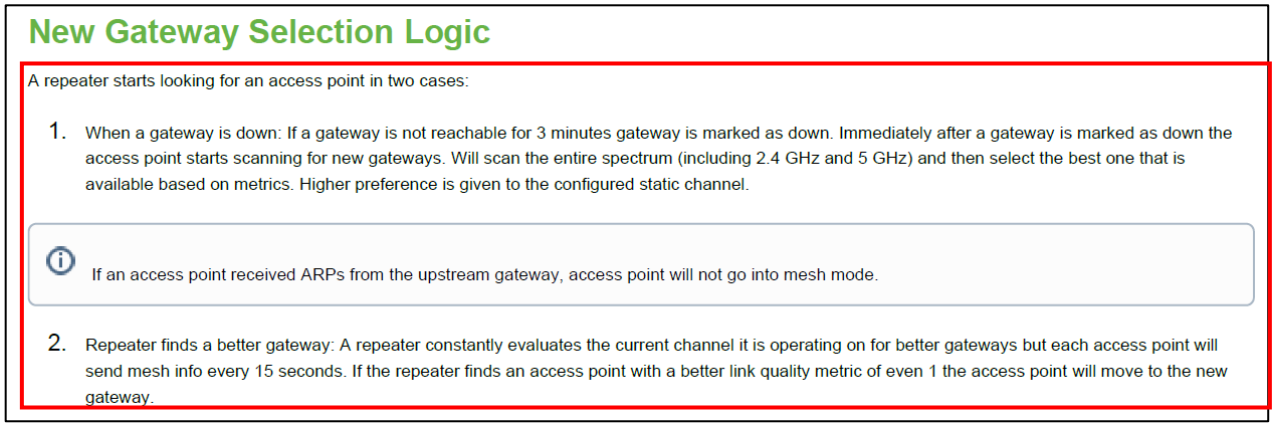
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of determining that a channel selection change is required due to a change in a performance requirement in the same way by monitoring network performance and identifying when an updated performance requirement indicates the current channel is inadequate, and achieves the same result by detecting that the first node must switch channels in response to the changed performance requirement. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products monitor network performance criteria and can tell when an updated performance requirement makes the current channel inadequate, thus detecting the need for a channel change due to the changed performance requirement in essentially the same way as the claimed invention. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

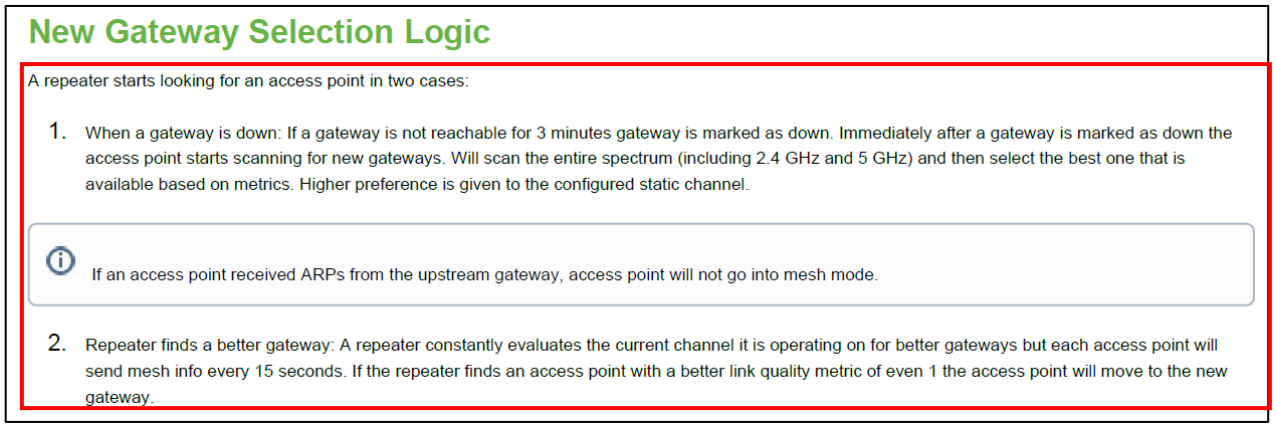
	Claim 7	
7	7. A method as recited in claim 1, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to congestion.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to congestion. Cisco Meraki and Catalyst CW Access Points channel selection process accounts for congestion and will change channels to alleviate noise and interference problems associated with congestion when needed. Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network determine that the first node requires a channel selection change by determining that the channel selection change is required due to congestion by considering link quality, number of hops, and interference based on their parent selection algorithm. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels The Meraki Mesh Algorithm's and the AP's focus on factors like link quality/signal strength, reducing the number of hops and interference helps to minimize and avoid congestion in the mesh network. For example, lower throughput is an indication of congestion. By choosing the "best" path to the wired network based on these criteria, and switching parents, the Meraki Mesh Algorithm addresses the level of network congestion experienced by an available parent node. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of determining that a channel selection change is

		required due to congestion in the same way by analyzing channel utilization (via Cisco's parent selection algorithms that consider metrics like interference and link quality) to decide when the first node's channel is overly congested and must be changed, and achieves the same result by triggering a channel switch whenever the current channel becomes too congested, thereby alleviating congestion in the mesh network. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products' parent selection algorithm takes congestion into account by measuring load values, and it automatically switches the first node's backhaul to a less-loaded channel when the current channel becomes overly congested, thereby performing the same congestion-triggered channel change as in the claim. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 8	
8	8. A method as recited in claim 1, wherein the channel selection change is made to avoid channel interference.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the channel selection change is made to avoid channel interference. Cisco Meraki and Catalyst CW Access Points channel selection process accounts for channel interference and will change channels to alleviate channel interference problems when needed. Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network determine that the first node requires a channel selection change by determining that the channel selection change is required due to interference by meshing automatically on the channels with the least interference possible using channel selection algorithms. Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of making a channel selection change to avoid channel interference in the same way by continually evaluating the interference on the current channel (using Cisco's Meraki Mesh Algorithm to monitor interference) and switching the first node to a cleaner channel when

		interference is excessive, and achieves the same result by maintaining communications on a channel with substantially reduced interference. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products continually evaluate interference on the current channel and switches the first node to a cleaner channel whenever the existing channel suffers excessive interference, achieving the same result of avoiding channel interference as the claimed method. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 9	
9	9. A method as recited in claim 1, wherein the channel selection change is made based on a change to network topology.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, wherein the channel selection change is made based on a change to network topology. Cisco Meraki and Catalyst CW Access Points implement this step by dynamically selecting a neighboring node to serve as a new parent for the first node. The Meraki Mesh Algorithm running on each node will set the parent to the "best" neighboring AP and establish a parent-child link, thereby selecting the second node as the parent node of the first node. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking If a repeater's current parent "is lost" (*e.g.*, a topology change), the AP detects the loss and scans other channels to find and connect to a new parent on a different channel. In this way, a channel selection change is performed when the network topology changes. A new and better neighbor is another example of a topology change.

		 Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of making a channel selection change based on a change in network topology in the same way by automatically detecting topology changes (for example, the loss of the first node's parent or the appearance of a new neighboring node) and scanning for an alternative parent on a different channel to reroute the first node's traffic, and achieves the same result by reestablishing the mesh link via a new parent-child connection after the topology change, thereby preserving network connectivity and performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because the Cisco Accused Products automatically detect changes in network topology (for example, if the first node's parent is lost or a new neighbor becomes available) and respond by selecting an alternate parent on a different channel to reroute the first node's traffic, thus making the same kind of channel selection change based on a topology change as described in the claim. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 10	
10	10. A method as recited in claim 1, further comprising repeating b)-d).	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network further comprising repeating steps b)-d).

Cisco Meraki and Catalyst CW Access Points repeatedly perform steps (b)–(d) as an ongoing adaptive process. In particular, the Meraki Mesh Algorithm's parent selection algorithm runs initially at network startup and then again whenever conditions require it (for example, when a parent is lost or during periodic optimization of the mesh). Thus, the sequence of steps (b)–(d) is continually repeated by the access points in the Cisco Meraki Wireless Mesh Network to maintain optimal routing.



Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of repeatedly performing the channel changing process as needed in the same way by using the Meraki Mesh Algorithm to continuously run Cisco's mesh routing algorithm (initially at network startup and again whenever conditions require, such as upon parent node failure or during periodic optimization) to re-execute the detection and channel-selection steps on an ongoing basis, and achieves the same result by continually reapplying the algorithm to adapt the mesh network and maintain optimal routing of traffic over time. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's Meraki Mesh Algorithm continuously runs the mesh routing algorithm such that it repeatedly performs the channel-change process (steps of detecting need, selecting new node, and selecting channel) whenever conditions require – for example, at startup, upon parent failures, or during periodic optimizations – which is effectively the same repetitive execution of steps (b)–(d) as claimed. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

		MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 11	
11	11. A method as recited in claim 1, further comprising receiving a monitoring signal that includes information about available nodes from which the second node is selected.	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network, in addition to the steps of claim 1 as shown above, employ and provide a method of controlling a wireless mesh network, further comprising receiving a monitoring signal that includes information about available nodes from which the second node is selected. Cisco Meraki and Catalyst CW Access Points receive and utilize "mesh probes" as signals carrying information about other available nodes to facilitate the dynamic selection of a new parent node. In particular, a repeater or child sends out link probe packets and receives link probe packets from nearby APs, which provide information (*e.g.*, identity and signal strength of neighboring nodes) that is used to identify and select the next parent node. **Mesh Probes** Each Meraki access point sends out link probe packets (known as mesh probes) at different bit rates and varying sizes. Because these packets are sent as broadcast frames, no ACK frames are needed from receiving stations. Four different types of probes at different data rates are sent in a batch of 15 seconds on both (2.4 /5 GHz) bands. All access points listen to the mesh probes and depending on the number of mesh probes correctly received, come up with a link quality metric as shown in dashboard. Source: Wireless Mesh Networking Cisco Meraki and Catalyst CW Access Points, using the Meraki Mesh Algorithm, send out periodic discovery frames to discover mesh neighbors. **Mesh Convergence Time** The Meraki mesh algorithm is designed for stationary Meraki access points with variable links. The algorithm sends out periodic discovery frames of varying sizes on the 2.4 GHz and 5 GHz data radios to discover mesh neighbors. More discussion on neighbor discovery and gateway selection can be found in Wireless Mesh Networking. After one neighbor is selected as the gateway, new routes are passively identified on the same operating channel. A mesh link may take a couple minutes to establish and may not always be used for every data flow if a new route with better link metrics becomes available. Source: Mesh Deployment Guide

		If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of receiving a monitoring signal that includes information about available nodes from which the second node is selected in the same way by utilizing the Meraki Mesh Algorithm's neighbor discovery mechanism in which mesh access points exchange signals (*e.g.*, the first node receives neighbor response messages containing identifiers, etc.) to learn which nodes are available, and achieves the same result by obtaining information about candidate neighboring nodes so that the second node can be identified and selected as the new parent node based on that information. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's Accused Products exchange neighbor discovery signals or mesh probes carrying information about other available access points (such as identifiers and signal strengths of nearby nodes), so that each node receives data about which neighboring nodes are available for selection – an equivalent mechanism to the claimed monitoring signal with information about available nodes for choosing the second node. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
	Claim 12	
12.pre	12. A computer program product for controlling a wireless mesh network, the computer program product being embodied in a computer readable medium and comprising computer instructions for:	*See* 1.Pre-1.d above, the contents of which are incorporated herein by reference. Cisco Meraki and Catalyst CW Access Points comprise a computer program product (Cisco's proprietary Meraki Mesh Algorithms and related software, *e.g.*, cloud-based software or a universal Meraki or IOS XE OS image or other firmware) for controlling a wireless mesh network, the computer program product being embodied in a computer readable medium (*e.g.*, the storage device or memory on the Meraki and Catalyst CW Access Points) with all of the claimed computer instructions, as shown in claim 1 *supra*. Source: Meraki Firmware Directory MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
12.a	a) configuring the wireless mesh network	*See* 1.a above, the contents of which are incorporated herein by reference.

	with a first configuration, the wireless mesh network comprising a plurality of relay nodes, each of the relay nodes being capable of relaying traffic from a child node associated with the relay node, each of the relay nodes including a first radio interface configurable to relay traffic to a parent relay node and a second dual purpose radio interface configurable to both relay traffic from the child relay node and to communicate with a client node that is not a relay node, and at least one of the relay nodes is configured in the first configuration to relay traffic from a child relay node to a parent relay node;	
12.b	b) determining that a first one of the plurality of relay nodes in the wireless mesh network requires a channel selection change;	*See* 1.b above, the contents of which are incorporated herein by reference.
12.c	c) dynamically selecting a second one of the	*See* 1.c above, the contents of which are incorporated herein by reference.

	plurality of relay nodes in the wireless mesh network to form a parent-child relationship with the first relay node; and	
12.d	d) making a channel selection change; wherein prior to being dynamically selected, the second relay node was a sibling node of the first relay node; and making a channel selection includes making a channel Selection change to meet a performance requirement.	*See* 1.d above, the contents of which are incorporated herein by reference.
	Claim 13	
13	13. A computer program product as recited in claim 12, wherein the second node is selected to be a parent node of the first node.	*See* claims 2 and 12 above, the contents of which are incorporated herein by reference.
	Claim 14	
14	14. A computer program product as recited in claim 12, wherein the second node is selected to be a child node of the first node.	*See* claims 3 and 12 above, the contents of which are incorporated herein by reference.

	Claim 15	
15	15. A computer program product as recited in claim 12, wherein making a channel selection includes making a channel selection change to meet a latency requirement.	*See* claims 4 and 12 above, the contents of which are incorporated herein by reference.
	Claim 16	
16	16. A computer program product as recited in claim 12, wherein making a channel selection includes making a channel selection change to meet a throughput requirement.	*See* claims 5 and 12 above, the contents of which are incorporated herein by reference.
	Claim 17	
17	17. A computer program product as recited in claim 12, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to a performance requirement change.	*See* claims 6 and 12 above, the contents of which are incorporated herein by reference.

	Claim 18	
18	18. A computer program product as recited in claim 12, wherein determining that the first node requires a channel selection change includes determining that the channel selection change is required due to congestion.	*See* claims 7 and 12 above, the contents of which are incorporated herein by reference.
	Claim 19	
19	19. A computer program product as recited in claim 12, wherein the channel selection change is made based on network topology.	*See* claims 9 and 12 above, the contents of which are incorporated herein by reference.
	Claim 20	
20	20. A computer program product as recited in claim 12, the computer program product further comprising computer instructions for repeating b)-d).	*See* claims 10 and 12 above, the contents of which are incorporated herein by reference.
	Claim 21	
21	21. A computer program product as recited in claim 12, the computer program product further comprising computer instructions for	*See* claims 11 and 12 above, the contents of which are incorporated herein by reference.

	receiving a monitoring signal that includes information about available nodes from which the second node is selected.	

Exhibit 3 – U.S. Patent No. 7,894,385 (Cisco Wireless Mesh Networks with Catalyst and Aironet Access Points)

Cisco Systems, Inc. ("Cisco") has and continues to make, use, sell, offer for sale, and/or import the Cisco Wireless Mesh Network including but not limited to Cisco's Catalyst and Aironet Access Points, with related controllers and software (the "Cisco Accused Products," *see* Exhibit 11), and has infringed and infringes at least Claims 1-3 (the "Asserted Claims") of U.S. Patent No. 7,894,385 (the "'385 Patent").[1]

Cisco has and continues to infringe (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '385 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringes and has infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringes and has infringed by testing, setting up, and using Cisco Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Wireless Mesh Network in the United States.

Cisco indirectly infringes and has infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the '385 Patent and make and use the infringing Cisco Wireless Mesh Network. Specifically, Cisco's Wireless Mesh Access Points, Design and Deployment Guides, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco's website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-3 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Wireless Mesh Network, and the infringing configuration and operation of the controller, access points, and software that comprise the infringing Cisco Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the '385 Patent since 2009 when Cisco executives

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g.*, https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

met with MeshDynamics to discuss partnering opportunities and the application that would result in the ’385 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’385 Patent, or in the alternative was willfully blind to the ’385 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringes and has infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringes the ’385 Patent. Cisco Accused Products come enabled as a mesh access points and automatically connect to the controller to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particularly suited for Cisco’s Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’385 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed and continue to infringe the Asserted Claims of the ’385 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics’s preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts

presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul
1.pre. A method of operating a mesh network having a plurality of mesh nodes, comprising:	Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul ("URWB") comprises a method of operating a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points). Cisco's Wireless Mesh Network supports mobile access points and routers. Cisco offers access points designed for mobile use cases. These access points ("APs") can be used in various settings, including vehicles like trains, buses, or emergency response vehicles. Cisco's Ultra Reliable Wireless Backhaul ("URWB") works with both fixed infrastructure and mobile assets, *e.g.*, moving vehicles. URWB can be implemented in a mesh topology, where each vehicle equipped with a URWB device acts as a node in the network. The vehicle can be equipped with Catalyst access points, including the Cisco Catalyst Heavy Duty and Industrial Wireless products. **URWB supports different architectures:** • **Mobility:** Connectivity between radios deployed on vehicles and fixed structures or just between vehicles is generally used to connect Automated Guided Vehicles (AGVs), autonomous mobile robots (AMRs), trains, and other moving assets. • **Fixed infrastructure:** Connectivity between radios attached to fixed structures – poles, towers, buildings, etc. – is generally used to support wireless backhaul, physical surveillance, and more. The layout and data communication can be implemented through Point-To-Point (PTP), Point-To-Multipoint (PTMP), and mesh topologies.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/ultra-wireless-backhaul-so.html

URWB allows multiple Catalyst APs to wirelessly connect and form a mesh network. One AP is designated as the root or parent (connected to the wired LAN), and others become mobile mesh nodes (connected wirelessly).

Achieving low latency with Cisco Ultra-Reliable Wireless Backhaul

Cisco Ultra-Reliable Wireless Backhaul (Cisco URWB) is a wireless WAN backhaul technology derived from Wi-Fi and designed to serve mobile network environments such as trains, buses, subways, remotely controlled cranes, AGVs, and AMRs. It can also be deployed as a fixed infrastructure, providing backhaul where fiber can't go or is too expensive. Cisco URWB is especially well suited for connectivity between mobile networks (such as vehicles) and a fixed infrastructure, optimized to quickly move the mobile device's communications within a fixed wireless infrastructure before breaking. Cisco URWB networks can be designed for Layer 2 use cases (such as AGVs or AMRs in a distribution center) as well as Layer 3 use cases (such as subways), allowing network mobility through multiple locations a few miles apart.

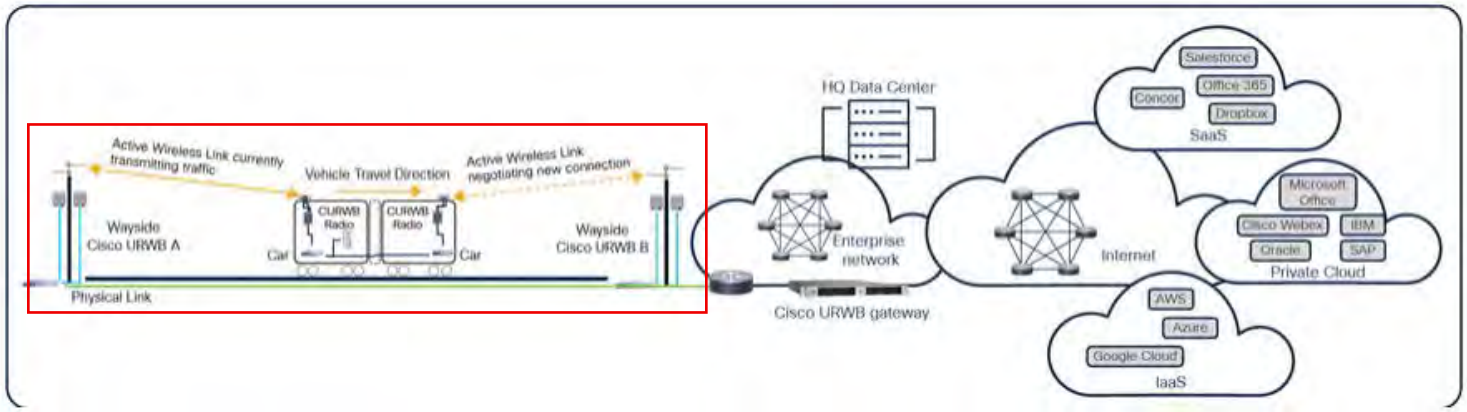
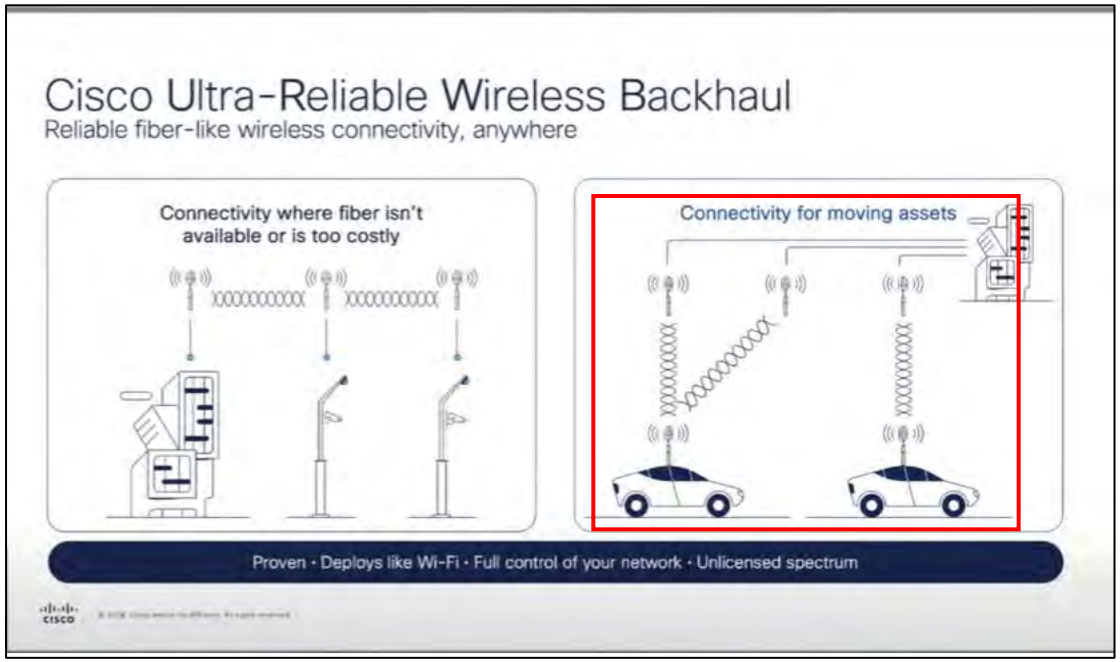


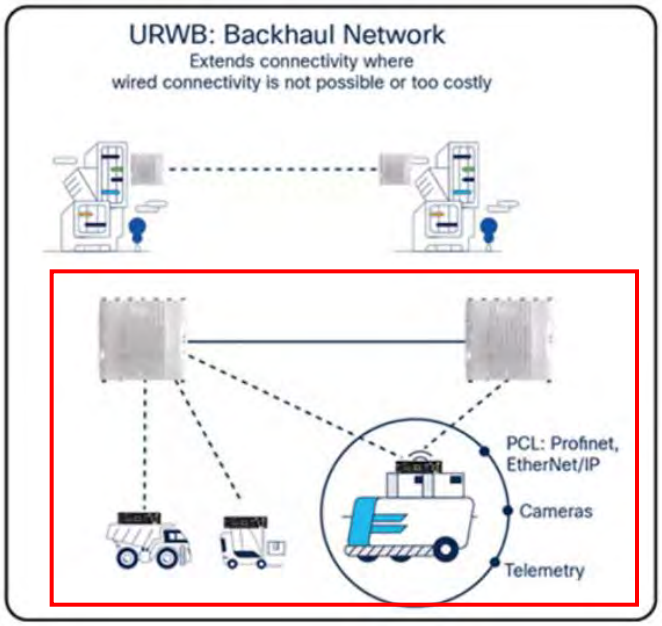
Figure 6.

Cisco Ultra Reliable Wireless Backhaul topology

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/arch-end-to-end-low-latency-wireless-networks.html



Source: https://www.youtube.com/watch?v=g9iJUmjy9Ug&list=PLinuRwpnsHacte6oSjeABKeotKLfPzNBN&t=86s

	 Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/elevating-wifi-intelligent-reliable-seamless-connectivity-so.html
1.a. for at least one mesh node of the mesh network, scanning a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node;	Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul ("URWB") comprises a method of operating a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty and Industrial Wireless Access Points) for at least one mesh node of the mesh network, scanning a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node. In URWB-enabled access points, such as the multi-radio IW9167E, there is a dedicated radio that can be used for scanning and monitoring the environment. Cisco's access point radio(s) can be a "dedicated" when being used for the purpose of scanning. This scanning radio can search for root or parent mesh nodes and optimize mesh path selection without disrupting client or backhaul traffic. URWB must know the boundary of a root or parent access point that it is currently connected to, and that the direction of a moving AP is soon going to lose that connection. This radio can scan ahead for alternative parent mesh points without dropping the current link.

Cisco Industrial Wireless Portfolio supporting URWB

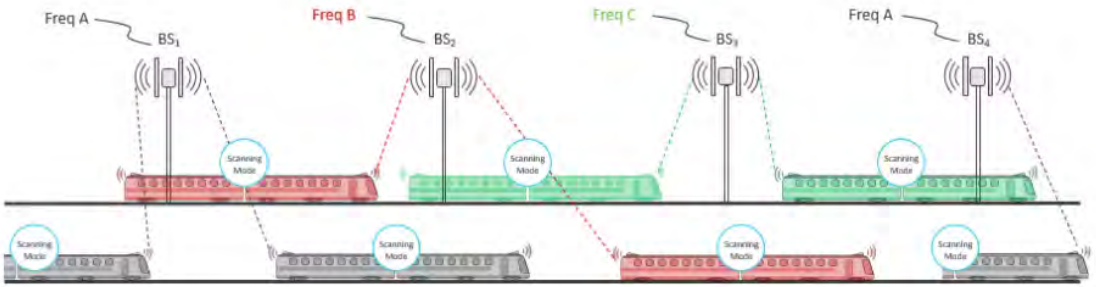
Table 1. Specifications summary for multi-radio Cisco Catalyst® Industrial Wireless products

	Catalyst IW9165E Rugged Access Point and Wireless Client	Catalyst IW9165D Heavy Duty Access Point	Catalyst IW9167I Heavy Duty Access Point	Catalyst IW9167E Heavy Duty Access Point	Catalyst IW9167E-HZ Heavy Duty Access Point
URWB mode	Standalone, Integrated with Wi-Fi technology	Standalone, Integrated with Wi-Fi technology	Integrated with Wi-Fi technology	Standalone, Integrated with Wi-Fi technology	Standalone, Integrated with Wi-Fi technology
Technology	2x2 dual radio up to 160 MHz	2x2 dual radio up to 160 MHz	4x4 tri-radio up to 160 MHz	4x4 tri-radio up to 160 MHz	4x4 tri-radio up to 160 MHz

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/ultra-wireless-backhaul-so.html

Overview of Fluidity frequency scan

Fluidity frequency scan is designed for scenarios involving high-density mobility environments using multiple frequencies on the infrastructure side to reduce self-interference and improve wireless channel capacity and performance. It helps to ensure continuous and seamless connectivity as these devices move across different network areas, facilitating smooth transitions between APs. When the current signal strength is weak, the device starts searching for a frequency with a better signal to maintain a stable connection. This feature is applicable in all high-density mobility environments, such as train-to-ground, mining, and port terminals.



For effective frequency scanning, a mobile device typically requires at least two radios:

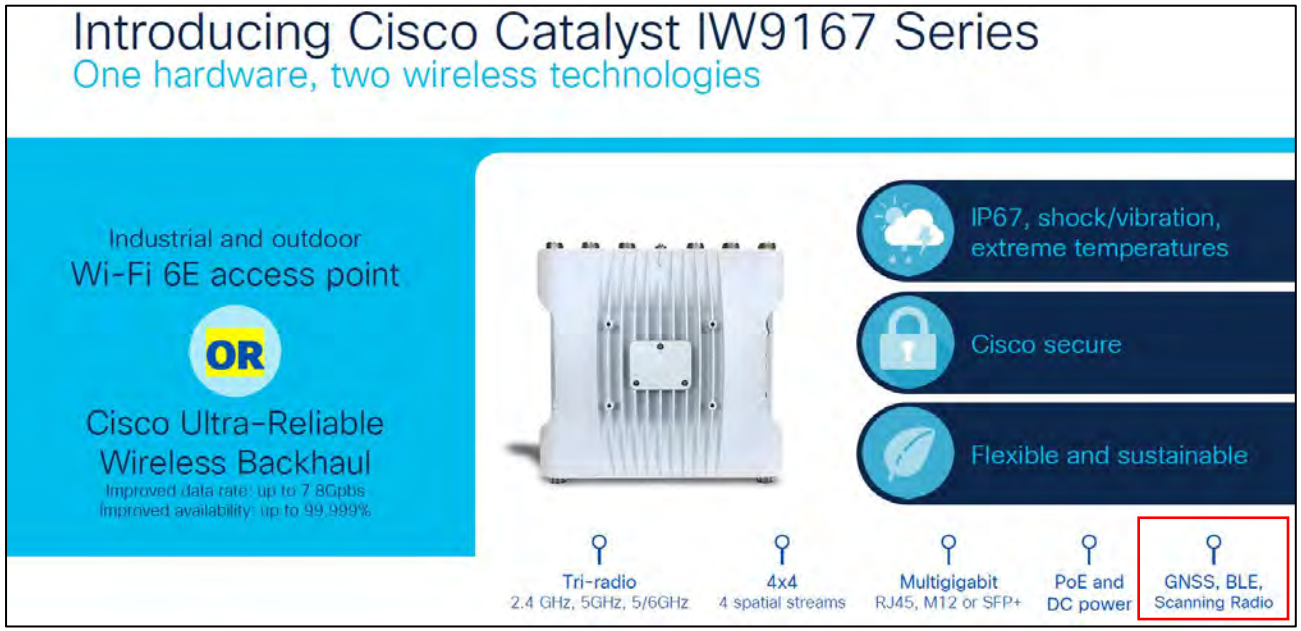
- Radio 1: Maintains the current connection with the network.
- Radio 2: Performs scanning for better available frequencies without disrupting the existing connection on Radio 1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/outdoor_industrial/iw9167/software/configuration/m-cisco-catalyst-iw9167e-iw9165e-and-iw9165dh-heavy-duty-access-point-cisco-ultra-reliable-wireless-backhaul-software-configuration-guide/m_auto-or-frequency-scan.html

Scanning radio	Dedicated radio for monitoring air space to perform advanced RF spectrum analysis and deliver features such as Cisco CleanAir®, Wireless Intrusion Prevention System (WIPS), etc.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-iw9167-series/cat-iw9167e-heavy-duty-ap-ds.html

The GNSS/BLE scanning radio helps to determine a new potential parent mesh node for connecting by enabling all mesh nodes to share a precise clock, allowing coordinated time-slot communication, and improves how APs connect and operate in mesh. This radio enables precise AP location and time sync for Time Division Multiple Access ("TDMA") for URWB and helps mesh APs coordinate in time and space, which is critical for scheduling and optimizing the mesh.



Source: https://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2022/pdf/BRKEWN-2832.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst Heavy Duty and Industrial Wireless Access Points perform the same function of determining a new potential parent or root mesh node for connecting with said at least one mesh node in the same way by using radios to regularly scan for and track potential new parents or roots in an RF environment and achieves the same result by determining the best parent or root mesh node for optimal network performance and to maintain connectivity while in motion. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's URWB networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic and to maintain connectivity when in motion by scanning for new parents.

	A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.b. wherein said at least one mesh node includes, in addition to said scanning radio, at least two relay radios in each mesh element and wherein said scanning radio and said at least two relay radios operate on different non-interfering channels;	Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul ("URWB") comprises a method of operating a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty and Industrial Wireless Access Points) wherein said at least one mesh node includes, in addition to said scanning radio, at least two relay nodes in each mesh element and wherein said scanning radio and said at least two relay nodes operate on different non-interfering channels. URWB also uses dual backhaul radios operating on different non-interfering channels (typically 5 GHz or 4.9 GHz bands) from the scanning radio(s), with MIMO (multiple-input, multiple-output) antennas. There is also a 2.4 GHz radio. Each of the 2.4GHz and 5GHz (or 6GHz) bands have several different non-interfering channels for radio operation. For instance, channels 1, 6 and 11 on the 2.4 GHz band do not interfere. There are 24 non-overlapping channels for 5-GHz, offering more flexibility and less congestion. The APs use at least the 2.4 GHz, 5 GHz, and/or 6 GHz bands to provide comprehensive monitoring and use of the wireless spectrum for various purposes, at different channels. **Cisco Industrial Wireless Portfolio supporting URWB** **Table 1. Specifications summary for multi-radio Cisco Catalyst® Industrial Wireless products** |	Catalyst IW9165E Rugged Access Point and Wireless Client	Catalyst IW9165D Heavy Duty Access Point	Catalyst IW9167I Heavy Duty Access Point	Catalyst IW9167E Heavy Duty Access Point	Catalyst IW9167E-HZ Heavy Duty Access Point
---	---	---	---	---	---
URWB mode	Standalone, Integrated with Wi-Fi technology	Standalone, Integrated with Wi-Fi technology	Integrated with Wi-Fi technology	Standalone, Integrated with Wi-Fi technology	Standalone, Integrated with Wi-Fi technology
Technology	2x2 dual radio up to 160 MHz	2x2 dual radio up to 160 MHz	4x4 tri-radio up to 160 MHz	4x4 tri-radio up to 160 MHz	4x4 tri-radio up to 160 MHz

Frequency	4.9 GHz, 5.1 to 6.9 GHz*	4.9 GHz, 5.1 to 6.9 GHz*	2.4 GHz, 4.9 GHz, 5.1 to 6.9 GHz*	2.4 GHz, 4.9 GHz, 5.1 to 6.9 GHz*	2.4 GHz, 4.9 GHz, 5.1 to 6.9 GHz*

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/ultra-wireless-backhaul-so.html

Overview of Cisco Catalyst IW9167E and IW9165 Access Points

Overview of Cisco Catalyst IW9167E

The Catalyst IW9167E access point provides reliable wireless connectivity for mission-critical applications in a state-of-the art platform to deliver a network that is more reliable and secure, with higher throughput, more capacity, and less device interference. The Catalyst IW9167E is Cisco's first outdoor Wi-Fi 6E ready Access Point supporting tri-radio and tri-band (2.4/5/6 GHz bands). The Catalyst IW9167E can operate in Wi-Fi (control and provisioning of wireless access points (CAPWAP)) mode or Ultra-Reliable Wireless Backhaul (URWB) mode and URWB software on Catalyst IW9167E designed to support the Cisco style parser.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/outdoor_industrial/iw9167/software/configuration/m-cisco-catalyst-iw9167e-iw9165e-and-iw9165dh-heavy-duty-access-point-cisco-ultra-reliable-wireless-backhaul-software-configuration-guide/chapter-1.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst Heavy Duty and Industrial Wireless Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance and to maintain connectivity when in motion by scanning for new parents, while maintaining a separate backhaul. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks using URWB deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents to maintain connectivity and while maintaining a separate

	backhaul. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.c. wherein said at least one mesh node is moving sufficiently rapidly that it may lose connectivity with its current parent mesh node, and wherein said dedicated scanning radio is utilized for discovery of potential new parent nodes and sampling of discovered potential new parent nodes is performed by an uplink relay radio of said at least one mesh node;	Cisco's Wireless Mesh Network comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty and Industrial Wireless Access Points) wherein said at least one mesh node is moving sufficiently rapidly that it may lose connectivity with its current mesh node, and wherein said dedicated scanning radio is utilized for discovery of potential new parent and sampling of discovered potential new parent nodes is performed by an uplink relay radio of said at least one mesh node. Cisco's Wireless Mesh Network supports mobile access points and routers. Cisco offers access points designed for mobile use cases. These APs can be used in various settings, including vehicles like trains, buses, or emergency response vehicles. Cisco's Ultra Reliable Wireless Backhaul ("URWB") works with both fixed infrastructure and mobile assets, *e.g.* moving vehicles. URWB can be implemented in a mesh topology, where each vehicle equipped with a URWB device acts as a node in the network ("at least one mesh node is moving sufficiently rapidly"). The vehicle can be equipped with Catalyst access points, including the Cisco Catalyst Heavy Duty and Industrial Wireless products. **URWB supports different architectures:** • **Mobility:** Connectivity between radios deployed on vehicles and fixed structures or just between vehicles is generally used to connect Automated Guided Vehicles (AGVs), autonomous mobile robots (AMRs), trains, and other moving assets. • **Fixed infrastructure:** Connectivity between radios attached to fixed structures – poles, towers, buildings, etc. – is generally used to support wireless backhaul, physical surveillance, and more. The layout and data communication can be implemented through Point-To-Point (PTP), Point-To-Multipoint (PTMP), and mesh topologies. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/ultra-wireless-backhaul-so.html As discussed *supra* claim element 1.a, Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points comprise a "scanning radio is utilized for discovery of potential new parent nodes" and to monitor the environment. URWB delivers a "make-before-break" handover, enabling moving vehicles to negotiate and establish a reliable parent connection ("sampling" using its uplink radio) with the next access point along their path before losing connectivity to

the current one. The availability of at least one AP is ensured ("sampling") before a handoff is made using an uplink radio. *See also infra* claim 1.d.

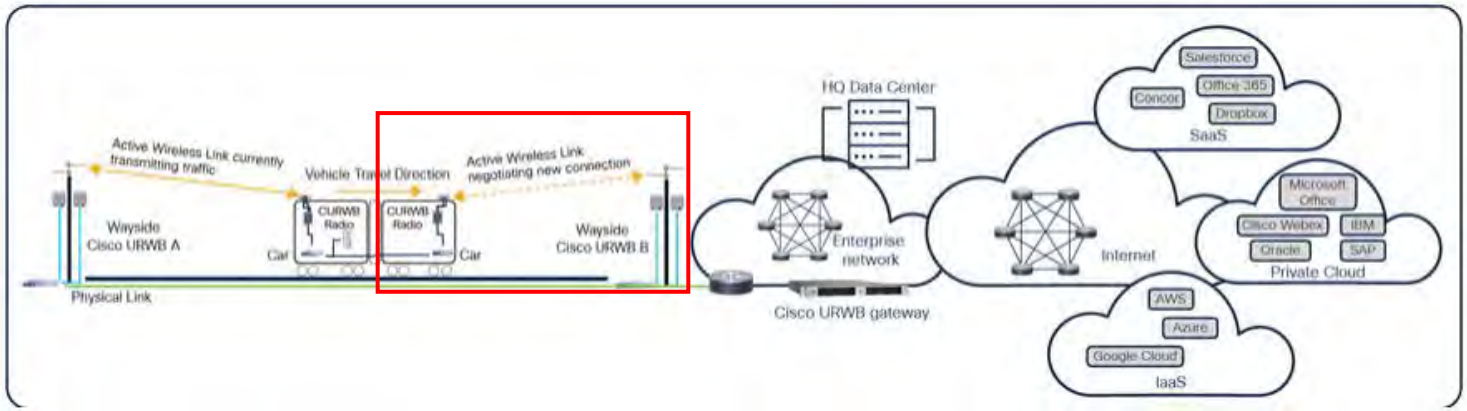
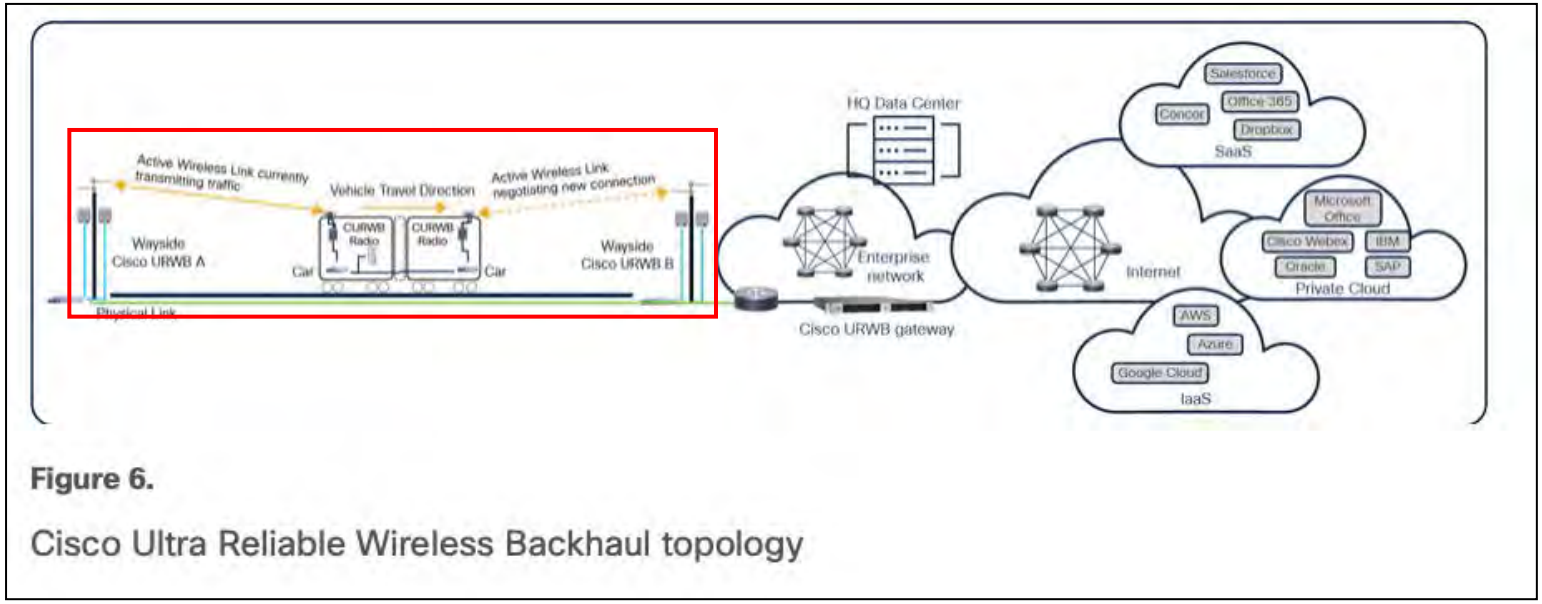


Figure 6.

Cisco Ultra Reliable Wireless Backhaul topology

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/arch-end-to-end-low-latency-wireless-networks.html; *see also* https://www.cisco.com/c/dam/en/us/td/docs/wireless/ultra-reliable-wireless-backhaul/device-software/cisco-urwb-fluidity-specs-usermanual.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst Heavy Duty and Industrial Wireless Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance and to maintain connectivity when in motion by scanning for new parents, while maintaining a separate backhaul. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks using URWB deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents to maintain connectivity when in motion and while maintaining a separate backhaul. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.d. wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node; and	Cisco's Wireless Mesh Network comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty and Industrial Wireless Access Points) wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node. URWB delivers a "make-before-break" handover, enabling moving vehicles to negotiate and establish a reliable connection ("sampling") with the next access point along their path before losing connectivity to the current one. When the mesh node uplink is negotiating, on information and belief, the parent node or current node holds the packets ("buffered") so no packets are lost during the handoff. As discussed *infra* claim 1.e, receiving and transmission times are coordinated and on information and belief, packets are buffered at the APs during times not slotted for receiving and transmission.  **Figure 6.** Cisco Ultra Reliable Wireless Backhaul topology Source: https://www.cisco.com/c/en/us/products/collateral/wireless/arch-end-to-end-low-latency-wireless-networks.html

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst Heavy Duty and Industrial Wireless Access Points perform the same function of sampling or negotiating with other nodes to determine a new potential parent mesh node for maintaining a network connection in the same way by using radios to regularly scan for potential new parents in an RF environment, while maintaining the traffic flow to the old parent before transfer to a new potential parent, and achieves the same result of ensuring the seamless transfer of data flow ("make-before-break" handover) and a good user experience without degrading performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks with URWB must solicit and learn all of the available APs to maintain a connection while moving. The "make-before-break" handover, like buffering, ensures that no packets are lost in the transition to a new RAP when moving. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.e. wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.	Cisco's Wireless Mesh Network comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty and Industrial Wireless Access Points) wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner. The current parent or fixed root access point (RAP) hosts URWB functions. The RAP contains intelligence to monitor moving node signal quality, advertise its availability and quality metrics (*e.g.*, SNR, latency), accept or reject roaming requests, and recalculate schedules and time slots dynamically. Mesh links are reserved, scheduled and optimized in real time to ensure stability even as one node moves. This is how URWB can deliver zero loss and seamless handoffs ("make-before-break" handover) in complex mobile use cases. Cisco®'s Ultra-Reliable Wireless Backhaul (URWB) technology addresses these complex fixed and mobile use cases, delivering advantages across the network: • Provides multigigabit high-throughput performance, connects moving assets, and links locations that need temporary connectivity. • Delivers ultra-low latency, zero loss, and seamless handoffs. • Extends 802.11 technology and uses unlicensed spectrum, reducing costs and complexity and making it easy to deploy pretty much anywhere. • Since URWB is based on the 802.11 standard, IT and OT teams find it easy to understand and deploy.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/ultra-wireless-backhaul-so.html

On information and belief, the RAPs utilizing URWB use a Time Division Multiple Access ("TDMA") or TDMA-like mechanism (using a similar Cisco proprietary protocol, such as "Fluidity" or Fluidmax mode, or Cisco's Multiprotocol Label Switching (MPLS) or Multipath Operations (MPO)). TDMA is a wireless scheduling method where the available airtime is divided into fixed time slots, and each device is assigned a slot for transmission ("sampling times are coordinated among multiple mesh nodes having a common current parent node"). The RAP acts like a TDMA scheduler and coordinates a unique time slot for each node for transmitting or receiving. The RAP ensures all nodes are time-synced (like a clock, "sends tokens to each of its children in a round-robin manner"), so they transmit only in their slots. All participating nodes—including mobile and fixed APs—are time-synchronized. This allows both old and new access points to coordinate when the mobile node will talk/listen without interference.

Configuring URWB Radio Mode

The wireless interfaces are configured to operate in a specific mode, or you can disable it. Once you configure the Radio mode, the device starts working as a Fluidity or Fixed infrastructure.

The following table shows the configuration of Radio mode on the device:

Table 1. Radio Mode Configuration

Radio Role	Radio Mode	Description
Fixed Infrastructure	Fixed Fluidmax primary Fluidmax secondary	P2P mode (point to point) P2MP (point to multipoint) mode (Fluidmax) and P2MP P2MP mode (Fluidmax) and P2MP
Mobility AP	Fluidity	Mobility mode
Mobility Client	Fluidity	Mobility mode

Source: https://www.cisco.com/c/en/us/td/docs/wireless/outdoor_industrial/iw9167/software/configuration/m-cisco-catalyst-iw9167e-iw9165e-and-iw9165dh-heavy-duty-access-point-cisco-ultra-reliable-wireless-backhaul-software-configuration-guide/m-configuring-curwb-radio-mode.html; *see also* https://community.cisco.com/t5/wireless-mobility-

knowledge-base/time-division-multiple-access-tdma/ta-p/3112403; https://www.cisco.com/c/dam/en/us/td/docs/wireless/ultra-reliable-wireless-backhaul/outdoor-devices/cisco-urwb-fm3200base-usermanual.pdf (referencing Fluidmax TDMA mode)

URWB ensures that a moving vehicle establishes a reliable connection with the next access point before losing connectivity to the current one. On information and belief, this transition relies on managing connection handoffs efficiently and coordination by the parent node, at least conceptually similar to how TDMA would allocate time slots to different users for uninterrupted communication. URWB's predictable communication, redundant paths and MPO, and dedicated time for packet transmission are TDMA-like approaches.

- URWB uses a specialized implementation of the Multiprotocol Label Switching (MPLS)-over-wireless links, created to overcome the limits of standard wireless protocols when transmitting any IP-compatible traffic with very low latency in a mobility context.
- URWB delivers a "make-before-break" handover, enabling moving vehicles to establish a reliable connection with the next access point along their path before losing connectivity to the current one.
- With Multipath Operations (MPO) technology, URWB can deliver uninterrupted connectivity to fast-moving devices by sending high-priority packets via redundant paths on uncorrelated frequencies at the same time to multiple access points. It can duplicate protected traffic up to eight times, exploiting time, spatial, and frequency diversity. This functionality, combined with cutting edge hardware capability, can further reduce latency and improve reliability, addressing both interference and hardware failures.

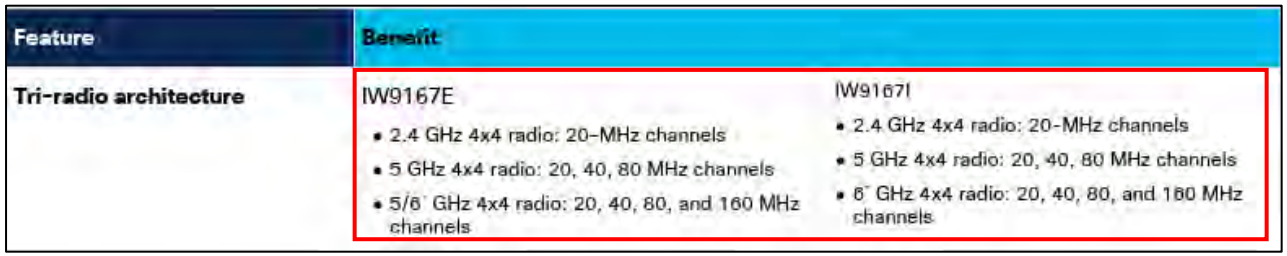
Source: https://www.cisco.com/c/en/us/products/collateral/wireless/ultra-reliable-wireless-backhaul/ultra-wireless-backhaul-so.html

Furthermore, URWB discloses this element because Cisco's URWB has the ability to switch to a plurality of operating modes (*e.g.*, fixed and fluidity) that requires the root node to coordinate when nodes should sample the presence of other candidate root nodes. For example, on information and belief, the root node would inform (*i.e.*, send a "token") to a node operating in fixed mode to perform sampling at a relatively slow rate and the root node would inform (*i.e.*, send a "token") to a node operating in fluidity mode to sample at a higher rate (because of the higher mobility use case, such as a train). This would be done when initially configuring a node for its mode of operation and, on information and belief, on an ongoing basis. The root node would send this information to the nodes (fixed and fluidity) separately (*i.e.*, "round-robin manner"). Also, with TDMA, the specific time slot made available for a particular node to transmit can be understood to be a "token," and the rotating availability of time slots for different nodes would be considered being done in a "round-robin manner."

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Heavy Duty and Industrial Wireless Access Points perform the same function of seamless parent node switching in a mobile mesh network in the same way by coordinating transmissions and achieves the same result of ensuring the seamless transfer of data flow ("make-before-break" handover) and a good user experience without degrading performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial. URWB time slots or "make-before-break" handover functionally ensure orderly, collision-free and uninterrupted communication. The time slot or transition serves the same function of preventing collisions and giving each node a clear opportunity and permission to transmit. The time slot or handover is also assigned and centrally controlled by the RAP, such that just like a token, only the node whose time slot or transition is active can transmit. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 2	**Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul**
2.pre. A mesh network comprising: a plurality of mesh nodes;	Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul ("URWB") comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points). *See supra* discussion for claim element 1.pre, the contents of which are incorporated herein by reference.
2.a. wherein at least one mesh node of the mesh network is configured to scan a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node;	*See supra* discussion for claim element 1.a., the contents of which are incorporated herein by reference.
2.b. wherein said at least one mesh node includes, in addition to the scanning radio, at least two relay radios in each	*See supra* discussion for claim element 1.b., the contents of which are incorporated herein by reference.

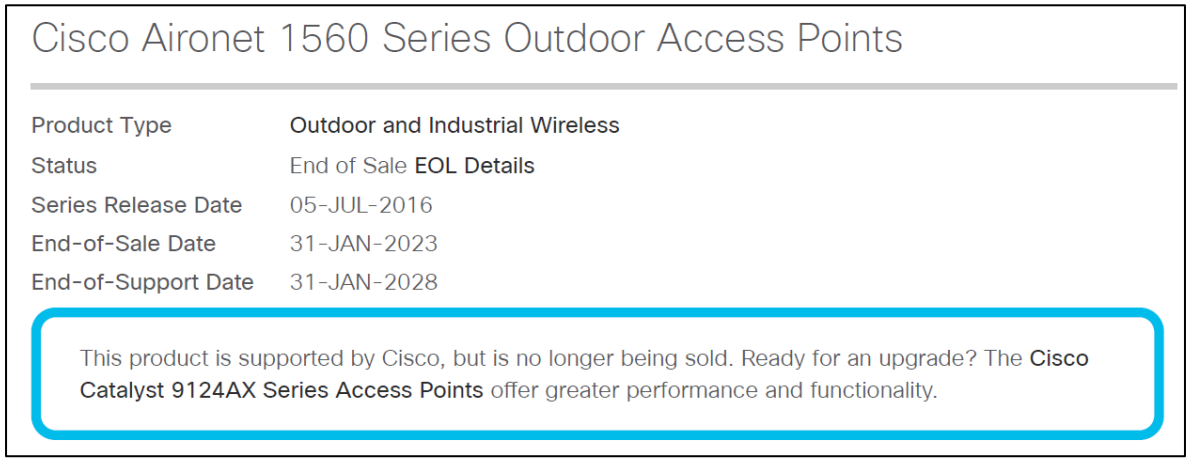
mesh element and wherein said scanning radio and said at least two relay radios operate on different non-interfering channels;	
2.c. wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node;	*See supra* discussion for claim element 1.d., the contents of which are incorporated herein by reference.
2.d. wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.	*See supra* discussion for claim element 1.e., the contents of which are incorporated herein by reference.
Claim 3	**Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul**
3.pre. A mesh network comprising: a plurality of mesh nodes;	Cisco's Wireless Mesh Network with Ultra Reliable Wireless Backhaul ("URWB") comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points). *See supra* discussion for claim element 1.pre, the contents of which are incorporated herein by reference.

3.a. wherein each node within said plurality of nodes comprises at least three radios further including	Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points in the Cisco Wireless Mesh Network with URWB comprise a mesh network wherein each node within said plurality of nodes comprises at least three radios. As shown above, the Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points are tri-radio, or include at least dual radios and a scanning radio. Each of the 2.4GHz and 5GHz (or 6GHz) bands have several different non-interfering channels for radio operation. For instance, channels 1, 6 and 11 on the 2.4 GHz band do not interfere. There are 24 non-overlapping channels for 5GHz, offering more flexibility and less congestion. The APs use at least the 2.4 GHz, 5 GHz, and/or 6 GHz bands to provide comprehensive monitoring and use of the wireless spectrum for various purposes, at different channels. *See supra* discussion for claim elements 1.a and 1.b., the contents of which are incorporated herein by reference. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst Heavy Duty Access Points and Industrial Wireless Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
3.b. a first relay radio operating on a first RF channel at a first point in time and dedicated to uplink connections to a single current parent node; a second relay radio operating on a second RF channel at a first point in	Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points in the Cisco Wireless Mesh Network comprise a first relay radio operating on a first RF channel at a first point in time and dedicated to uplink connections to a single current parent node and a second relay radio operating a second RF channel at a first point in time dedicated to downlink connections to zero or more child nodes. As shown above, and in Cisco's documentation, the Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points have dedicated backhaul radios (uplink to the parent and network and downlink to zero or more child nodes) operating on different channels. Cisco's access point radio(s) can be a "dedicated" when being used for the purpose of uplink to the parent node or downlink to zero or more child nodes. Cisco Catalyst Heavy Duty Access Points and Industrial Wireless Access Points are tri-radio, or include at least dual radios and a scanning radio. Each of the 2.4GHz and 5GHz (or 6GHz) bands have several different non-interfering channels for radio operation. For instance,

time dedicated to downlink connections to zero or more child nodes; and	channels 1, 6 and 11 on the 2.4 GHz band do not interfere. There are 24 non-overlapping channels for 5GHz, offering more flexibility and less congestion.  Source: https://www.industrialnetworking.com/pdf/Cisco-Catalyst-IW9167E.pdf *See supra* discussion for claim elements 1.a and 1.b, the contents of which are incorporated herein by reference. If these claim elements are not literally infringed, they are infringed under the doctrine of equivalents. Cisco's Catalyst Heavy Duty Access Points and Industrial Wireless Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed elements are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

3.c. a dedicated scanning radio operating on a third RF channel at a first point in time and configured to scan a Radio Frequency (RF) environment to discover new potential parent mesh nodes for connection with said first relay radio;	*See supra* discussion for claim element 1.a, the contents of which are incorporated herein by reference.
3.d. said first, second, and third RF channels are different from each other	*See supra* discussion for claim elements 1.a, 1.b, 3.a, and 3.b, the contents of which are incorporated herein by reference.
3.e. wherein after said dedicated scanning radio discovers a potential new parent node, said first relay radio samples the RF link to said potential new parent node using its uplink radio and concurrent with said sampling, packets to be sent to the mesh node from its current parent node are buffered by the current parent node, and packets to be sent from the mesh node to its current parent node are buffered by the mesh node wherein sampling times are coordinated among multiple mesh nodes having a common	*See supra* discussion for claim elements 1.a, 1.c, 1.d. and 1.e. the contents of which are incorporated herein by reference.

current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.	
Claim 2	**Cisco's Wireless Mesh Network with Catalyst and Aironet Access Points**
2.pre. A mesh network comprising: a plurality of mesh nodes;	Cisco's Wireless Mesh Network (a component of the Cisco Unified Wireless Network solution) comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Series Access Points). Cisco Catalyst 9100 and Aironet Access Points support Wi-Fi capabilities. The Catalyst series is a follow-up from the Cisco Aironet series access points, and the Aironet and Catalyst Access Points contain the same mesh networking capabilities, as shown below in Cisco's documentation and guides that apply to both product lines: Q. Which Cisco access points support Wi-Fi 6? A. Cisco Catalyst 9100 enterprise-class access points support Wi-Fi 6 capabilities, extending the power of intent-based networking with hardware and software innovations. The Cisco Catalyst 9100 APs offer resiliency and security while enabling intelligence at the network edge. Q. What is the portfolio transition from the Cisco Aironet Series to the Cisco Catalyst 9100 platform? A. The Cisco Catalyst 9100 Access Points consist of the Cisco Catalyst 9105, 9115, 9120, 9130, 9136, 9162, 9164, 9166 and 9166D1, and outdoor 9163E Series. These access points are the follow-ups to the Cisco Aironet 1815, 1850, 2800, and 3800 Series Access Points, respectively. The Cisco Catalyst 9100 Access Points come equipped with Wi-Fi 6/6E capabilities, a better industrial design, and improved RF performance, and deliver reliability, security, and intelligence at scale. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html



Cisco Aironet 1560 Series Outdoor Access Points

Product Type	Outdoor and Industrial Wireless
Status	End of Sale EOL Details
Series Release Date	05-JUL-2016
End-of-Sale Date	31-JAN-2023
End-of-Support Date	31-JAN-2028

This product is supported by Cisco, but is no longer being sold. Ready for an upgrade? The Cisco Catalyst 9124AX Series Access Points offer greater performance and functionality.

Source: https://www.cisco.com/c/en/us/support/wireless/aironet-1560-series/series.html

Mesh deployments are supported on the Cisco Catalyst 9100 Access Points:

Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points?

A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

In addition to the Wave 1 and Wave 2 access points, mesh deployments are supported on the Cisco Catalyst 9100 Access Points, including, *e.g.*, Catalyst 9124 Access Points, Catalyst 9130 Access Points and "AX" access points:

Families of mesh APs

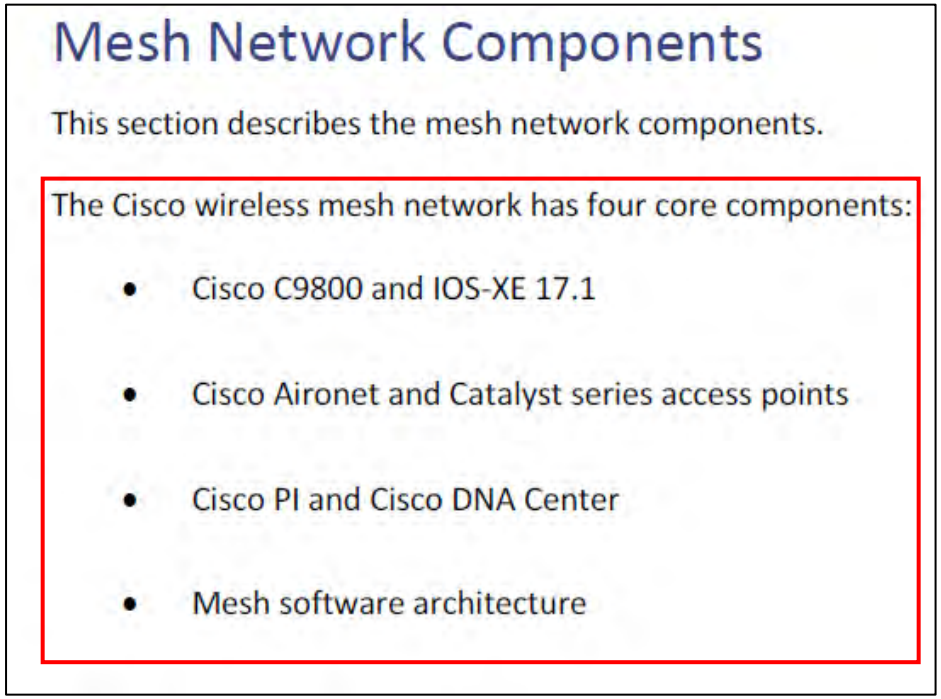
The two families of mesh APs available are the Cisco Wave 1 APs and the Cisco Wave 2 APs + the AX APs. These two families of mesh APs have different capabilities and features, and they are not fully compatible with each other. The Cisco Wave 1 APs do not support the same feature set as the Cisco Wave 2 + AX APs.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg/mesh-access-points.html

> From this release, mesh support is included in the Cisco Catalyst 9130AX Series Access Points. All traditional capabilities of mesh are included in the Cisco Catalyst 9130AX Series APs operating in Cisco IOS XE Dublin 17.12.1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html

Cisco Catalyst 9100 Access Points (with the Cisco Aironet series access points) can be deployed in a Cisco Wireless Mesh Network with a controller and software, such as the Catalyst 9800 Controller and Cisco DNA and Prime Infrastructure. Cisco Aironet and Catalyst series access points are both identified as components of a Cisco wireless mesh network:

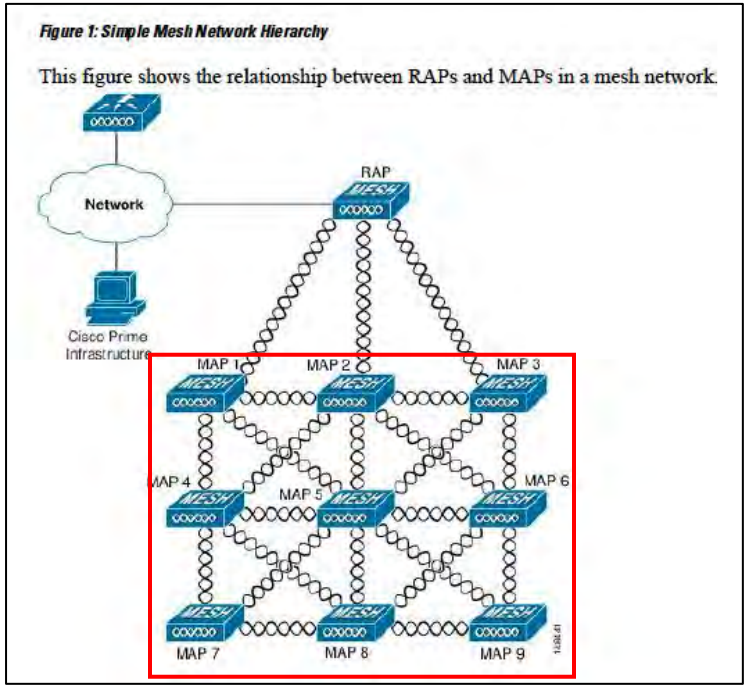


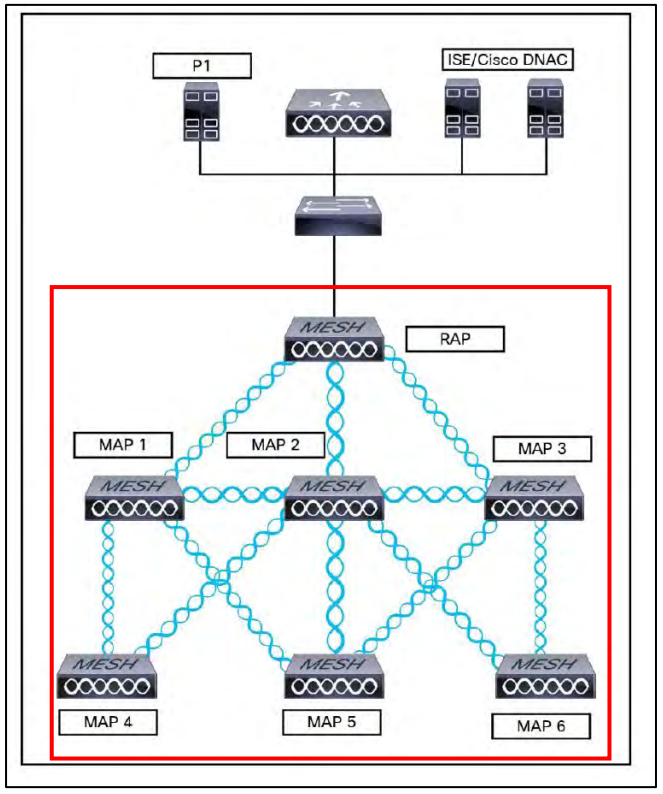
Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network.

- Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series.
- Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

In Cisco's Wireless Mesh Network, Cisco Catalyst 9100 and Aironet Access Points ("plurality of mesh nodes" as shown in red below) are organized in a mesh network.





	Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
2.a. wherein at least one mesh node of the mesh network is configured to scan a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise of a method for operating a mesh network having a plurality of mesh nodes comprising for at least one mesh node of the mesh network, scanning a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node. Cisco's Catalyst 9100 and Aironet Access Points support multiple radios, including one or more radios for client access, one or more radios for wireless backhaul, and a scanning radio. Cisco's access point radio(s) can be a "dedicated scanning radio" when being used for the purpose of scanning to discover neighbors or other mesh nodes ("new potential parent mesh node") but may subsequently be used for other purposes such as backhaul or sampling. With multiple radios (*e.g.*, tri-radio) one radio can be dedicated to the purpose of scanning. For example, Cisco Catalyst 9130AX Wi-Fi 6 access point is tri-radio (2.4 GHz + 5 GHz with dual-5GHz using Flexible Radio Assignment) that supports mesh. One radio can be used for client access, the others used flexibly for backhaul or scanning at different channels. **Flexible Radio Assignment (FRA) with tri-radio mode** FRA allows the access points to intelligently determine the operating mode of serving radios based on the RF environment and traffic demands. The access points can operate in the following modes: • Dual-radio mode: One 8x8 5 GHz and one 4x4 2.4 GHz. One radio will serve clients on the 5-GHz band, while the other serves clients on the 2.4-GHz band. • Tri-radio mode: * Dual 4x4 5 GHz and one 4x4 2.4 GHz. With two 4x4 5-GHz and one 4x4 2.4-GHz radios (tri-radio) inside the access point, client device capacity can be increased on demand. The access point's default mode is dual radio with 8x8 5 GHz and 4x4 2.4 GHz. It has the ability to split the 8x8 radio into two separate 4x4 5-GHz radios through software, thereby enabling the benefits of FRA while allowing the 2.4-GHz radio to remain active. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html

Cisco Access Points with Tri-Radio

This topic describes the Tri-Radio feature for Cisco Access Points (APs).

Access Points with three radios are designed for high density environments. The APs by default run one dedicated 2.4-GHz 4x4 mode radio and one 5-GHz 8x8 mode radio. In the default mode, the radios are managed by the Flexible Radio Assignment (FRA), and the Dual Radio Mode is in the disabled state indicating that the radios have either been assigned as client serving 8x8 radio or have not yet been evaluated by FRA.

When you enable the dual radio mode setting, the 8x8 radio is split to two independent 5-GHz 4x4 radios. In this mode, slot 1 and slot 2 are active independent 4x4 radio interfaces. They can serve different user groups with different assigned channels.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/cisco_access_points_with_tri_radio.pdf

There is also a Cisco RF ASIC (Software-Defined Radio), as described in Cisco's datasheet that functions a dedicated monitoring radio, enabling spectrum scanning off-channel to find other neighboring mesh nodes and to facilitate parent discovery ("dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node").

Cisco RF ASIC	The Cisco RF Application-Specific Integrated Circuit (ASIC) is a fully integrated Software-Defined Radio (SDR) that can perform advanced RF spectrum analysis and delivers features such as Cisco CleanAir, Wireless Intrusion Prevention System (WIPS), FastLocate, and Dynamic Frequency Selection (DFS) detection.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html

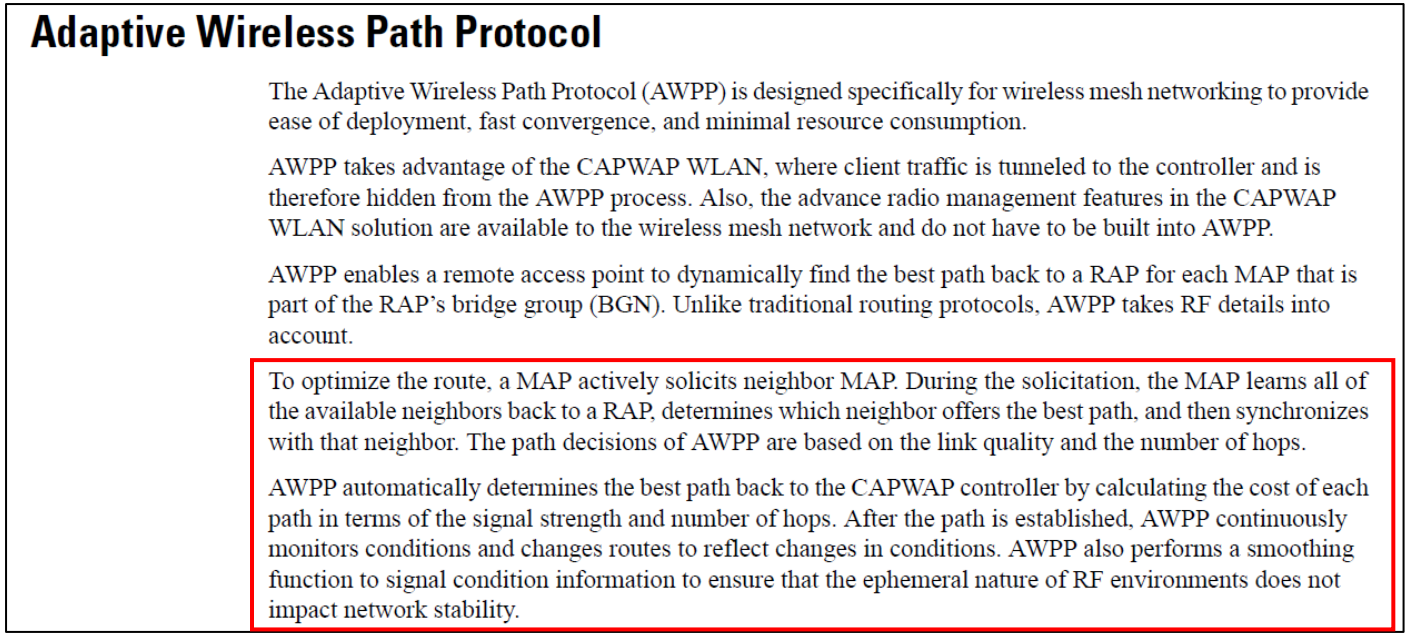
- Off-Channel Radio Resource Management (RRM) Scanning and Neighbor Discovery Protocol (NDP) Transmission: RF ASIC performs off-channel scanning to analyze channel utilization. In typical RRM scenarios, when the radio is off channel, it cannot receive from or transmit to clients on the channel it was serving. In a busy network, this scanning is often interrupted periodically to prevent potential intermittent delays or packet loss on time sensitive applications. The RF ASIC radio offloads this function by continuously monitoring conditions and enhancing the performance of the access points' client-serving radios.

Source: https://blogs.cisco.com/networking/cisco-rf-asic-named-best-enterprise-wi-fi-network-technology-by-the-wireless-broadband-alliance

In addition, while 5 GHz is often used for the primary backhaul (communication between access points), 2.4 GHz can also be used for backhaul and scanning, for example, when Universal Backhaul Access (UBA) is enabled.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-6/b_mesh_86/m_mesh-network-components.html

Catalyst 9100 and Aironet Access Points in Cisco Wireless Mesh Networks use Adaptive Wireless Path Protocol (AWPP) to determine the best path through the mesh access points (MAPs) to a Root Access Point (RAP). This includes selection of a parent mesh node and establishing a parent-child relationship. AWPP takes RF details into account. During solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path and then synchronizes with that neighbor. On information and belief, the 5GHz backhaul radio, for example, may also be used by AWPP for scanning to find a new potential parent mesh node.



Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf at 12; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

AWPP follows a process in selecting parents for a MAP with a radio backhaul, including generating a list of channels with neighbors by passive scanning in the *scan* state, which is a subset of all backhaul channels. The channels with the neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the best neighbor. Parent maintenance and optimization occurs in the maintain state.

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The Mesh Background Scanning feature allows for faster mesh convergence and sampling potential new parent nodes. The child MAP periodically scans its neighbors across all off-channels to connect with a better potential parent.

This new Mesh Background Scanning and Auto parent selection will further improve convergence times and parent selection reliability and stability - a MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time .

A child MAP will switch between optimal parents based on periodic evaluation of best neighbors across all off-channels. Best next parent is identified with highest adjustedEase value of another off-channel neighbor compared to current on-channel parent's stickyEase.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Deploy multiple RAPs in each BGN

When deploying a mesh network, there should be multiple paths for each access point back to a WLC. Multiple paths can be added by having multiple RAPs per mesh tree. If a RAP fails and goes offline, other mesh access points will join another RAP with the same bridge group name (BGN) and still have a path back to the WLC.

For best results, follow these simple recommendations:

Ensure that RAPs are configured on different channels to reduce or avoid co-channel interference. MAPs will use background scanning to identify each RAP.

RAPs should be on the same VLAN/subnet to prevent mesh AP address renegotiation on parent change that could delay total mesh convergence time.

Ensure that MAPs have background scanning enabled, to facilitate new parent discovery.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/technical-reference/c9800-best-practices.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining a new potential parent mesh node for connecting with said at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in

	the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.b. wherein said at least one mesh node includes, in addition to the scanning radio, at least two relay radios in each mesh element and wherein said scanning radio and said at least two relay radios operate on different non-interfering channels;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise of a method for operating a mesh network having a plurality of mesh nodes, comprising wherein said at least one mesh node includes, in addition to said scanning radio, at least two relay radios in each mesh element and wherein said scanning radio and said at least two relay radios operate on different non-interfering channels. Cisco's Catalyst 9100 and Aironet Access Points support multiple relay radios, including one or more radios for client access and one or more radios for wireless backhaul, and a scanning radio, that operate on different non-interfering channels by operating at different frequencies with different channels (*e.g.*, 2.4GHz, 5GHz channels). For example, Cisco Catalyst 9130AX Wi-Fi 6 access point is tri-radio (2.4 GHz + 5 GHz with dual-5GHz backhaul using Flexible Radio Assignment). One relay radio can be used for client access, the other used flexibly for backhaul or scanning at two different channels. Each of the 2.4GHz and 5GHz bands have several different non-interfering channels for radio operation. For instance, channels 1, 6 and 11 on the 2.4 GHz band do not interfere. The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands, and there are 24 non-overlapping channels, offering more flexibility and less congestion. **Flexible Radio Assignment (FRA) with tri-radio mode** FRA allows the access points to intelligently determine the operating mode of serving radios based on the RF environment and traffic demands. The access points can operate in the following modes: • Dual-radio mode: One 8x8 5 GHz and one 4x4 2.4 GHz. One radio will serve clients on the 5-GHz band, while the other serves clients on the 2.4-GHz band. • Tri-radio mode: * Dual 4x4 5 GHz and one 4x4 2.4 GHz. With two 4x4 5-GHz and one 4x4 2.4-GHz radios (tri-radio) inside the access point, client device capacity can be increased on demand. The access point's default mode is dual radio with 8x8 5 GHz and 4x4 2.4 GHz. It has the ability to split the 8x8 radio into two separate 4x4 5-GHz radios through software, thereby enabling the benefits of FRA while allowing the 2.4-GHz radio to remain active.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html

Cisco Access Points with Tri-Radio

This topic describes the Tri-Radio feature for Cisco Access Points (APs).

Access Points with three radios are designed for high density environments. The APs by default run one dedicated 2.4-GHz 4x4 mode radio and one 5-GHz 8x8 mode radio. In the default mode, the radios are managed by the Flexible Radio Assignment (FRA), and the Dual Radio Mode is in the disabled state indicating that the radios have either been assigned as client serving 8x8 radio or have not yet been evaluated by FRA.

When you enable the dual radio mode setting, the 8x8 radio is split to two independent 5-GHz 4x4 radios. In this mode, slot 1 and slot 2 are active independent 4x4 radio interfaces. They can serve different user groups with different assigned channels.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/cisco_access_points_with_tri_radio.pdf

As discussed above, these radios are also in addition to the Cisco RF ASIC (Software-Defined Radio) scanning radio, as described in Cisco's datasheet that functions a dedicated monitoring radio, enabling spectrum scanning.

Cisco RF ASIC	The Cisco RF Application-Specific Integrated Circuit (ASIC) is a fully integrated Software-Defined Radio (SDR) that can perform advanced RF spectrum analysis and delivers features such as Cisco CleanAir, Wireless Intrusion Prevention System (WIPS), FastLocate, and Dynamic Frequency Selection (DFS) detection.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the

	Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.c. wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a mesh network wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node. Catalyst 9100 and Aironet Access Points in Cisco Wireless Mesh Networks use AWPP to determine the best path through the MAPs to a RAP. This includes selection of a parent mesh node and establishing a parent-child relationship. AWPP takes RF details into account. During solicitation, the MAP learns all of the available neighbors back to a RAP, receives a token to commence sampling while packets are buffered, and observes the downlink-uplink connectivity ("samples potential new parent nodes"), determines which neighbor offers the best path and then synchronizes with that neighbor to either switch to a new parent that offers a better path, or decides to stay with its current parent.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf at 12; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

AWPP follows a process in selecting parents for a MAP with a radio backhaul, including generating a list of channels with neighbors by passive scanning in the scan state, which is a subset of all backhaul channels. The channels with the neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the best neighbor. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides to sample potential new parents.

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The Mesh Background Scanning feature allows for faster mesh convergence and soliciting and selecting potential new parent nodes. A child MAP will switch between optimal parents based on evaluation across all off-channels.

This new Mesh Background Scanning and Auto parent selection will further improve convergence times and parent selection reliability and stability - a MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time .

A child MAP will switch between optimal parents based on periodic evaluation of best neighbors across all off-channels. Best next parent is identified with highest adjustedEase value of another off-channel neighbor compared to current on-channel parent's stickyEase.

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

	Parent Loss Detection/ Keep Alive Timers	Channel Scan/Seek	DHCP/CAPWAP Information	Time per hop (sec)
Standard	21 / 3 sec	Scan/Seek all 2.4 and 5 Ghz Channels	Renew / Restart CAPWAP	48.6*
Fast	7 / 3 sec	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	20.5*
Very Fast	4 / 1.5 sec	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	15.9*
CCN/BG Scan Fast/VF	4 / 3 sec for 50ms	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	8-10 sec

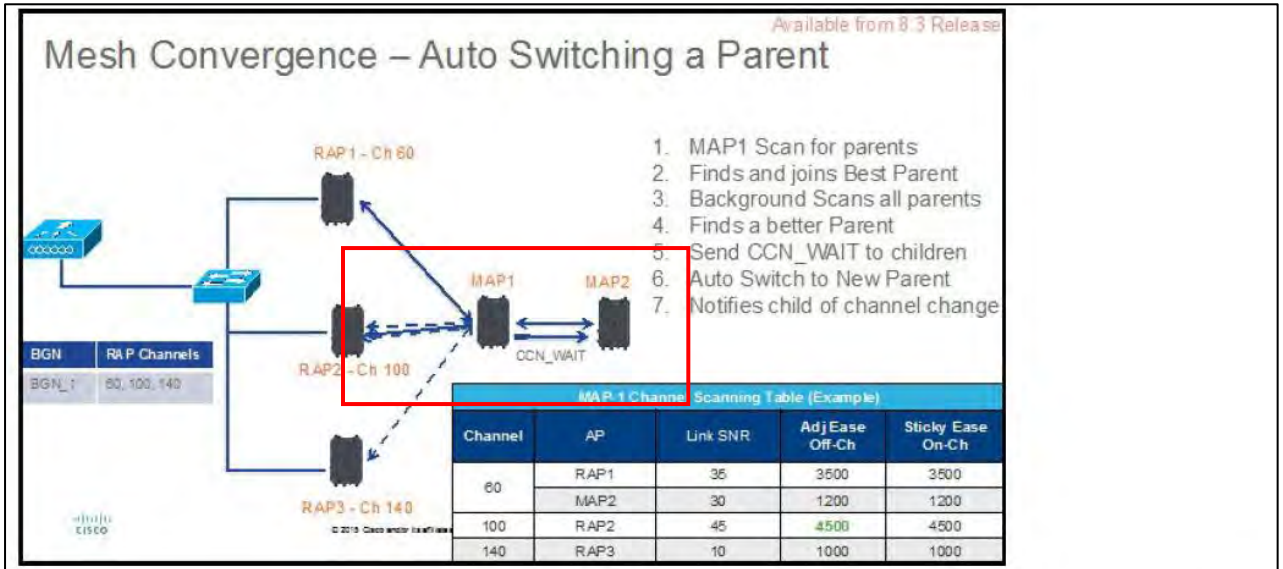
During this process, when a MAP goes off-channel periodically, it transmits NDReq broadcast packets on the selected off-channel and receives NDResp packs from all neighbor.

Non–DFS channel scan

- A MAP goes off-channel periodically, transmits NDReq broadcast packets on the selected off-channel, and shall receive NDResp packets from all 'reachable' neighbors
- Off-channel scan periodicity will occur every 3 seconds and stay for a maximum of 50 milliseconds per off-channel
- NDReq has to be transmitted every 10 milliseconds to send at least 4 messages within 50 milliseconds dwell time to hear better from each neighbor

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

During the process, when the access points go “off-channel,” packets to be sent to a current parent node are buffered by the current parent node, and packets to be sent from the child to its current parent are buffered, as shown in the illustration below. On information and belief, details regarding packet transfer during parent node switching will be shown in Cisco’s proprietary code.



The illustration below shows Mesh convergence and Parent Auto Switching when the "ease" value dictates switching to a better parent is prompted by new "ease value" even if the original parent is still available.

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/support/docs/wireless/wireless-lan-controller-software/214393-configure-rrm-to-defer-off-channel-scann.html (describing delays associated with buffering the client traffic while the APs are performing off-channel scanning)

On information and belief, Cisco’s proprietary mesh network operates the same or similar to 802.11 Passive/Active Scanning, where client radios look for the best access points to connect to. The client radio sends a null data frame to the access point, indicating that the client radio is going into sleep mode, and the access point well then starts buffering packets destined to the client radio. The client radio also buffers packets destined to the access point, and then the client radio goes “off channel” to scan other channels for possible neighboring access points. On information and belief, the access points in a Cisco Wireless Mesh Network operate in the same or similar as shown in Cisco’s proprietary code.

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of sampling other nodes to determine a new potential parent mesh node for connecting with said at least one mesh node in the same way by using radios to regularly scan for and broadcast potential new parents in an RF environment, while maintaining the traffic flow to the old parent before transfer to a new potential parent, and achieves the same result of ensuring the seamless transfer of data flow and a good user experience without degrading performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks solicit and the MAP learns all of the available neighbors back to a RAP to determine the best path. Packets must be buffered to prevent packet loss when a node goes off channel or the parent is switched. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.d. wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a mesh network wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner. In a Cisco Wireless Mesh Network (802.11 Wi-Fi Mesh), Catalyst 9100 and Aironet Series Access Point parent nodes (RAP or MAP) send out beacon frames and other control frames to children (MAPs) to manage how resources are allocated between the parent and child nodes. These frames provide essential information including network parameters (SSID, supported rates, channel information, etc.), parent node status (whether the node is available to accept new child nodes), and link quality (such as signal strength), etc. This could include information elements with Action frames, like Channel Switch Announcement, that can be used to tell a child to go off-channel and scan, which is not a part of the beacon frame. Other control frames, *e.g.*, trigger frames and PS-Poll frames may be used to coordinate or schedule a child/client to go off-channel. The time interval between two consecutive beacon frames is called the beacon interval and is 100ms by default. Source: https://documentation.meraki.com/MR/Wireless_Troubleshooting/Analyzing_Wireless_Packet_Captures On information and belief, the common current parent sends tokens with, or in, the beacon messages (or other control frames) that coordinate the sampling times in a round robin manner, such that each of the children can sample at appropriate times and processing efficiencies are maintained. The implementation of such tokens is enabled in Cisco's proprietary code and may take the form of flags, contention free notices, request/response messages, or other notifications or permissions.

	On information and belief, Cisco's proprietary mesh network operates the same or similar to Target Wake Time ("TWT") in Wi-Fi 6. TWT is a feature in Wi-Fi 6 that allows devices to negotiate specific times to wake up and communicate with the access point, optimizing power consumption and improving network efficiency. Access points negotiate a "TWT agreement" specifying when the device will send or receive data. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of seamless parent node switching in a mesh network in the same way by coordinating sampling times with periodic transmissions and exchanges of broadcast packets (which may be set by or sent at the beacon interval), and achieves the same result of ensuring the seamless transfer of data flow and a good user experience without degrading performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks achieve coordinated and seamless transfer of data flows during parent switching. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to..
Claim 3	**Cisco's Wireless Mesh Network with Catalyst and Aironet Access Points**
3.pre. A mesh network comprising: a plurality of mesh nodes;	Cisco's Wireless Mesh Network (a component of the Cisco Unified Wireless Network solution) comprises a plurality of Wi-Fi mesh nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Series Access Points). Cisco Catalyst 9100 and Aironet Access Points support Wi-Fi capabilities. The Catalyst series is a follow-up from the Cisco Aironet series access points, and the Aironet and Catalyst Access Points contain the same mesh networking capabilities, as shown below in Cisco's documentation and guides that apply to both product lines: Q. Which Cisco access points support Wi-Fi 6? A. Cisco Catalyst 9100 enterprise-class access points support Wi-Fi 6 capabilities extending the power of intent-based networking with hardware and software innovations. The Cisco Catalyst 9100 APs offer resiliency and security while enabling intelligence at the network edge.

Q. What is the portfolio transition from the Cisco Aironet Series to the Cisco Catalyst 9100 platform?

A. The Cisco Catalyst 9100 Access Points consist of the Cisco Catalyst 9105, 9115, 9120, 9130, 9136, 9162, 9164, 9166 and 9166D1, and outdoor 9163E Series. These access points are the follow-ups to the Cisco Aironet 1815, 1850, 2800, and 3800 Series Access Points, respectively. The Cisco Catalyst 9100 Access Points come equipped with Wi-Fi 6/6E capabilities, a better industrial design, and improved RF performance, and deliver reliability, security, and intelligence at scale.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

Cisco Aironet 1560 Series Outdoor Access Points

Product Type	Outdoor and Industrial Wireless
Status	End of Sale **EOL Details**
Series Release Date	05-JUL-2016
End-of-Sale Date	31-JAN-2023
End-of-Support Date	31-JAN-2028

This product is supported by Cisco, but is no longer being sold. Ready for an upgrade? The **Cisco Catalyst 9124AX Series Access Points** offer greater performance and functionality.

Source: https://www.cisco.com/c/en/us/support/wireless/aironet-1560-series/series.html

Mesh deployments are supported on the Cisco Catalyst 9100 Access Points:

Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points?

A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

In addition to the Wave 1 and Wave 2 access points, mesh deployments are supported on the Cisco Catalyst 9100 Access Points, including, *e.g.*, Catalyst 9124 Access Points, Catalyst 9130 Access Points and "AX" access points:

Families of mesh APs

The two families of mesh APs available are the Cisco Wave 1 APs and the Cisco Wave 2 APs + the AX APs. These two families of mesh APs have different capabilities and features, and they are not fully compatible with each other. The Cisco Wave 1 APs do not support the same feature set as the Cisco Wave 2 + AX APs.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg/mesh-access-points.html

From this release, mesh support is included in the Cisco Catalyst 9130AX Series Access Points. All traditional capabilities of mesh are included in the Cisco Catalyst 9130AX Series APs operating in Cisco IOS XE Dublin 17.12.1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html

Cisco Catalyst 9100 Access Points (with the Cisco Aironet series access points) can be deployed in a Cisco Wireless Mesh Network with a controller and software, such as the Catalyst 9800 Controller and Cisco DNA and Prime Infrastructure. Cisco Aironet and Catalyst series access points are both identified as components of a Cisco wireless mesh network:

Mesh Network Components

This section describes the mesh network components.

The Cisco wireless mesh network has four core components:

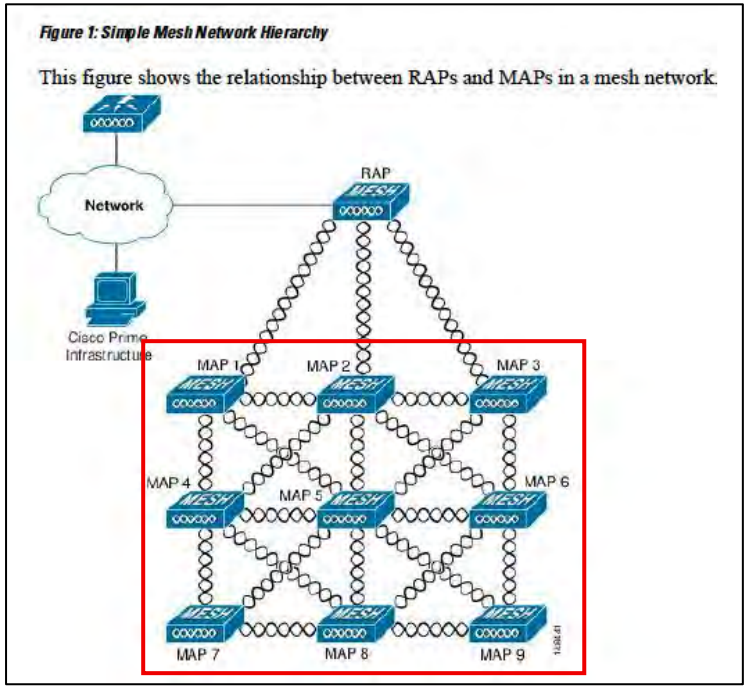
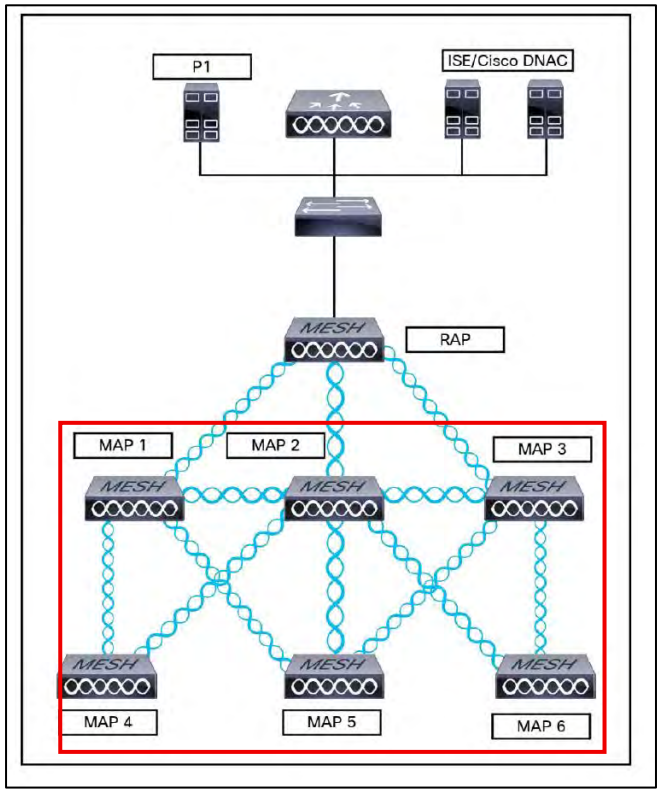
- Cisco C9800 and IOS-XE 17.1
- Cisco Aironet and Catalyst series access points
- Cisco PI and Cisco DNA Center
- Mesh software architecture

Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network.

- Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series.
- Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Cisco Wireless Mesh Networks organize the wireless access points (“mesh nodes”) in a mesh network.

	  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
3.a. wherein each node within said plurality of nodes comprises at least three radios further including:	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a mesh network wherein each node within said plurality of nodes comprises at least three radios. As shown *supra* in claim 2, Cisco's Catalyst 9100 and Aironet Access Points comprise at least three radios. For example, Cisco Catalyst 9130AX Wi-Fi 6 access point is tri-radio (2.4 GHz + 5 GHz with dual-5GHz backhaul using Flexible Radio Assignment). One radio can be used for client access, the other used flexibly for backhaul or scanning at two different channels. Each of the 2.4GHz and 5GHz bands have several different non-interfering channels for radio operation. For instance, channels 1, 6 and 11 on the 2.4 GHz band do not interfere. The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands, and there are 24 non-overlapping channels, offering more flexibility and less congestion.

Flexible Radio Assignment (FRA) with tri-radio mode	FRA allows the access points to intelligently determine the operating mode of serving radios based on the RF environment and traffic demands. The access points can operate in the following modes: • Dual-radio mode: One 8x8 5 GHz and one 4x4 2.4 GHz. One radio will serve clients on the 5-GHz band, while the other serves clients on the 2.4-GHz band. • Tri-radio mode: * Dual 4x4 5 GHz and one 4x4 2.4 GHz. With two 4x4 5-GHz and one 4x4 2.4-GHz radios (tri-radio) inside the access point, client device capacity can be increased on demand. The access point's default mode is dual radio with 8x8 5 GHz and 4x4 2.4 GHz. It has the ability to split the 8x8 radio into two separate 4x4 5-GHz radios through software, thereby enabling the benefits of FRA while allowing the 2.4-GHz radio to remain active.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html

Cisco Access Points with Tri-Radio

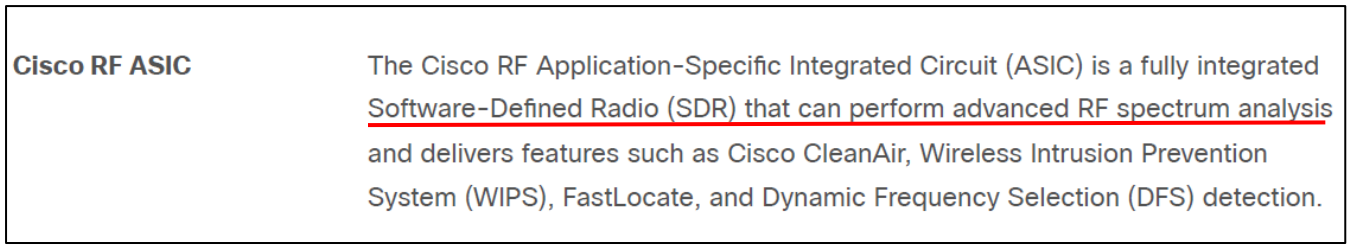
This topic describes the Tri-Radio feature for Cisco Access Points (APs).

Access Points with three radios are designed for high density environments. The APs by default run one dedicated 2.4-GHz 4x4 mode radio and one 5-GHz 8x8 mode radio. In the default mode, the radios are managed by the Flexible Radio Assignment (FRA), and the Dual Radio Mode is in the disabled state indicating that the radios have either been assigned as client serving 8x8 radio or have not yet been evaluated by FRA.

When you enable the dual radio mode setting, the 8x8 radio is split to two independent 5-GHz 4x4 radios. In this mode, slot 1 and slot 2 are active independent 4x4 radio interfaces. They can serve different user groups with different assigned channels.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/cisco_access_points_with_tri_radio.pdf

There is also a Cisco RF ASIC (Software-Defined Radio), as described in Cisco's datasheet that functions a dedicated monitoring radio, enabling spectrum scanning and finding a potential new parent for connection on the mesh nodes uplink radio.

	 If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
3.b. a first relay radio operating on a first RF channel at a first point in time and dedicated to uplink connections to a single current parent node; a second relay radio operating on a second RF channel at a first point in time dedicated to downlink connections to	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a mesh network with a first relay radio operating on a first RF channel at a first point in time and dedicated to uplink connections to a single current parent node and a second relay radio operating on a second RF channel at a first point in time dedicated to downlink connections to zero or more child nodes. Cisco's Catalyst 9100 and Aironet Access Points support multiple relay radios, including one or more radios for wireless backhaul and uplink/downlink connection to the parent node. Cisco's access point radio(s) can be a "dedicated" when being used for the purpose of uplink to the parent node or downlink to zero or more child nodes. Different channels are available for operating the uplink connection, including MU-MIMO and OFDMA technology for splitting channels for uplink to a parent and downlink connections to zero or more child nodes. For example, Cisco Catalyst 9130AX Wi-Fi 6 access point is tri-radio (2.4 GHz + 5 GHz with dual-5GHz using Flexible Radio Assignment) that supports mesh. Radios are used flexibly for backhaul (uplink and downlink) at different channels.

zero or more child nodes; and

Cisco Catalyst 9130AX Series Wireless Access Point Features

The C9130AX series AP is a wireless controller-based product, and supports:

- Four radios, a dual-band 5 GHz (8x8) flexible radio with 2.4 GHz and 4x4 5 GHz, a single-band 5 GHz radio, and an omni IoT radio that can be used with a BLE device.
- Four dual-band and four single-band 5 GHz integrated antennas on the C9130AXI AP models (C9130AXI-x and C9130AXI-EWC-x).

Note The 'x' in the model numbers represents the regulatory domain. For information on supported regulatory domains, see theAP Model Numbers and Regulatory Domains.

- Integrated internal antennas that are omni directional in azimuth, for both 2.4 GHz and 5 GHz bands.
- External antennas on the C9130AXE AP model (C9130AXE-x and C9130AXE-EWC-x).
- Multiuser Multiple-Input Multiple-Output (MU-MIMO) technology for uplink and downlink.
- Orthogonal Frequency Division Multiple Access (OFDMA)-based scheduling for both uplink and downlink.
- Multigigabit Ethernet (mGig)

Source: https://www.cisco.com/c/en/us/td/docs/wireless/access_point/9130ax/quick/guide/ap9130ax-getstart.html

Cisco Access Points with Tri-Radio

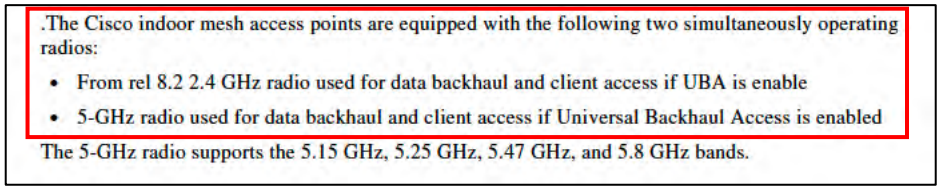
This topic describes the Tri-Radio feature for Cisco Access Points (APs).

Access Points with three radios are designed for high density environments. The APs by default run one dedicated 2.4-GHz 4x4 mode radio and one 5-GHz 8x8 mode radio. In the default mode, the radios are managed by the Flexible Radio Assignment (FRA), and the Dual Radio Mode is in the disabled state indicating that the radios have either been assigned as client serving 8x8 radio or have not yet been evaluated by FRA.

When you enable the dual radio mode setting, the 8x8 radio is split to two independent 5-GHz 4x4 radios. In this mode, slot 1 and slot 2 are active independent 4x4 radio interfaces. They can serve different user groups with different assigned channels.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/cisco_access_points_with_tri_radio.pdf

In addition, while 5 GHz is often used for the primary backhaul (communication between access points), both the 2.4 GHz and 5GHz radios are configurable to be used for backhaul, *e.g.*, uplink and downlink at different frequencies, for example, when Universal Backhaul Access (UBA) is enabled.



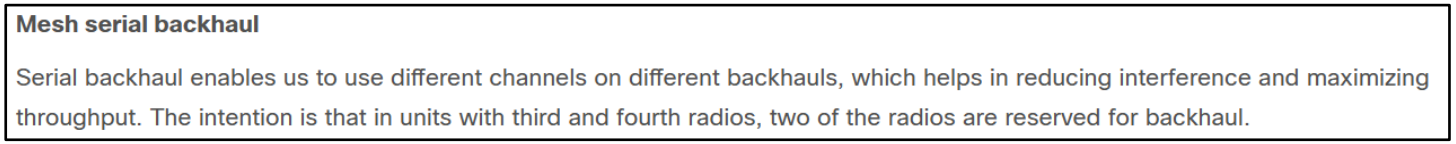
.The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios:

- From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable
- 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled

The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands.

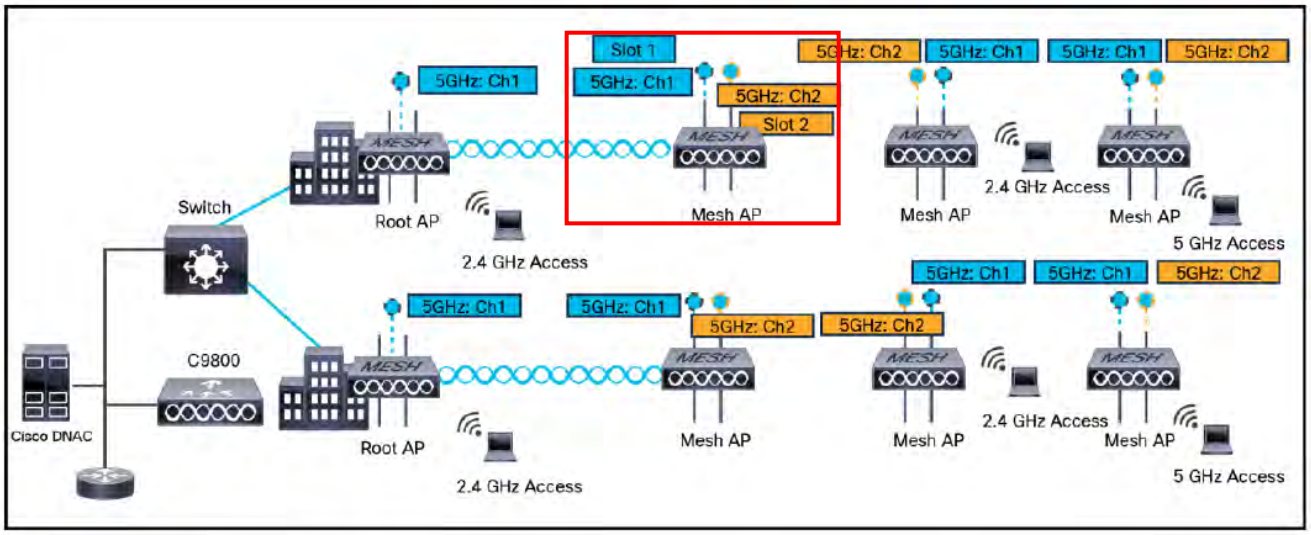
Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-6/b_mesh_86/m_mesh-network-components.html; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/technotes/8-5/cisco-enterprise-mobility-design-guide-8-5.pdf

With Mesh Serial Backhaul, two of the radios are reserved for backhaul. The Mesh Serial Backhaul feature in a mesh access point (MAP) allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio, and use different channels.



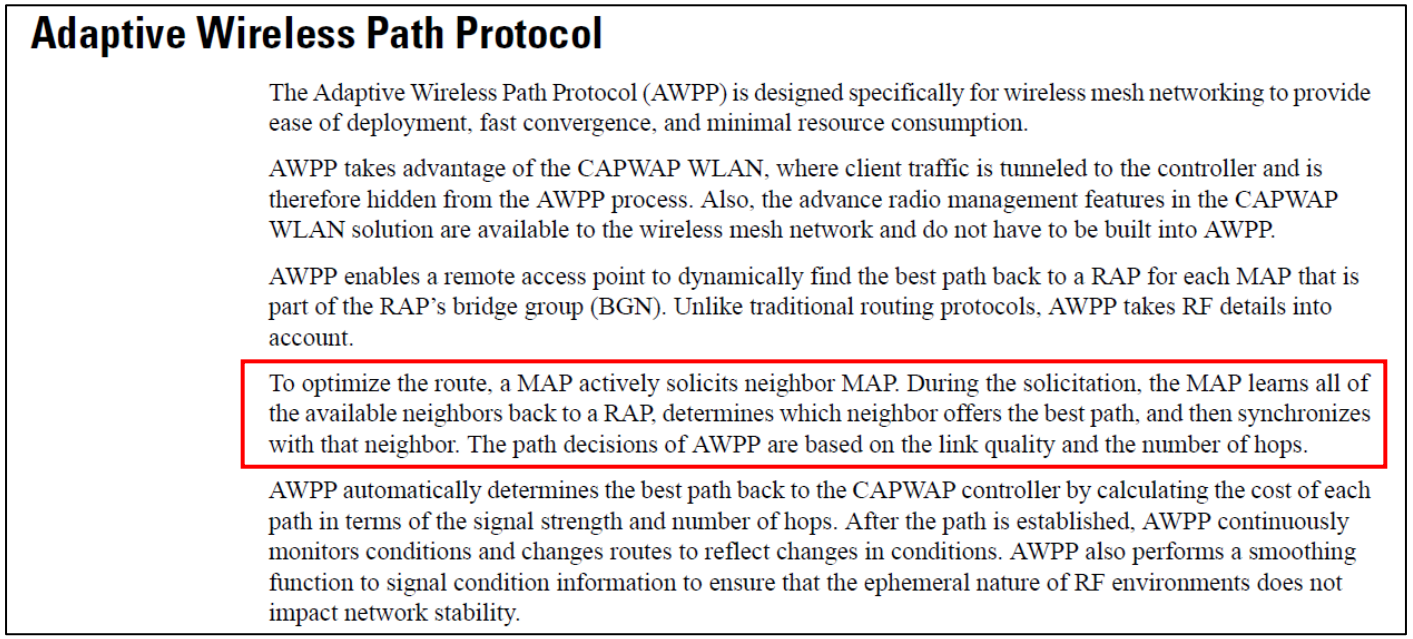
Mesh serial backhaul

Serial backhaul enables us to use different channels on different backhauls, which helps in reducing interference and maximizing throughput. The intention is that in units with third and fourth radios, two of the radios are reserved for backhaul.

	 Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-7/config-guide/b_wl_17_7_cg/m_mesh_ewlc.html#mesh-backhaul-24ghz-5ghz If these claim elements are not literally infringed, they are infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed elements are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
3.c. a dedicated scanning radio operating on a third RF channel at a first point in time and	*See supra* discussion for claim elements 2.a and 3.a, the contents of which are incorporated herein by reference.

configured to scan a Radio Frequency (RF) environment to discover new potential parent mesh nodes for connection with said first relay radio;	
3.d. said first, second, and third RF channels are different from each other	*See supra* discussion for claim elements 2.a, 2.b, 3.a and 3.b, the contents if which are incorporated herein by reference.
3.e. wherein after said dedicated scanning radio discovers a potential new parent node, said first relay radio samples the RF link to said potential new parent node using its uplink radio and concurrent with said sampling, packets to be sent to the mesh node from its current parent node are buffered by the current parent node, and packets to be sent from the mesh node to its current parent node are buffered by the mesh node wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a mesh network wherein after said dedicated scanning radio discovers a potential new parent node, said first relay radio samples the RF link to said potential new parent node using its uplink radio and concurrent with said sampling, packets to be sent to the mesh node from its current parent node are buffered by the current parent node, and packets to be sent from the mesh nod e to its current parent node are buffered by the mesh node wherein sampling times are coordinated among multiple mesh nodes having a common parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner. Catalyst 9100 and Aironet Access Points in Cisco Wireless Mesh Networks use AWPP to determine the best path through the MAPs to a RAP. This includes selection of a parent mesh node and establishing a parent-child relationship. AWPP takes RF details into account. During solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path and then synchronizes with that neighbor.

children in a round-robin manner.



Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf at 12; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

As shown *supra* in claim elements 2.a-2.d (which are incorporated herein by reference), a dedicated scanning radio discovers neighbors and potential new parent nodes. AWPP follows a sampling process in selecting parents for a MAP with a radio backhaul, including generating a list of channels with neighbors by passive scanning in the scan state, which is a subset of all backhaul channels. The channels with the neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the best neighbor.

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

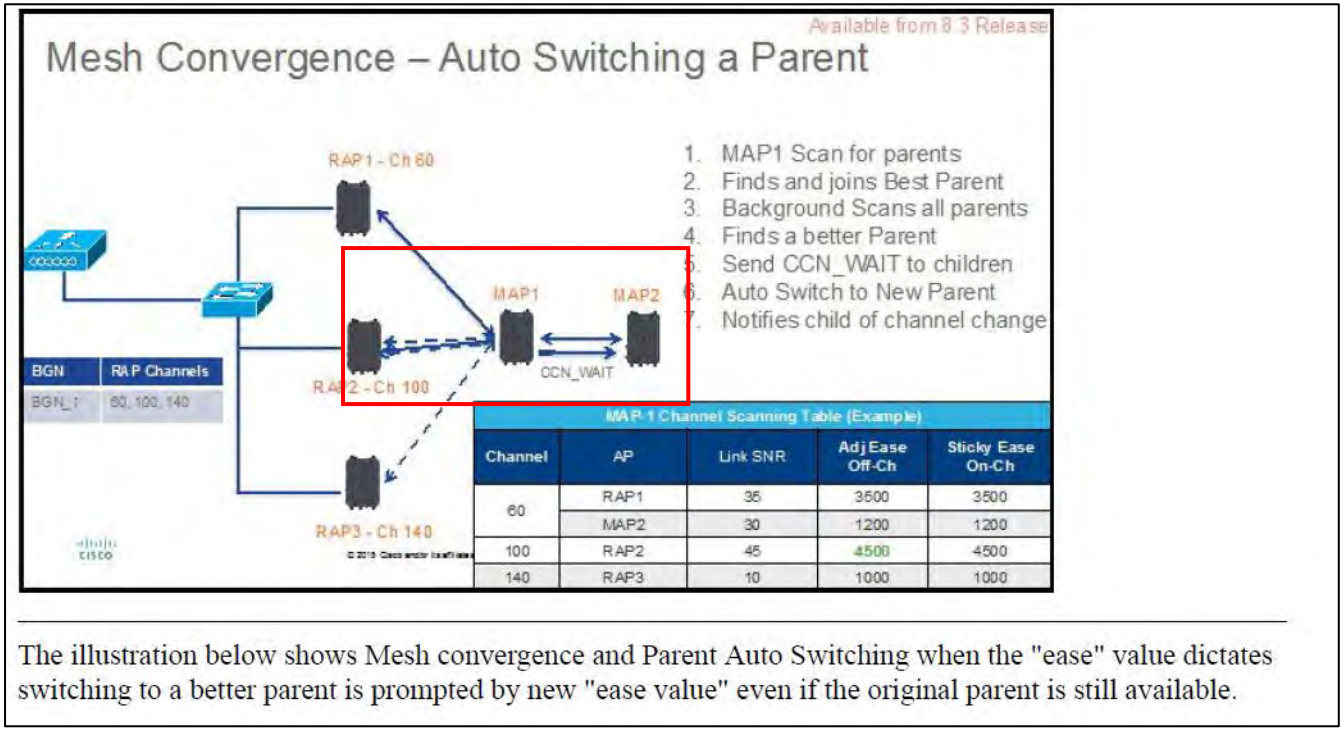
Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

As shown *supra* in claim elements 2.a-2.d (which are incorporated herein by reference), during the process, when the access points go "off-channel," packets to be sent to a child node from the current parent node are buffered by the current parent, and packets to be sent from the child node to its current parent are buffered by the child, as shown in the illustration below. On information and belief, details regarding packet transfer during parent node switching will be shown in Cisco's proprietary code.



The illustration below shows Mesh convergence and Parent Auto Switching when the "ease" value dictates switching to a better parent is prompted by new "ease value" even if the original parent is still available.

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/support/docs/wireless/wireless-lan-controller-software/214393-configure-rrm-to-defer-off-channel-scann.html (describing delays associated with buffering the client traffic while the APs are performing off-channel scanning)

As shown *supra* in claim elements 2.a-2.d (which are incorporated herein by reference), sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Exhibit 4 – U.S. Patent No. 7,894,385 (Cisco's Wireless Mesh Networks with Meraki and Catalyst CW Access Points)

Cisco Systems, Inc. ("Cisco") has and continues to make, use, sell, offer for sale, and/or import the Cisco's Meraki Wireless Mesh Networks including but not limited to Cisco's Meraki and Catalyst CW Access Points, with related cloud architecture and software (the "Cisco Accused Products," *see* Exhibit 11), and has infringed and infringes at least Claims 2-3 (the "Asserted Claims") of U.S. Patent No. 7,894,385 (the "'385 Patent").[1]

Cisco has and continues to infringe (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '385 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringes and has infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringes and has infringed by testing, setting up, and using Cisco Meraki Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Meraki Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Meraki Wireless Mesh Network in the United States. As indicated in the Meraki Mesh Networking Guide, each Meraki access point has unique mesh characteristics independent of general technology specifications listed on the datasheet. Cisco recommends a Meraki Systems Engineer as a resource to ensure that an access point's capabilities match the desired design targets.[3] The Meraki Trial Program allows customers to run a real world proof of concept to test the design, which is arranged by a Cisco Meraki account representative. Meraki is also cloud-based architecture, and the infringing mesh networking methods and systems are controlled and put into service by Cisco.[4] Cisco knows how its customers are configuring the specific access points sold, as customers must register to the Meraki cloud and use the invoice or serial number of each purchased access point.[5]

---

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g*., https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

[3] *See, e.g*., https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

[4] *See Centillion Data Systems, LLC v. Qwest Communications International, Inc.*, 631 F.3d 1279 (Fed. Cir. 2011); *NTP, Inc. v. Research in Motion, Ltd*., 418 F. 3d 1282, 1317-1318 (Fed. Cir. 2005); *see also SiRF Tech. Inc. v. ITC*, No. 2009-1262 (Fed. Cir. Apr. 12, 2010) (direct infringement where the claims do not specifically require action by its customers or end users).

[5] *See, e.g*., https://documentation.meraki.com/MR/MR_Installation_Guides/CW9163E-MR_Installation_Guide.

Cisco indirectly infringes and has infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the ’385 Patent and make and use the infringing Cisco Meraki Wireless Mesh Network. Specifically, Cisco’s Meraki Mesh Deployment Guides, Meraki Mesh Networking Guide, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco’s website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Meraki Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 2-3 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Meraki Wireless Mesh Network, and the infringing configuration and operation of the cloud architecture, access points, and software that comprise the infringing Cisco Meraki Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Meraki Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the ’385 Patent since 2009 when Cisco executives met with MeshDynamics to discuss partnering opportunities and the application that would result in the ’385 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Meraki Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’385 Patent, or in the alternative was willfully blind to the ’385 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringes and has infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Meraki Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringes the ’385 Patent. Cisco Accused Products come enabled as a mesh access points and automatically connect to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particularly suited for Cisco’s Meraki Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Meraki Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’385 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed and continue to infringe the Asserted Claims of the '385 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics's preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 2	Cisco Meraki Wireless Mesh Networks with Meraki and Catalyst CW Access Points
2.pre. A mesh network comprising: a plurality of mesh nodes;	Cisco's Meraki Wireless Mesh Network comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Meraki and Catalyst CW Access Points). The Meraki cloud-managed wireless series is a line of enterprise-grade WLAN access points ("APs") that leverage the Cisco Meraki cloud-managed architecture. Both Cisco Meraki and Catalyst CW Access Points support Wi-Fi capabilities and can be paired with the Meraki cloud platform. Meraki AP models start with MR, *e.g.* MR33, MR36,

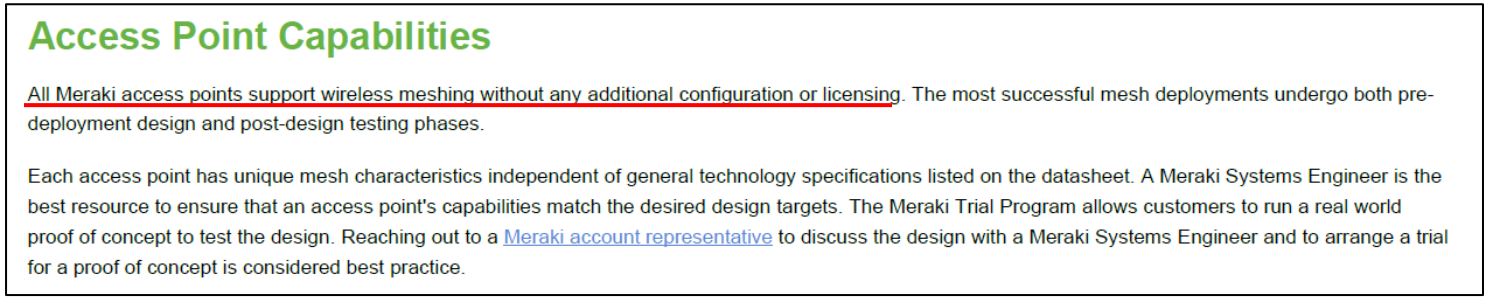
MR46E, and the Catalyst APs with CW, *e.g.* CW9162, CW9164, CW9166, are dual-mode APs which can be used with the Catalyst 9800 controller or the Meraki cloud architecture.

Source: Meraki Family Datasheet; Meraki Wi-Fi Models

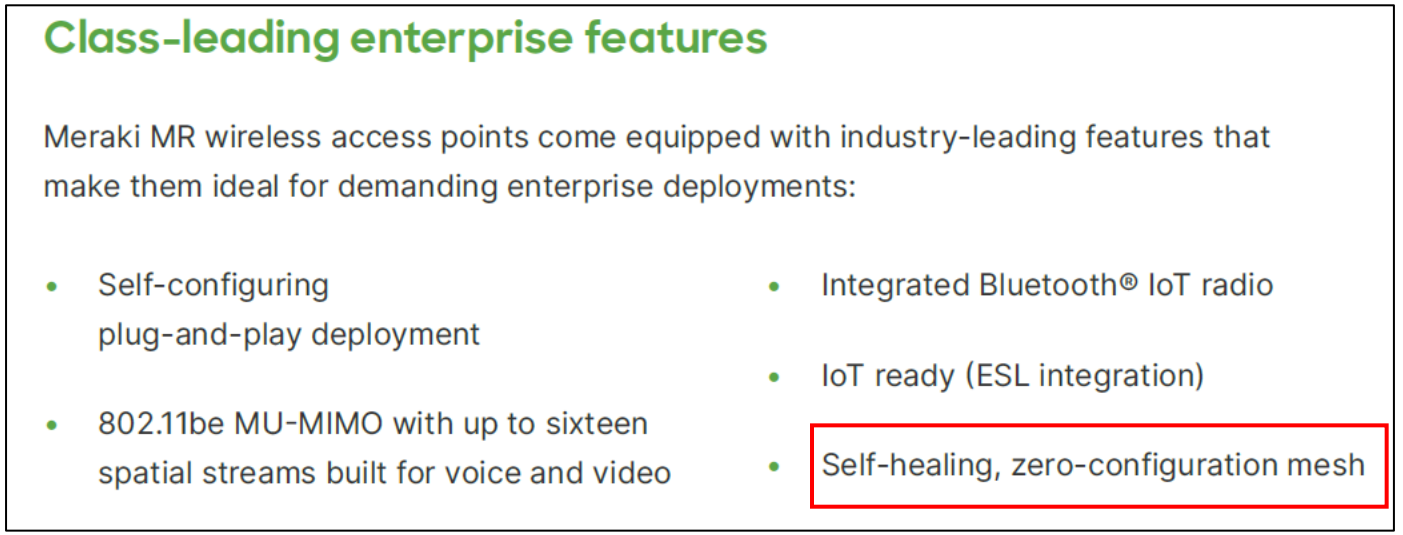
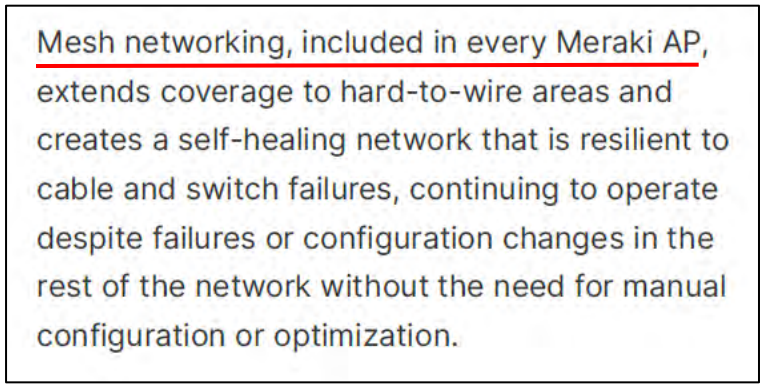


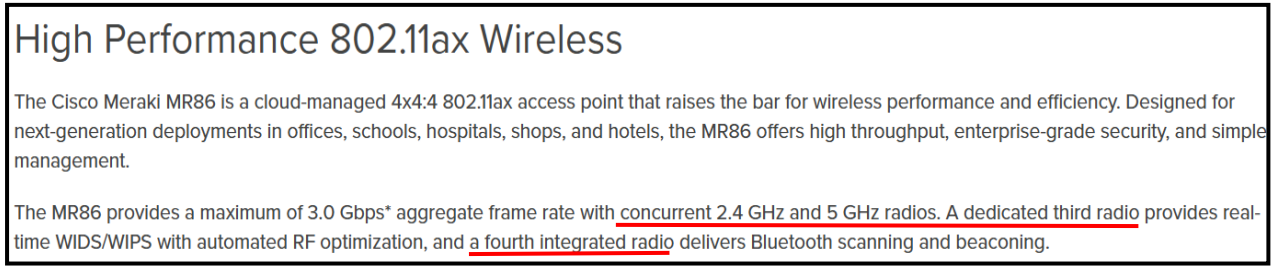
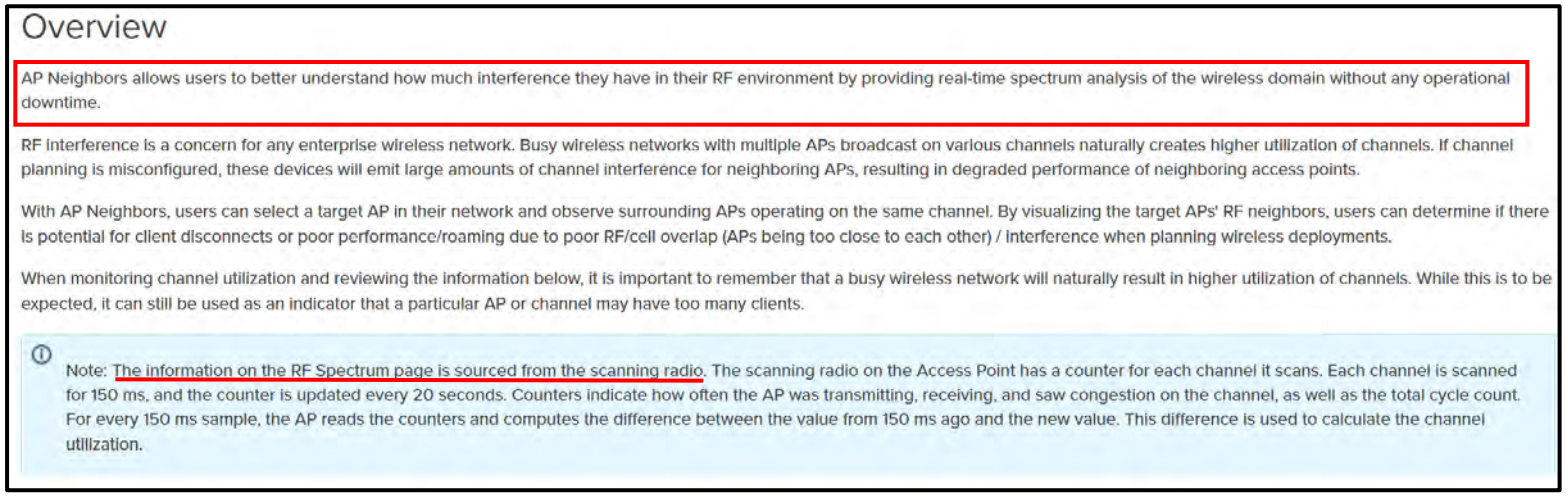
Source: Wireless Mesh Networking; *see also* Meraki Best Practices - Wireless Mesh Networking

All Meraki APs support wireless meshing without any additional configuration ("zero configuration") or licensing.



Source: Mesh Deployment Guide

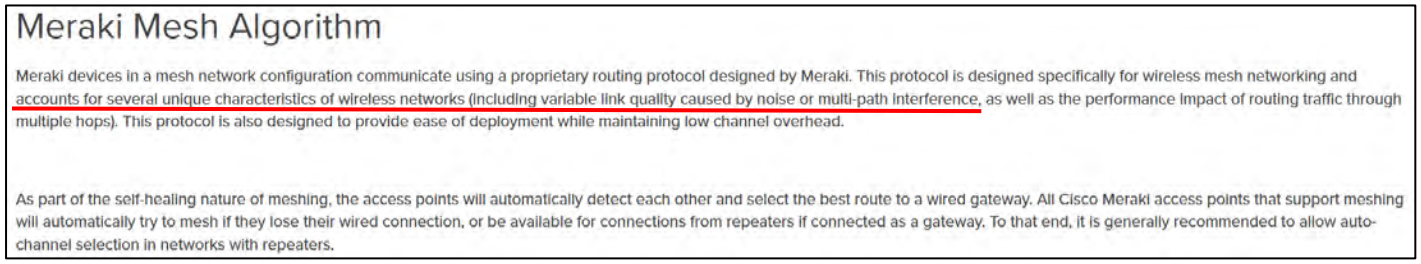
	 Source: Meraki Family Datasheet  Source: Meraki Family Datasheet
2.a. wherein at least one mesh node of the mesh network is configured to scan a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network comprise wherein at least one mesh node of the mesh network, scanning a Radio Frequency (RF) environment using a dedicated scanning radio to determine a new potential parent mesh node for connecting with said at least one mesh node. APs in Cisco's Meraki Wireless Mesh Network have dedicated scanning radios. In addition, Cisco's access point radio(s) can be a "dedicated scanning radio" when being used for the purpose of scanning to discovery neighbor or mesh

potential parent mesh node for connecting with said at least one mesh node;	nodes (“new potential parent mesh node”) but may subsequently be used for other purposes such as backhaul or sampling. For example, the MR86 datasheet identifies a dedicated scanning radio in addition to the 2.4/5GHz radios. This radio enables spectrum scanning to find other neighboring mesh nodes to determine interference and to facilitate parent/gateway discovery.  Source: MR86 Datasheet  Source: AP Neighbors Meraki indoor access points also feature a radio dedicated to full-time scanning for neighbors and interference.

Indoor APs feature a radio dedicated to full-time scanning, rogue AP containment, and automatic RF optimization. With Air Marshal, it's possible to set up a real-time wireless intrusion detection and prevention system (WIDS/WIPS) with user-defined threat remediation policies and intrusion alarms, enabling secure wireless environments without complex setup or systems integration. Auto RF eliminates the need for manual RF configuration by scanning the environment for utilization, interference, and other metrics, and computing the optimal channel and power settings for every AP in the network. Meraki WLANs are fully HIPAA and PCI compliant.

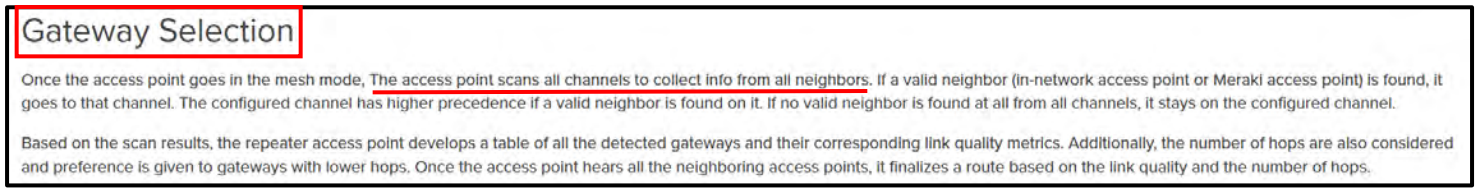
Source: Meraki Family Datasheet

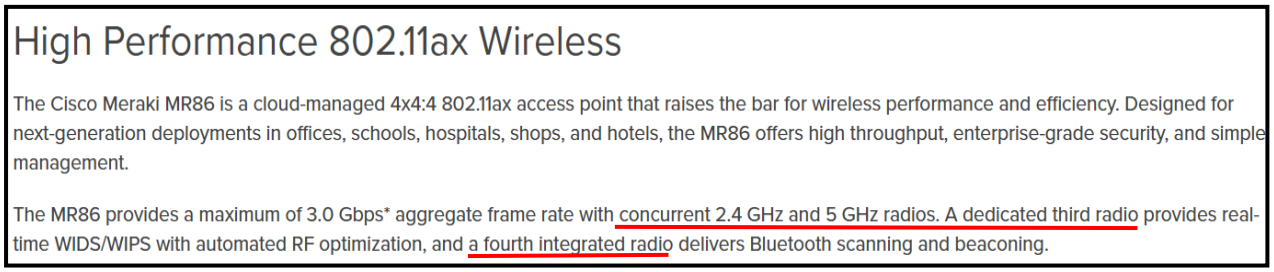
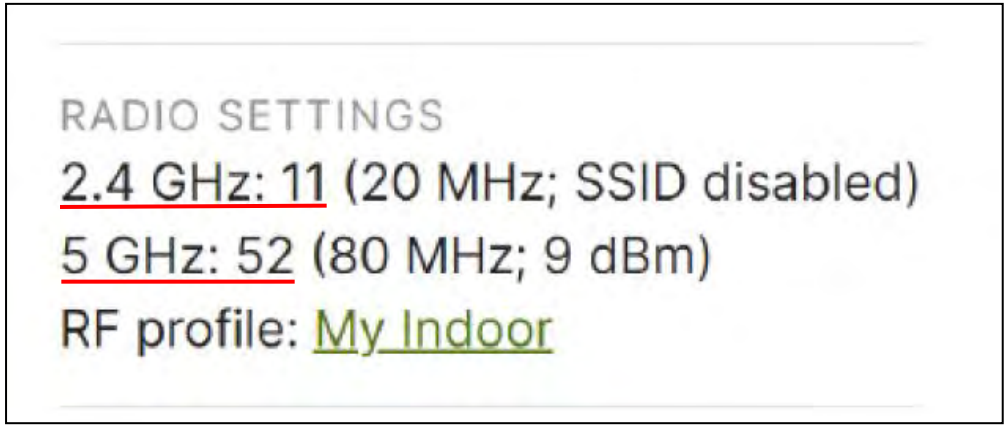
In the Cisco Meraki Mesh Network, the dedicated scanning radio also aids determination of a new potential parent mesh node for connecting with said at least one mesh node by providing input to the Meraki Mesh Algorithm. For example, the protocol is designed for mesh networking and accounts for several characteristics of the network, including multi-path interference. The access points will automatically detect each other and select the best route to a wired gateway or parent.



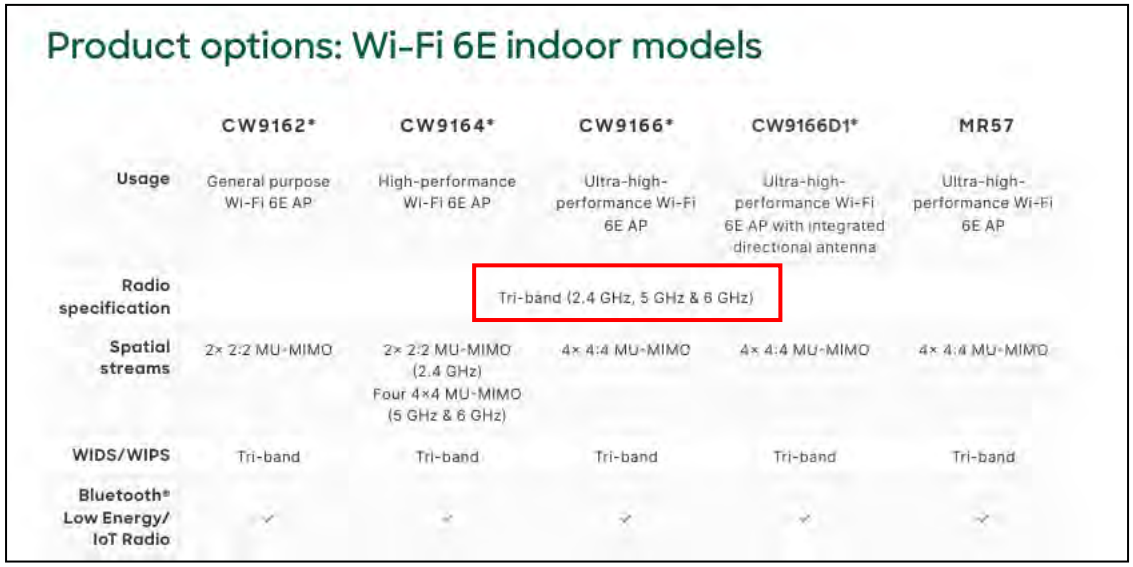
Source: Wireless Mesh Networking

The scanning radio also scans to find all neighbors to select the best parent or gateway.

	 Source: Wireless Mesh Networking APs in Cisco's Meraki Wireless Mesh Network mesh automatically ("determines a new potential parent mesh node") on channels with the least interference as possible using channel selection algorithms. Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco Meraki and Catalyst CW Access Points perform the same function of collecting information to determine a new potential parent mesh node for connecting with said at least one mesh node in the same way by using radios to regularly scan for and track potential new gateways or parents in an RF environment and achieves the same result by determining the best gateway or parent mesh node for optimal network performance based upon the information collected. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new gateways or parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.b. wherein said at least one mesh node includes, in addition to the scanning radio, at least two relay radios in each mesh element and	Cisco's Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network comprise a mesh network wherein said at least one mesh node includes, in addition to said scanning radio, at least two relay radios in each mesh element and wherein said scanning radio and said at least two relay radios operate on different non-interfering channels.

wherein said scanning radio and said at least two relay radios operate on different non-interfering channels;	Cisco's Meraki and Catalyst CW Access Points support multiple relay radios, including one or more radios for client access and one or more radios for wireless backhaul, and a scanning radio, that operate on different non-interfering channels by operating at different frequencies with different channels (*e.g.*, 2.4GHz, 5GHz channels). As shown *supra*, APs in Cisco's Meraki Wireless Mesh Network have dedicated scanning radios operating at a different frequency than the relay radios. For example, the MR86 datasheet identifies a dedicated scanning radio in addition to the two 2.4/5GHz relay radios. The Cisco Meraki MR86 access point also is dual-radio (2.4 GHz + 5 GHz) that can be used flexibly for backhaul, client access or scanning at different channels. Each of the 2.4GHz and 5GHz bands have several different non-interfering channels for radio operation. For instance, channels 1, 6 and 11 on the 2.4 GHz band do not interfere. The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands, and there are 24 non-overlapping channels, offering more flexibility and less congestion.  Source: MR86 Datasheet  Source: Wireless Mesh Networking

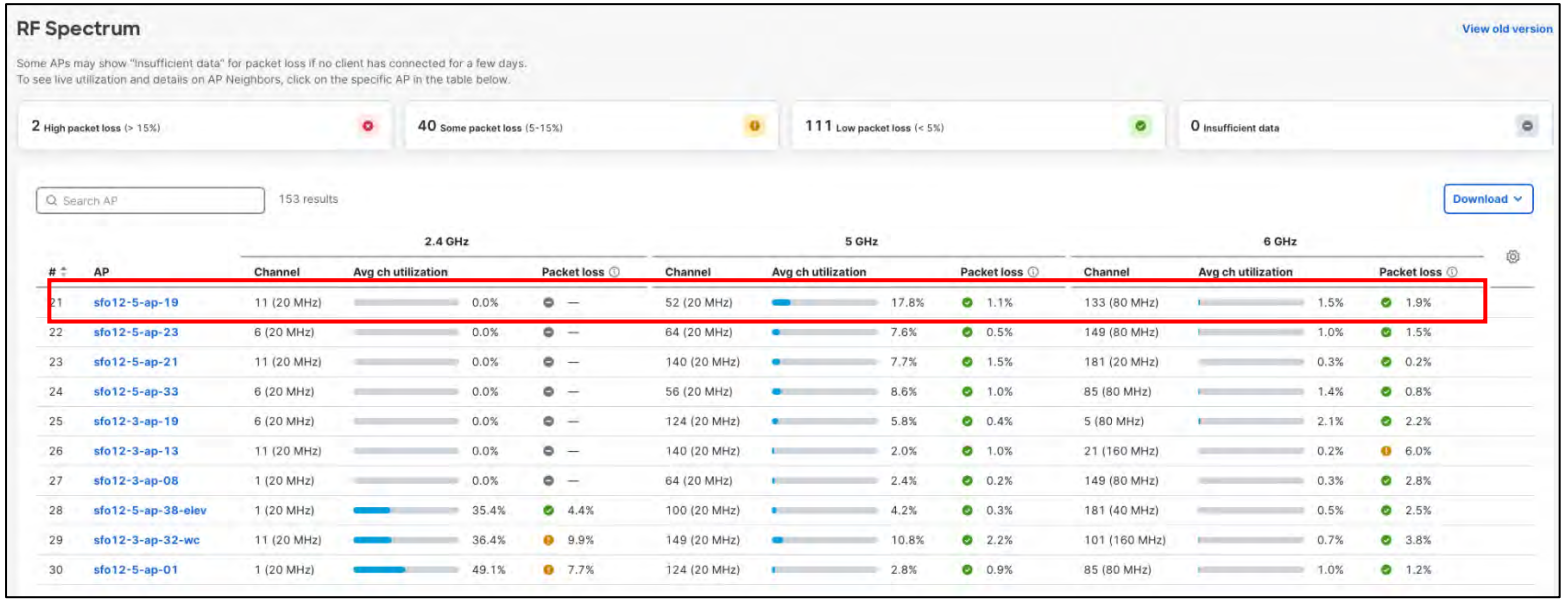
Additional radios are tri-band, with a scanning radio, and each radio is operating on a different frequency and channel.



Product options: Wi-Fi 6E indoor models

	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: Meraki Family Datasheet



RF Spectrum

View old version

Some APs may show "Insufficient data" for packet loss if no client has connected for a few days.
To see live utilization and details on AP Neighbors, click on the specific AP in the table below.

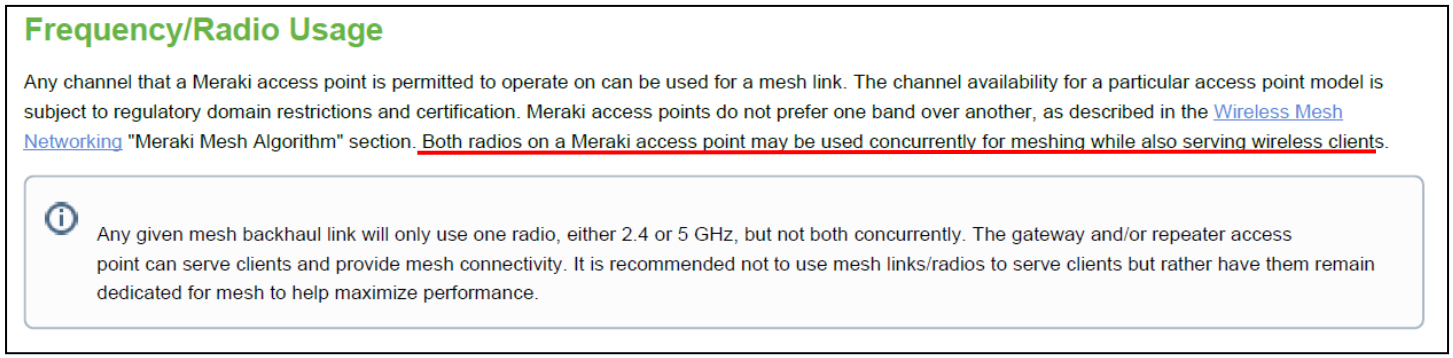
2 High packet loss (> 15%) | 40 Some packet loss (5-15%) | 111 Low packet loss (< 5%) | 0 Insufficient data

Search AP | 153 results | Download

#	AP	2.4 GHz Channel	2.4 GHz Avg ch utilization	2.4 GHz Packet loss	5 GHz Channel	5 GHz Avg ch utilization	5 GHz Packet loss	6 GHz Channel	6 GHz Avg ch utilization	6 GHz Packet loss
21	sfo12-5-ap-19	11 (20 MHz)	0.0%	—	52 (20 MHz)	17.8%	1.1%	133 (80 MHz)	1.5%	1.9%
22	sfo12-5-ap-23	6 (20 MHz)	0.0%	—	64 (20 MHz)	7.6%	0.5%	149 (80 MHz)	1.0%	1.5%
23	sfo12-5-ap-21	11 (20 MHz)	0.0%	—	140 (20 MHz)	7.7%	1.5%	181 (20 MHz)	0.3%	0.2%
24	sfo12-5-ap-33	6 (20 MHz)	0.0%	—	56 (20 MHz)	8.6%	1.0%	85 (80 MHz)	1.4%	0.8%
25	sfo12-3-ap-19	6 (20 MHz)	0.0%	—	124 (20 MHz)	5.8%	0.4%	5 (80 MHz)	2.1%	2.2%
26	sfo12-3-ap-13	11 (20 MHz)	0.0%	—	140 (20 MHz)	2.0%	1.0%	21 (160 MHz)	0.2%	6.0%
27	sfo12-3-ap-08	1 (20 MHz)	0.0%	—	64 (20 MHz)	2.4%	0.2%	149 (80 MHz)	0.3%	2.8%
28	sfo12-5-ap-38-elev	1 (20 MHz)	35.4%	4.4%	100 (20 MHz)	4.2%	0.3%	181 (40 MHz)	0.5%	2.5%
29	sfo12-3-ap-32-wc	11 (20 MHz)	36.4%	9.9%	149 (20 MHz)	10.8%	2.2%	101 (160 MHz)	0.7%	3.8%
30	sfo12-5-ap-01	1 (20 MHz)	49.1%	7.7%	124 (20 MHz)	2.8%	0.9%	85 (80 MHz)	1.0%	1.2%

Source: AP Neighbors

The relay radios may be used concurrently for meshing while also serving wireless clients (*i.e.*, for wirelessly communicating with another node or a client device in the wireless mesh network).

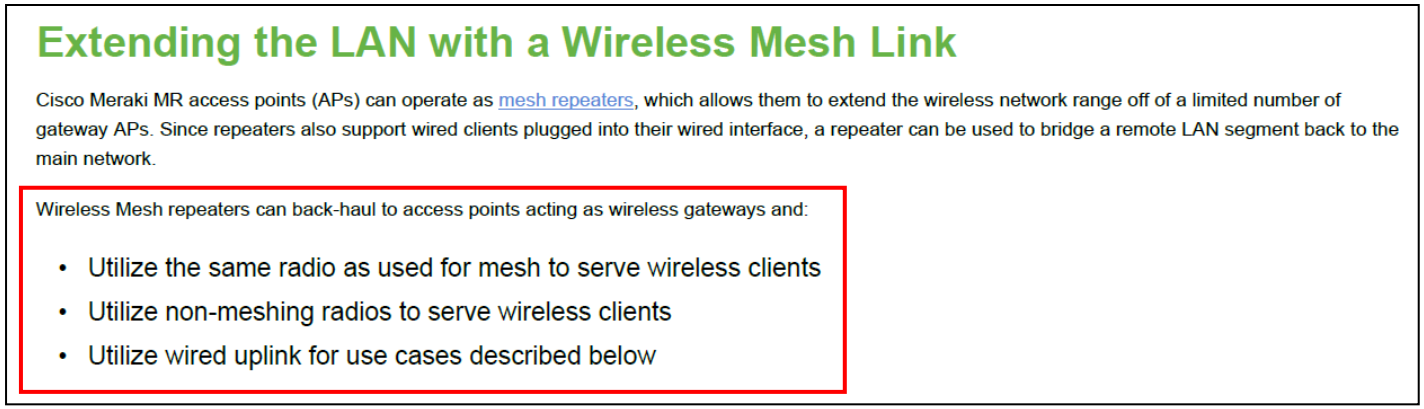


Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: Mesh Deployment Guide



Extending the LAN with a Wireless Mesh Link

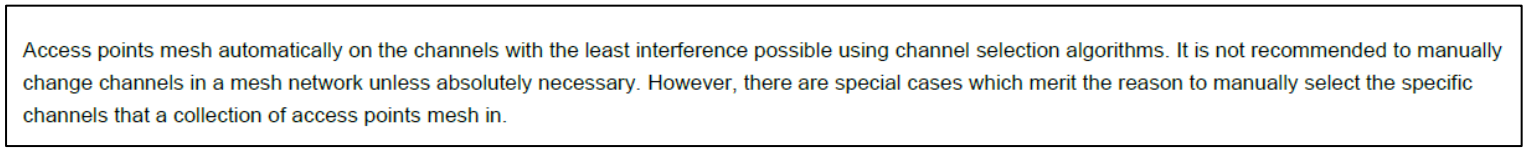
Cisco Meraki MR access points (APs) can operate as mesh repeaters, which allows them to extend the wireless network range off of a limited number of gateway APs. Since repeaters also support wired clients plugged into their wired interface, a repeater can be used to bridge a remote LAN segment back to the main network.

Wireless Mesh repeaters can back-haul to access points acting as wireless gateways and:

- Utilize the same radio as used for mesh to serve wireless clients
- Utilize non-meshing radios to serve wireless clients
- Utilize wired uplink for use cases described below

Source: Extending the LAN with a Wireless Mesh Link

APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference as possible using channel selection algorithms.



Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in.

	Source: Manually Changing Channels If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents or gateways in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node or gateway for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents or gateways. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.c. wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node;	Cisco's Meraki and Catalyst CW Access Points in the Cisco Wireless Mesh Network comprise a mesh network wherein while said at least one mesh node samples potential new parent nodes, packets to be sent to said at least one mesh node from its current parent node are buffered by the current parent node, and packets to be sent from said at least one mesh node to its current parent are buffered by said at least one mesh node. Cisco Meraki and Catalyst CW Access Points in Cisco Wireless Mesh Networks use the Meraki Mesh Algorithm to determine the best path through the repeaters to a gateway. This includes selection of a gateway mesh node and establishing a parent-child relationship. The Meraki Mesh Algorithm takes RF details into account. During this process, the repeater learns all of the available neighbors back to a gateway, receives a token to commence sampling while packets are buffered and observes the downlink-uplink connectivity ("samples potential new parent nodes"), determines which neighbor offers the best path and then synchronizes with that neighbor to either switch to a new parent that offers a better path or decides to stay with its current parent or go to another path. Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

To find the best path and update the routing tables, including the best parent back to the gateway or a new mesh gateway, each Meraki AP sends out link probe packets (known as mesh probes) at different bit rates and varying sizes. Four different types of probes at different data rates are sent in a batch of 15 seconds on both 2.4/5GHz bands (*i.e.*, the initial discovery and processes involve off-channel scanning). These mesh probes are used to come up with a link quality metric ("samples potential new parent nodes"). The algorithm sends out periodic discovery frames of varying sizes to discover mesh neighbors. After one neighbor is selected as the gateway, new routes are passively identified on the same operating channel. A mesh link, however, may take a couple of minutes to establish, and may not always be used for every data flow if a new route with better link metrics becomes available. At this time, on information and belief, to establish the link or to sample a link, the packets to be sent to the mesh node or repeater from its current parent node or gateway are buffered by the parent or gateway, and packets to be sent from the one mesh node or repeater are buffered by the mesh node or repeater.

Mesh Probes

Each Meraki access point sends out link probe packets (known as mesh probes) at different bit rates and varying sizes. Because these packets are sent as broadcast frames, no ACK frames are needed from receiving stations. Four different types of probes at different data rates are sent in a batch of 15 seconds on both (2.4 /5 GHz) bands. All access points listen to the mesh probes and depending on the number of mesh probes correctly received, come up with a link quality metric as shown in dashboard.

Source: Wireless Mesh Networking

Mesh Convergence Time

The Meraki mesh algorithm is designed for stationary Meraki access points with variable links. The algorithm sends out periodic discovery frames of varying sizes on the 2.4 GHz and 5 GHz data radios to discover mesh neighbors. More discussion on neighbor discovery and gateway selection can be found in Wireless Mesh Networking. After one neighbor is selected as the gateway, new routes are passively identified on the same operating channel. A mesh link may take a couple minutes to establish and may not always be used for every data flow if a new route with better link metrics becomes available.

Source: Mesh Deployment Guide

When a repeater changes gateways and there are existing traffic flows, it will also ensure the seamless transfer of data flow and a good user experience. The repeater does not move traffic flows immediately, but uses the old route for 5 seconds before being transferred to the new route. In this way, the Meraki Mesh Network further buffers packets to ensure traffic is not disrupted during the sampling process to ensure the AP is connected to the new gateway.

Gateway Change Process

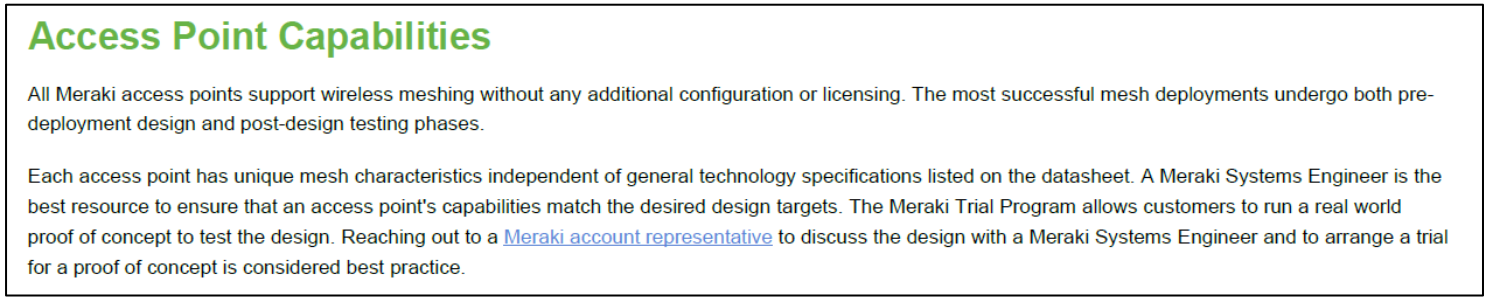
In most cases, when the repeater changes gateway there are existing traffic flows that need to be considered and the repeater needs to ensure that the new gateway does not provide a degraded performance. To ensure seamless transfer of data flow and a good user experience, a repeater does not move traffic flows immediately. Existing traffic flows use the old route for 5 seconds before being transferred to the new route using the new gateway.

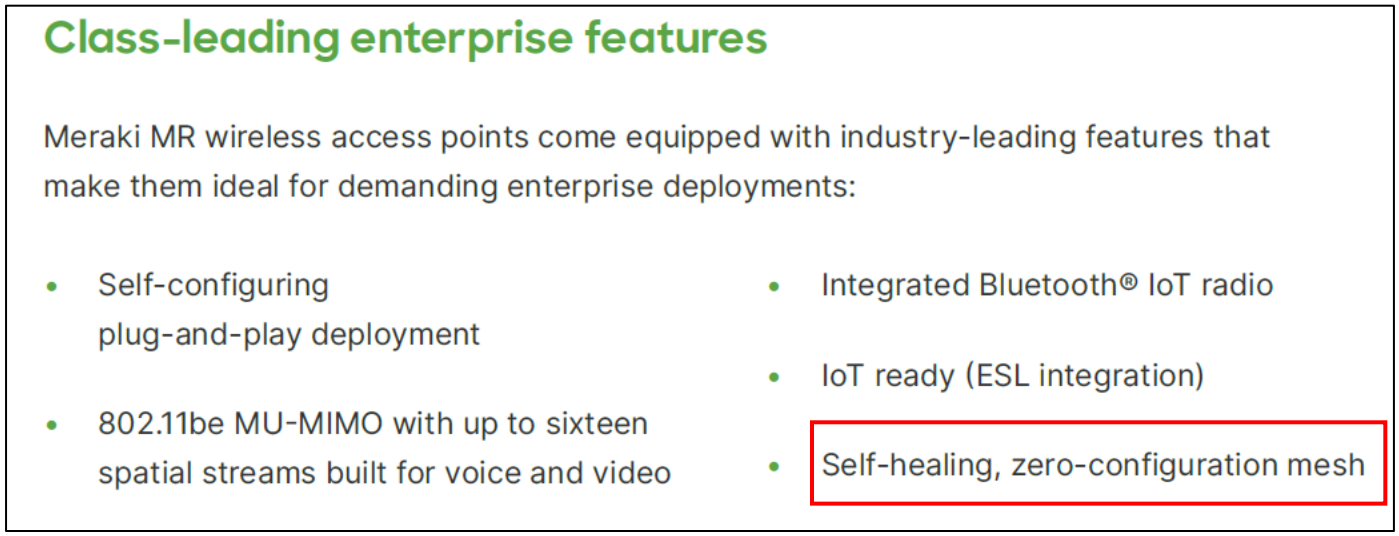
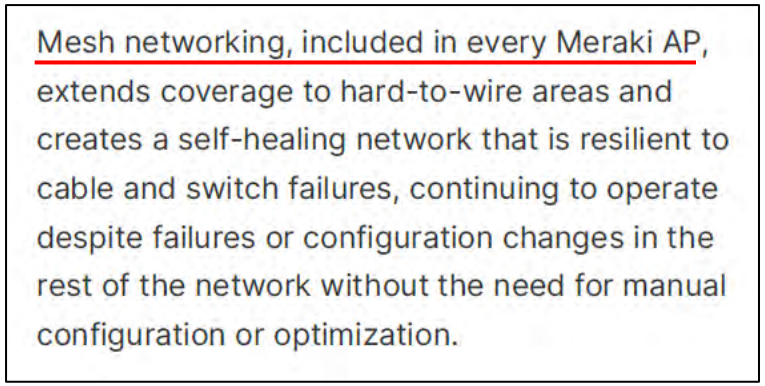
Source: Wireless Mesh Networking

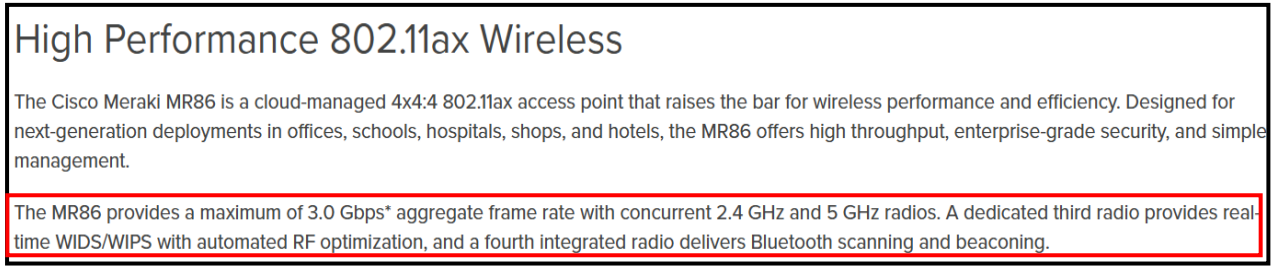
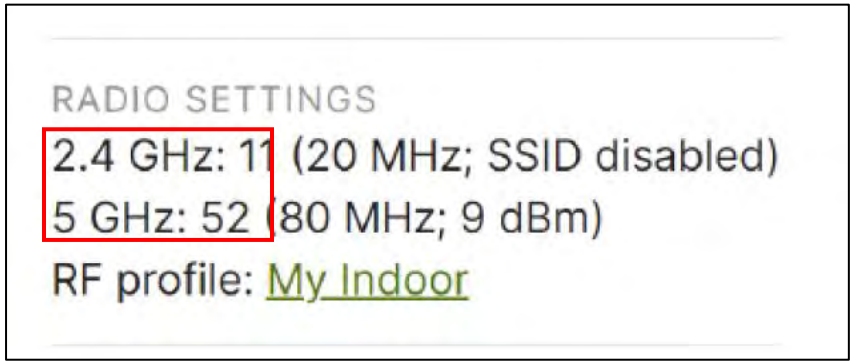
On information and belief, Cisco's proprietary mesh network operates the same or similar to 802.11 Passive/Active Scanning, where client radios look for the best access points to connect to. The client radio sends a null data frame to the

	access point, indicating that the client radio is going into sleep mode, and the access point well then starts buffering packets destined to the client radio. The client radio also buffers packets destined to the access point, and then the client radio goes "off channel" to scan other channels for possible neighboring access points. On information and belief, the access points in a Cisco Meraki Wireless Mesh Network operate in the same or similar as shown in Cisco's proprietary code. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of sampling other nodes to determine a new potential parent mesh node for connecting with said at least one mesh node in the same way by using radios to regularly scan for and broadcast potential new parents in an RF environment, while maintaining the traffic flow to the old parent before transfer to a new potential parent, and achieves the same result of ensuring the seamless transfer of data flow and a good user experience without degrading performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks solicit and the repeater learns all of the available neighbors back to a gateway to determine the best path while ensuring no packets are lost in switching to a new parent or gateway. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.d. wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.	Cisco's Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network comprise a mesh network wherein sampling times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner. As discussed above, each Meraki AP periodically sends out mesh probes and packets as broadcast frames. **Mesh Probes** Each Meraki access point sends out link probe packets (known as mesh probes) at different bit rates and varying sizes. Because these packets are sent as broadcast frames, no ACK frames are needed from receiving stations. Four different types of probes at different data rates are sent in a batch of 15 seconds on both (2.4 /5 GHz) bands. All access points listen to the mesh probes and depending on the number of mesh probes correctly received, come up with a link quality metric as shown in dashboard. Source: Wireless Mesh Networking

	In a Cisco Meraki Wireless Mesh Network (802.11 Wi-Fi Mesh), gateway or parent mesh nodes send out beacon frames and other control frames to repeaters or children to manage how resources are allocated between the parent and child nodes. These frames provide essential information including network parameters (SSID, supported rates, channel information, etc.), parent node status (whether the node is available to accept new child nodes), and link quality (such as signal strength), etc. This could include information elements with Action frames, like Channel Switch Announcement, that can be used to tell a child to go off-channel and scan, which is not a part of the beacon frame. Other control frames, *e.g.*, trigger frames and PS-Poll frames may be used to coordinate or schedule a child/client to go off-channel. The time interval between two consecutive beacon frames is called the beacon interval and is 100ms by default. Source: https://documentation.meraki.com/MR/Wireless_Troubleshooting/Analyzing_Wireless_Packet_Captures On information and belief, the common current parent sends tokens with, or in, the beacon messages (or other control frames) that coordinate the sampling times in a round robin manner, such that each of the children can sample at appropriate times and processing efficiencies are maintained. The implementation of such tokens is enabled in Cisco's proprietary code and may take the form of flags, contention free notices, request/response messages, or other notifications or permissions. On information and belief, Cisco's proprietary mesh network operates the same or similar to Target Wake Time ("TWT") in Wi-Fi 6. TWT is a feature in Wi-Fi 6 that allows devices to negotiate specific times to wake up and communicate with the access point, optimizing power consumption and improving network efficiency. Access points negotiate a "TWT agreement" specifying when the device will send or receive data. Source: Unlocking High Efficiency with Wi-Fi 6 | The Meraki Blog If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of seamless parent node or gateway switching in a mesh network in the same way by coordinating sampling times with periodic transmissions and exchanges of broadcast packets (which may be set by or sent at the beacon interval), and achieves the same result of ensuring the seamless transfer of data flow and a good user experience without degrading performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks achieve coordinated and seamless transfer of data flows during parent switching. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 3	
3.pre. A mesh network comprising: a plurality of mesh nodes;	Cisco's Meraki Wireless Mesh Network comprises a mesh network having a plurality of mesh nodes (*e.g.*, Cisco Meraki and Catalyst CW Access Points). The Meraki cloud-managed wireless series is a line of enterprise-grade WLAN access points ("APs") that leverage the Cisco Meraki cloud-managed architecture. Both Cisco Meraki and Catalyst CW Access Points support Wi-Fi capabilities and can be paired with the Meraki cloud platform. Meraki AP models start with MR, *e.g.* MR33, MR36, MR46E, and the Catalyst APs with CW, *e.g.* CW9162, CW9164, CW9166, are dual-mode APs which can be used with the Catalyst 9800 controller or the Meraki cloud architecture. Source: Meraki Family Datasheet; Meraki Wi-Fi Models  Source: Wireless Mesh Networking; *see also* Meraki Best Practices - Wireless Mesh Networking All Meraki APs support wireless meshing without any additional configuration ("zero configuration") or licensing. 

	Source: Mesh Deployment Guide  Source: Meraki Family Datasheet  Source: Meraki Family Datasheet
3.a. wherein each node within said plurality of nodes comprises at least	Cisco Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network comprise a mesh network wherein each node within said plurality of nodes comprises at least three radios.

three radios further including:	As shown *supra* in claim 2, Cisco Meraki and Catalyst CW Access Points comprise at least three radios. For example, Cisco Meraki MR86 access point is a cloud-managed 4x4:4 802.11ax access point and has concurrent 2.4GHz and 5GHz radios. A dedicated third radio provides RF optimization, while a fourth delivers Bluetooth scanning and beaconing. Each of the 2.4GHz and 5GHz bands have several different non-interfering channels for radio operation. For instance, channels 1, 6 and 11 on the 2.4 GHz band do not interfere. The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands, and there are 24 non-overlapping channels, offering more flexibility and less congestion.  Source: MR86 Datasheet APs in the Cisco Meraki Wireless Mesh Network have multiple radios, including 2.4 and 5 GHz, at different channels and frequencies. Both radios may be used concurrently for meshing while also serving wireless clients.  Source: Wireless Mesh Networking

Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: Mesh Deployment Guide

Extending the LAN with a Wireless Mesh Link

Cisco Meraki MR access points (APs) can operate as mesh repeaters, which allows them to extend the wireless network range off of a limited number of gateway APs. Since repeaters also support wired clients plugged into their wired interface, a repeater can be used to bridge a remote LAN segment back to the main network.

Wireless Mesh repeaters can back-haul to access points acting as wireless gateways and:

- Utilize the same radio as used for mesh to serve wireless clients
- Utilize non-meshing radios to serve wireless clients
- Utilize wired uplink for use cases described below

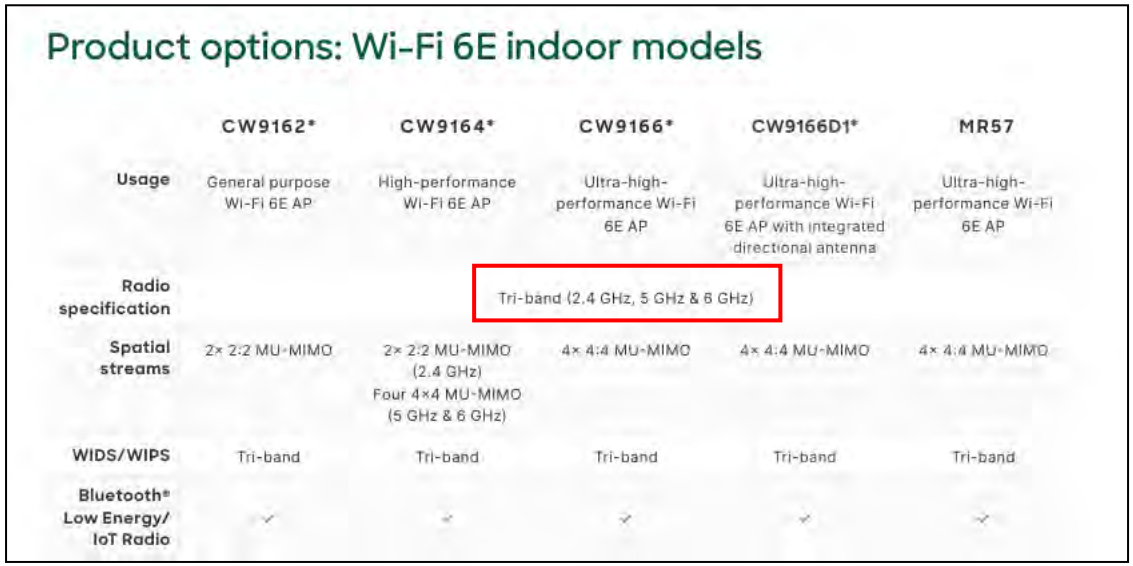
Source: Extending the LAN with a Wireless Mesh Link

Meraki indoor access points also feature an additional radio dedicated to full-time scanning.

Indoor APs feature a radio dedicated to full-time scanning, rogue AP containment, and automatic RF optimization. With Air Marshal, it's possible to set up a real-time wireless intrusion detection and prevention system (WIDS/WIPS) with user-defined threat remediation policies and intrusion alarms, enabling secure wireless environments without complex setup or systems integration. Auto RF eliminates the need for manual RF configuration by scanning the environment for utilization, interference, and other metrics, and computing the optimal channel and power settings for every AP in the network. Meraki WLANs are fully HIPAA and PCI compliant.

Source: Meraki Family Datasheet

Additional radios are tri-band, with a scanning radio, and each radio is operating on a different frequency and channel.



Product options: Wi-Fi 6E indoor models

	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: Meraki Family Datasheet

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed element are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
3.b. a first relay radio operating on a first RF channel at a first point in time and dedicated to uplink connections to a single current parent node; a second relay radio operating on a second RF channel at a first point in time dedicated to downlink connections to zero or more child nodes; and	Cisco's Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network comprise a mesh network with a first relay radio operating on a first RF channel at a first point in time and dedicated to uplink connections to a single current parent node; a second relay radio operating on a second RF channel at a first point in time dedicated to downlink connections to zero or more child nodes. In Cisco's Meraki Wireless Mesh Network, APs can be in one of two states: gateway ("parent") or repeater ("parent" or "child"). Gateway APs are connected directly to the wired network, granting it an uplink to the Internet, while repeaters connect to the best gateway (and path through other repeaters) using the Meraki Mesh Algorithm. **Mesh Network Components** In a mesh network, access points can be in one of two states: Gateway, or Repeater. **Gateways** Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

Meraki Mesh Algorithm

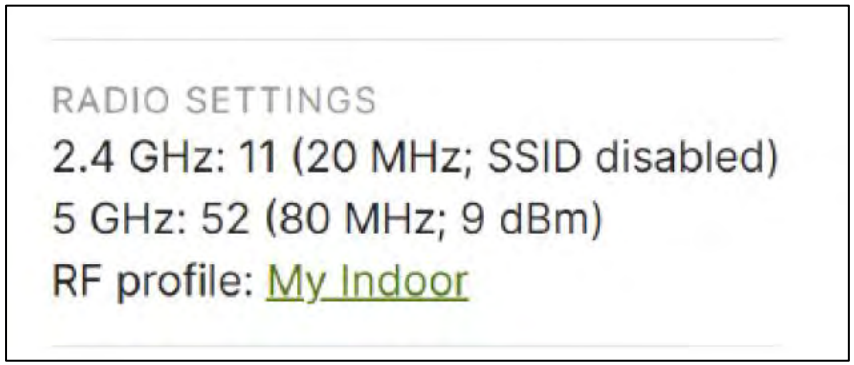
Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

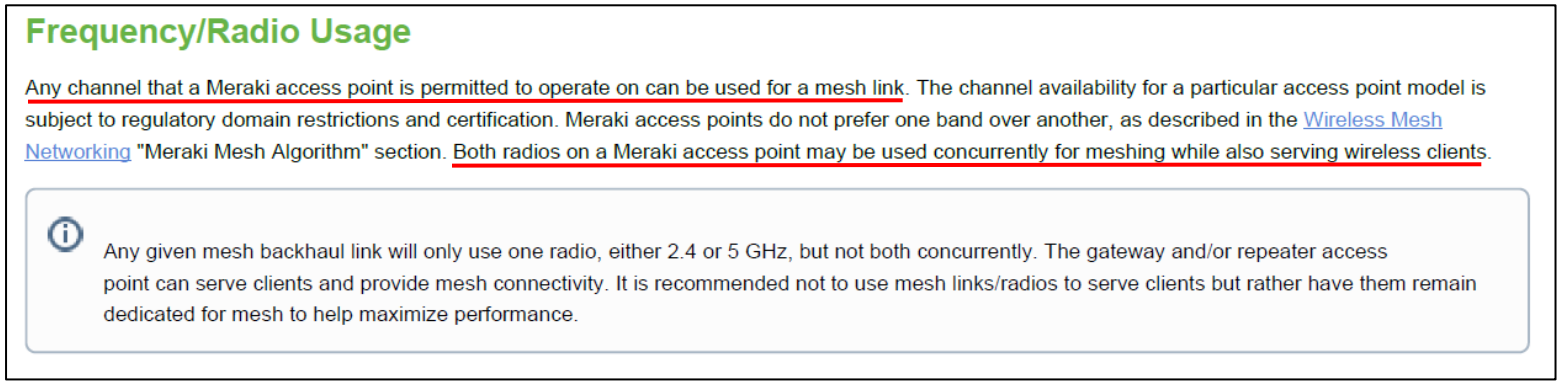
Source: Wireless Mesh Networking

Cisco Meraki and Catalyst CW Access Points support multiple radios. Cisco's access point radio(s) can be a "dedicated" when being used for the purpose of uplink to the parent node or downlink to zero or more child nodes. Different channels are available for operating the uplink connection to a parent or gateway and downlink connections to zero or more child nodes or repeaters. For example, any channel that a Meraki access point is permitted to operate on can be used for a mesh link (2.4 GHz + 5 GHz).

APs in the Cisco Meraki Wireless Mesh Network have multiple radio interfaces, including 2.4 and 5 GHz. Both radios may be used concurrently for meshing (a dedicated uplink to a parent or gateway and downlink to repeater or child) while also serving wireless clients.



Source: Wireless Mesh Networking

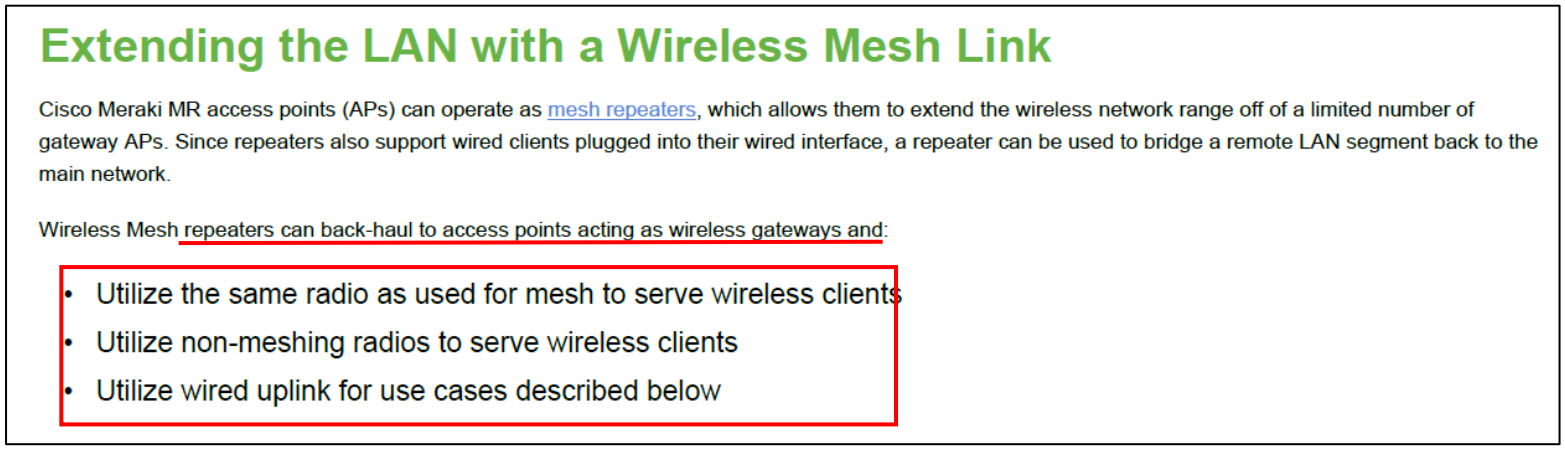


Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: Mesh Deployment Guide



Extending the LAN with a Wireless Mesh Link

Cisco Meraki MR access points (APs) can operate as mesh repeaters, which allows them to extend the wireless network range off of a limited number of gateway APs. Since repeaters also support wired clients plugged into their wired interface, a repeater can be used to bridge a remote LAN segment back to the main network.

Wireless Mesh repeaters can back-haul to access points acting as wireless gateways and:

- Utilize the same radio as used for mesh to serve wireless clients
- Utilize non-meshing radios to serve wireless clients
- Utilize wired uplink for use cases described below

Source: Extending the LAN with a Wireless Mesh Link

If these claim elements are not literally infringed, they are infringed under the doctrine of equivalents. Cisco's Meraki and Catalyst CW Access Points perform the same function of determining a new potential parent mesh node for connecting with at least one mesh node in the same way by using radios to regularly scan for and track potential new parents in an RF environment, with radios operating at different frequencies and channels, and achieves the same result by determining the best parent mesh node for optimal network performance. Moreover, any differences between the functionality of the Cisco Accused Products and the claimed elements are insignificant or insubstantial because Cisco's mesh networks deploy multiple dual-radio (or tri-radio) mesh access points to find the best path through the network for data traffic by scanning for new parents. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
3.c. a dedicated scanning radio operating on a third RF channel at a first point in time and configured to scan a Radio Frequency (RF) environment to discover new potential parent mesh nodes for connection with said first relay radio;	*See supra* discussion for claim elements 2.a and 3.a, the contents of which are incorporated herein by reference.
3.d. said first, second, and third RF channels are different from each other	*See supra* discussion for claim elements 2.a, 2.b, 3.a, and 3.b, the contents of which are incorporated herein by reference.
3.e. wherein after said dedicated scanning radio discovers a potential new parent node, said first relay radio samples the RF link to said potential new parent node using its uplink radio and concurrent with said sampling, packets to be sent to the mesh node from its current parent node are buffered by the current parent node, and packets to be sent from the mesh node to its current parent node are buffered by the mesh node wherein sampling	Cisco's Meraki and Catalyst CW Access Points in the Cisco Meraki Wireless Mesh Network comprise a mesh network wherein after said dedicated scanning radio discovers a potential new parent node, said first relay radio samples the RF link to said potential new parent node using its uplink radio and concurrent with said sampling, packets to be sent to the mesh node from its current parent node are buffered by the current parent node, and packets to be sent from the mesh nod e to its current parent node are buffered by the mesh node wherein sampling times are coordinated among multiple mesh nodes having a common parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner. Access points in the Meraki mesh will detect other APs and determine the best path (*i.e.*, "parent" repeater or gateway). **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

times are coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner.	Source: Wireless Mesh Networking As shown *supra* in claim elements 2.a-2.d (which are incorporated herein by reference), a dedicated scanning radio discovers neighbors and potential new parent nodes or gateways. The repeater will sample the potential parent node by sending mesh probes using its uplink radio (*i.e.*, to select the best route to the gateway). These mesh probes are used to come up with a link quality metric. **Mesh Probes** Each Meraki access point sends out link probe packets (known as mesh probes) at different bit rates and varying sizes. Because these packets are sent as broadcast frames, no ACK frames are needed from receiving stations. Four different types of probes at different data rates are sent in a batch of 15 seconds on both (2.4 /5 GHz) bands. All access points listen to the mesh probes and depending on the number of mesh probes correctly received, come up with a link quality metric as shown in dashboard. Source: Wireless Mesh Networking **Mesh Convergence Time** The Meraki mesh algorithm is designed for stationary Meraki access points with variable links. The algorithm sends out periodic discovery frames of varying sizes on the 2.4 GHz and 5 GHz data radios to discover mesh neighbors. More discussion on neighbor discovery and gateway selection can be found in Wireless Mesh Networking. After one neighbor is selected as the gateway, new routes are passively identified on the same operating channel. A mesh link may take a couple minutes to establish and may not always be used for every data flow if a new route with better link metrics becomes available. Source: Mesh Deployment Guide As shown *supra* in claim elements 2.a-2.d (which are incorporated herein by reference), concurrent with this sampling, packets to the mesh node or repeater from its current gateway or parent are buffered at the parent node, and packets to be sent from the mesh node or repeater to its current parent node or gateway are buffered by the mesh node. Sampling times are also coordinated among multiple mesh nodes having a common current parent node whereby the common current parent node sends tokens to each of its children in a round-robin manner. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Exhibit 5 – U.S. Patent No. 7,885,243 (Cisco Wireless Mesh Networks with Catalyst and Aironet Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Wireless Mesh Network, including but not limited to Cisco's Catalyst and Aironet Access Points with related controllers and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1-13 (the "Asserted Claims") of U.S. Patent No. 7,885,243 (the "'243 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '243 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringed by testing, setting up, and using Cisco Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Wireless Mesh Network in the United States.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the '243 Patent and make and use the infringing Cisco Wireless Mesh Network. Specifically, Cisco's Wireless Mesh Access Points, Design and Deployment Guides, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco's website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-13 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Wireless Mesh Network, and the infringing configuration and operation of the controller, access points, and software that comprise the infringing Cisco Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the '243 Patent since 2009 when Cisco executives

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g.*, https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

met with MeshDynamics to discuss partnering opportunities and the application that would result in the ’243 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’243 Patent, or in the alternative was willfully blind to the ’243 Patent and the infringement, following the MeshDynamics meeting in 2009.

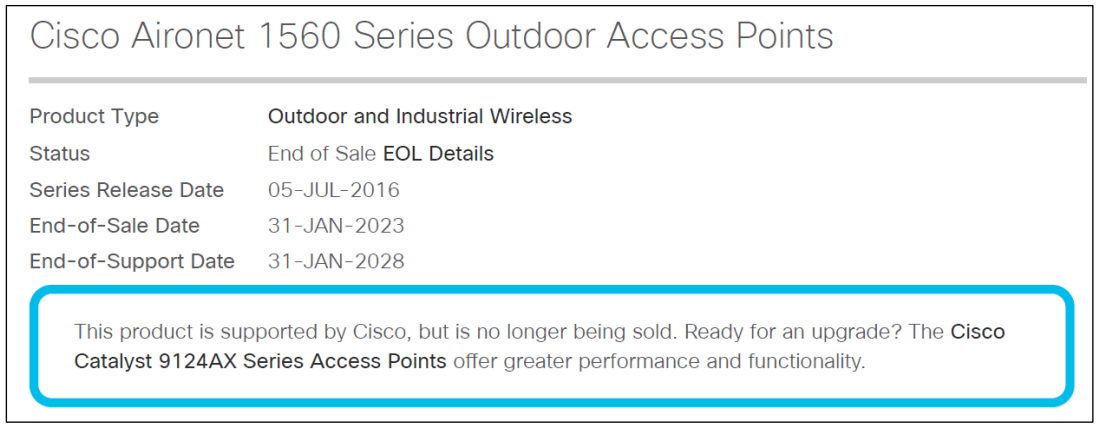
Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’243 Patent. Cisco Accused Products come enabled as a mesh access point and automatically connect to the controller to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particular suited for Cisco’s Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’243 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the ’243 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics’s preliminary infringement analysis of the Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts presently known to it. To the best of MeshDynamics’s knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco's Wireless Mesh Network with Catalyst and Aironet Access Points
(1pre) A wireless mesh network comprising:	Cisco's Wireless Mesh Network (a component of the Cisco Unified Wireless Network solution) comprises Wi-Fi relay or mesh nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Access Points) and additional components, including controllers and software. Cisco Catalyst 9100 and Aironet Access Points support Wi-Fi capabilities. The Catalyst series is a follow-up from the Cisco Aironet series access points, and the Aironet and Catalyst Access Points contain the same mesh networking capabilities, as shown below in Cisco's documentation and guides that apply to both product lines: Q. Which Cisco access points support Wi-Fi 6? A. Cisco Catalyst 9100 enterprise-class access points support Wi-Fi 6 capabilities, extending the power of intent-based networking with hardware and software innovations. The Cisco Catalyst 9100 APs offer resiliency and security while enabling intelligence at the network edge. Q. What is the portfolio transition from the Cisco Aironet Series to the Cisco Catalyst 9100 platform? A. The Cisco Catalyst 9100 Access Points consist of the Cisco Catalyst 9105, 9115, 9120, 9130, 9136, 9162, 9164, 9166 and 9166D1, and outdoor 9163E Series. These access points are the follow-ups to the Cisco Aironet 1815, 1850, 2800, and 3800 Series Access Points, respectively. The Cisco Catalyst 9100 Access Points come equipped with Wi-Fi 6/6E capabilities, a better industrial design, and improved RF performance, and deliver reliability, security, and intelligence at scale. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html



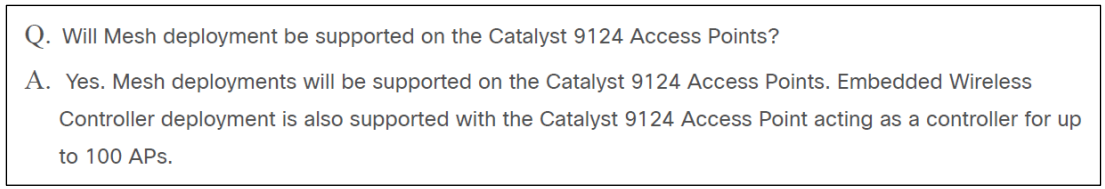
Cisco Aironet 1560 Series Outdoor Access Points

Product Type	Outdoor and Industrial Wireless
Status	End of Sale EOL Details
Series Release Date	05-JUL-2016
End-of-Sale Date	31-JAN-2023
End-of-Support Date	31-JAN-2028

This product is supported by Cisco, but is no longer being sold. Ready for an upgrade? The **Cisco Catalyst 9124AX Series Access Points** offer greater performance and functionality.

Source: https://www.cisco.com/c/en/us/support/wireless/aironet-1560-series/series.html

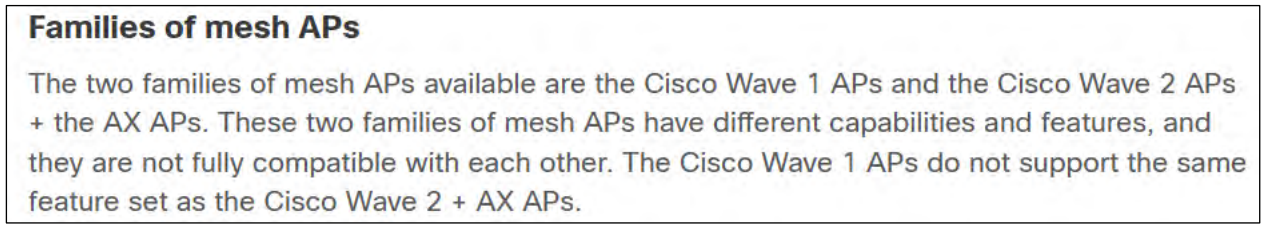
In addition to the Wave 1 and Wave 2 access points, mesh deployments are supported on the Cisco Catalyst 9100 Access Points, including, *e.g.*, Catalyst 9124 Access Points, Catalyst 9130 Access Points and "AX" access points:



Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points?

A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html



Families of mesh APs

The two families of mesh APs available are the Cisco Wave 1 APs and the Cisco Wave 2 APs + the AX APs. These two families of mesh APs have different capabilities and features, and they are not fully compatible with each other. The Cisco Wave 1 APs do not support the same feature set as the Cisco Wave 2 + AX APs.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html; see also https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg/mesh-access-points.html

From this release, mesh support is included in the Cisco Catalyst 9130AX Series Access Points. All traditional capabilities of mesh are included in the Cisco Catalyst 9130AX Series APs operating in Cisco IOS XE Dublin 17.12.1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html

Cisco Catalyst 9100 and Aironet Access Points can be deployed in a Cisco Wireless Mesh Network with a controller and software, such as the Catalyst 9800 Controller and Cisco DNA and Prime Infrastructure:

Mesh Network Components

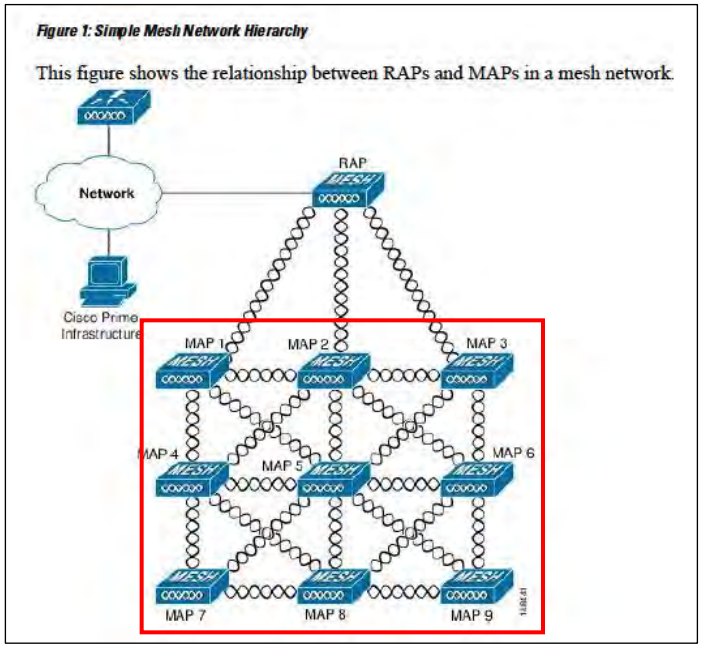
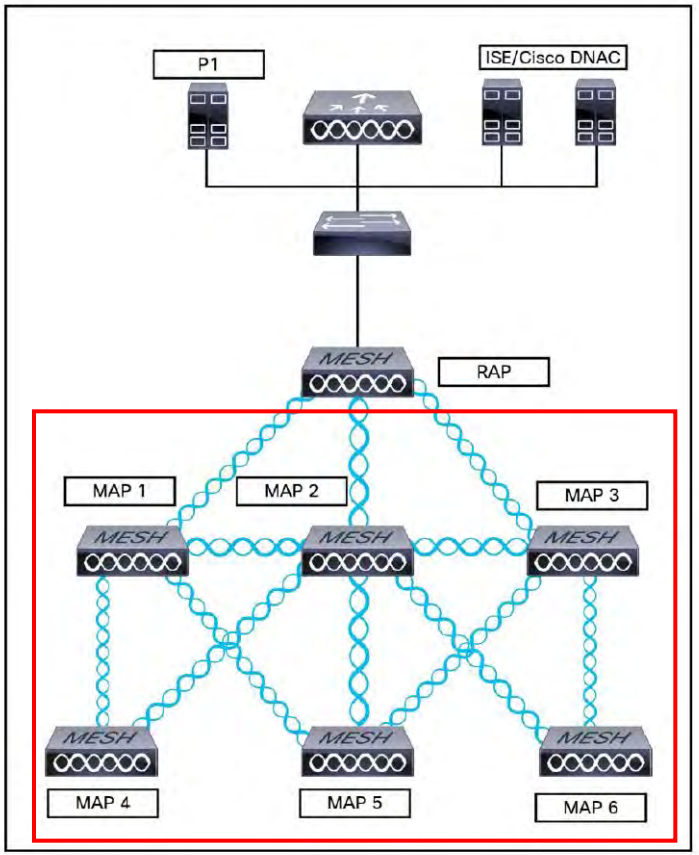
This section describes the mesh network components.

The Cisco wireless mesh network has four core components:

- Cisco C9800 and IOS-XE 17.1
- Cisco Aironet and Catalyst series access points
- Cisco PI and Cisco DNA Center
- Mesh software architecture

Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network.

- Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series.
- Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series.

	Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh Cisco's Wireless Mesh Network comprises one or more, or a plurality of, access points, *e.g.*, Cisco Catalyst 9100 and Aironet Access Points, as shown in red below.   Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
(1a) an access server wherein the access server sets one or more functioning parameters of	Cisco's Wireless Mesh Network comprises an access server (*e.g.*, Catalyst 9800 series wireless controllers or other Cisco Wireless LAN Controllers (WLCs)) wherein the access server sets one or more functioning parameters of the wireless mesh network.

the wireless mesh network;	**Mesh network components** The Cisco wireless mesh network has four core components: • Cisco Catalyst 9800 Series and Cisco IOS XE. • Cisco Aironet and Catalyst access points. • Cisco Prime Infrastructure and Cisco DNA Center. • Mesh software architecture. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh **Mesh access point roles** Mesh networking employs Cisco Aironet outdoor mesh APs and indoor mesh APs along with the Cisco WLC and Cisco Prime Infrastructure and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh Embedded wireless controller (*e.g.*, access server) mesh deployment is also supported in the Catalyst 9100 Access Points: Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points? A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html The access server provides the framework for access points to communicate and choose optimal paths, including setting the functioning parameters of the wireless mesh network. For example, Cisco sets or recommends the number of hops or number of mesh access points (MAPs) per root access point (RAP).

Note With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

- Number of backhaul hops is limited to eight but we recommend three to four hops.

  The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.

- Number of MAPs per RAP.

  There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

The controller can also configure mesh parameters, including setting global mesh parameters such as the maximum range between RAP and MAPs, and setting the data rate.

Configuring Global Mesh Parameters

This section provides instructions to configure the mesh access point to establish a connection with the controller including:

- Setting the maximum range between RAP and MAP (not applicable to indoor MAPs).
- Enabling a backhaul to carry client traffic.
- Defining if VLAN tags are forwarded or not.
- Defining the authentication mode (EAP or PSK) and method (local or external) for mesh access points including security settings (local and external authentication).

You can configure the necessary mesh parameters using either the GUI or the CLI. All parameters are applied globally.

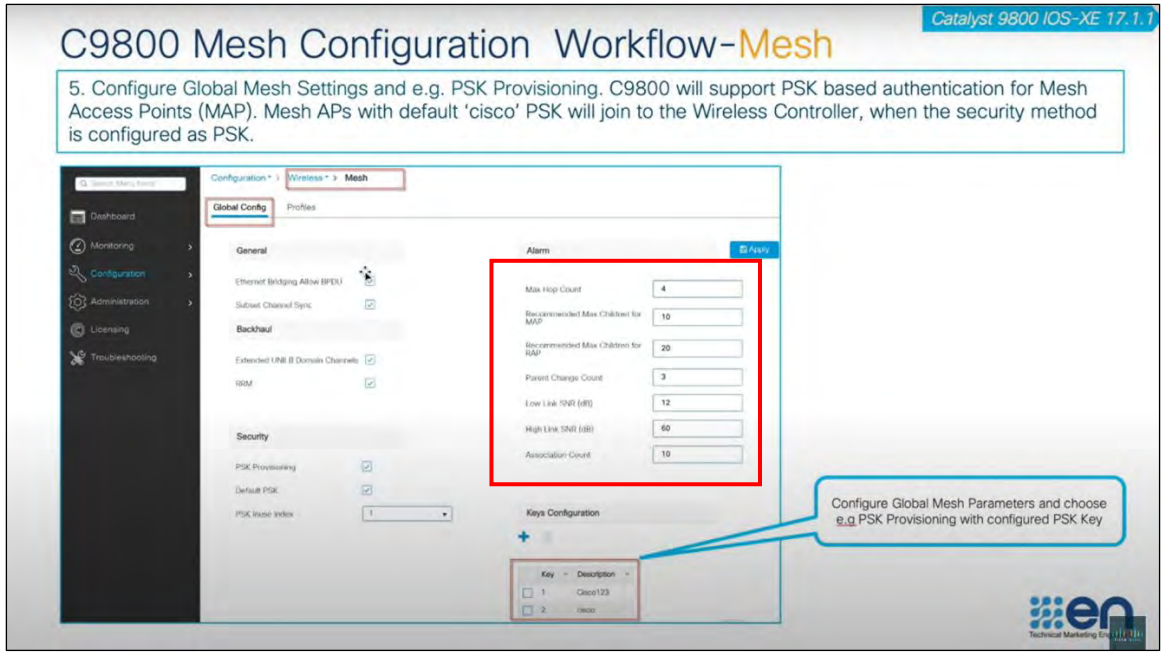
Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/connecting-mesh-aps.html

Configuring Wireless Backhaul Data Rate

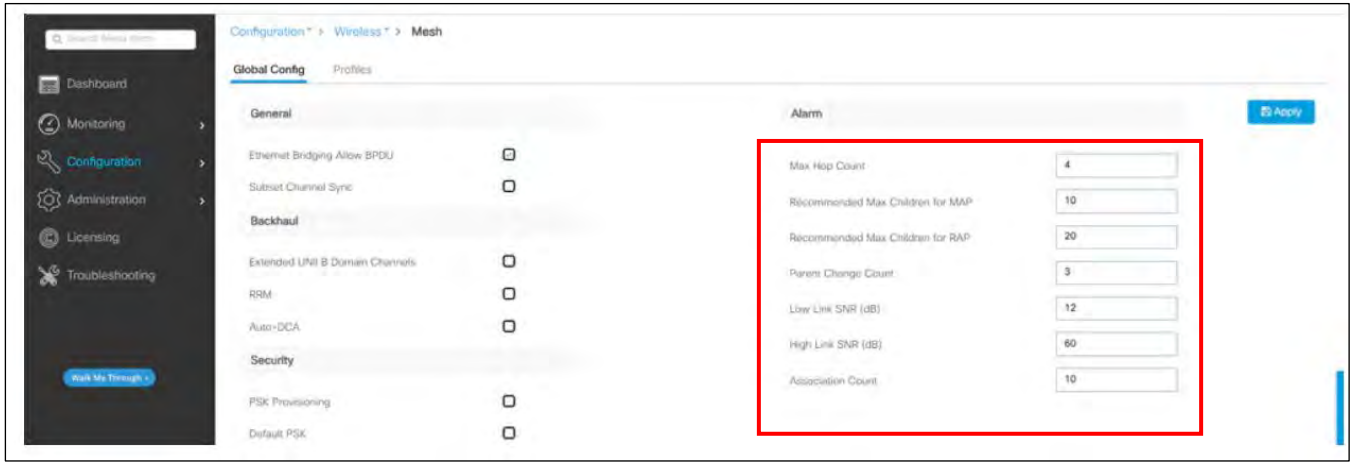
Backhaul is used to create only the wireless connection between the access points. The backhaul interface vary between 802.11a/n/ac/ax rates depending upon the access point. The rate selection is important for effective use of the available RF spectrum. The rate can also affect the throughput of client devices, and throughput is an important metric used by industry publications to evaluate vendor devices.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-10/config-guide/b_cg810.pdf

The setting of the global parameters at the controller, including max hop count and low and high link signal-to-noise ratio ("SNR") are shown in the screenshots below:



Source: https://www.youtube.com/watch?v=ESNFzlGTueU

	 Source: https://www.cisco.com/c/en/us/support/docs/wireless/embedded-wireless-controller-on-catalyst-9124ax-access-points/221557-configure-point-to-point-mesh-link-with.html If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of establishing one or more functioning parameters of the wireless mesh network in the same way by using a controller to set recommended or global parameters for the wireless mesh network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. To the extent Cisco argues the controller is not the claimed "access server," for example, a person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1b) one or more root nodes connected to said access server and an external network;	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprise one or more root nodes (*e.g.*, "Root APs" or RAPs) connected to said access server (*e.g.*, Cisco WLC) and an external network. Cisco Catalyst 9100 and Aironet Access Points are RAPs and MAPs in the Cisco Wireless Mesh Network. The RAPs have a wired connection to the controller and external network via the controller, and the MAPs have wireless connections to the controller via the RAP or another MAP.

For Mesh mode, an AP should be configured with Bridge mode. Access points within a mesh network operate in one of the following two ways:

- Root AP (RAP)
- Mesh AP (MAP)

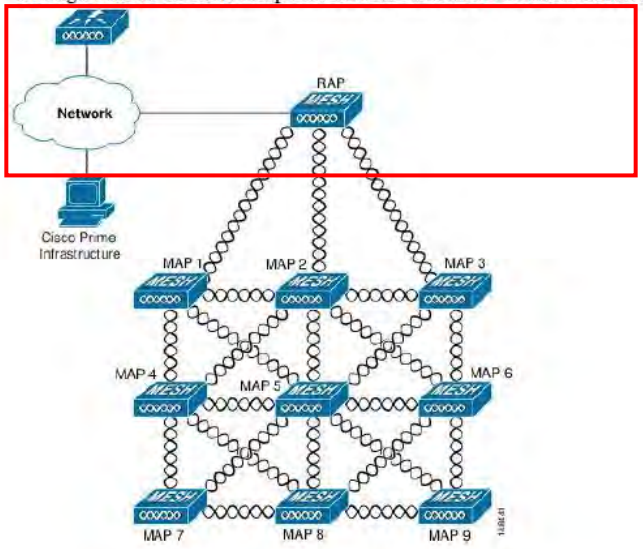
Note: All APs are configured and shipped as mesh APs. To use an APs as a RAP, you must reconfigure the mesh AP as a RAP. In all mesh networks, ensure that there is at least one RAP.

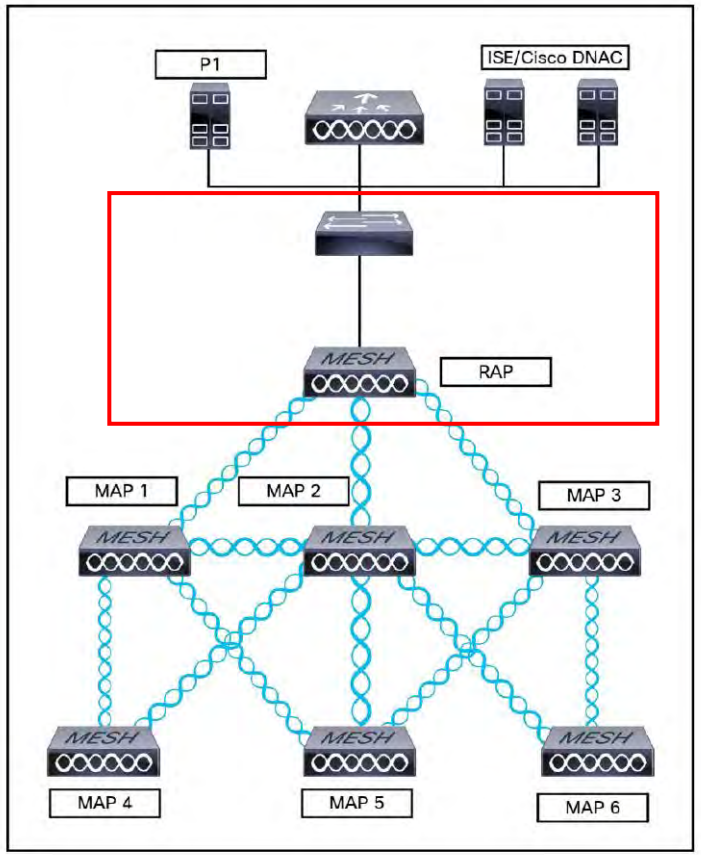
While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Figure 1: Simple Mesh Network Hierarchy

This figure shows the relationship between RAPs and MAPs in a mesh network.





	Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh In Cisco's Wireless Mesh Network, Ethernet bridging also allows RAPs to connect to wired networks while also acting as a bridge for wireless traffic from MAPs. Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-4/b_mesh_84/Mesh_Deployment_Modes.pdf If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of connecting the mesh network to the external network in the same way by using a root access point and achieves the same result of creating a wireless mesh network. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The RAP connects the wireless mesh network to the wired network, providing access to the external network, such as the Internet. The RAP is directly connected via a wired connection to a Cisco WLC. This controller is typically located within the wired network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused product to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1c) one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprise one or more AP nodes (*e.g.*, MAPs as shown above) wherein each AP node is in wireless two-way data communication with an associated parent node (*e.g.*, RAPs or MAPs) wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria (*e.g.*, an ease calculation criteria) or the associated parent node is a root node (*e.g.*, RAP) within wireless communication range. The Cisco Catalyst 9100 and Aironet Access Points (*e.g.*, MAPs) are in wireless two-way data communication with a parent node. A parent can be either the RAP itself or another MAP. MAPs communicate with each other and the RAP using wireless 802.11 backhaul channels, and 802.11 request/response protocol.

associated parent node is an available parent node meeting one or more communication criteria or the associated parent node is a root node within wireless communication range;	For Mesh mode, an AP should be configured with Bridge mode. Access points within a mesh network operate in one of the following two ways: • Root AP (RAP) • Mesh AP (MAP) **Note:** All APs are configured and shipped as mesh APs. To use an APs as a RAP, you must reconfigure the mesh AP as a RAP. In all mesh networks, ensure that there is at least one RAP. While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. **Mesh neighbors, parents, and children** Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73). • A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP. ◦ Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected. • A child AP selects the parent AP as its best route back to the RAP. • A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh The parent node is selected from all available parent nodes using Adaptive Wireless Path Protocol (AWPP) and wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria (*e.g.*, AWPP ease calculation criteria) or the associated parent node is a root node (*e.g.*, RAP) within wireless communication range. Cisco Catalyst 9100 and Aironet Access Points in Cisco Wireless Mesh Networks use AWPP to determine the best path through the mesh access points back to the RAP and controller. This includes selection of parents, establishing a parent-child relationship, and meeting one or more communication criteria to choose the best parent (*e.g.*, AWPP ease calculation criteria), or the associated parent node is a RAP within wireless communication range.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

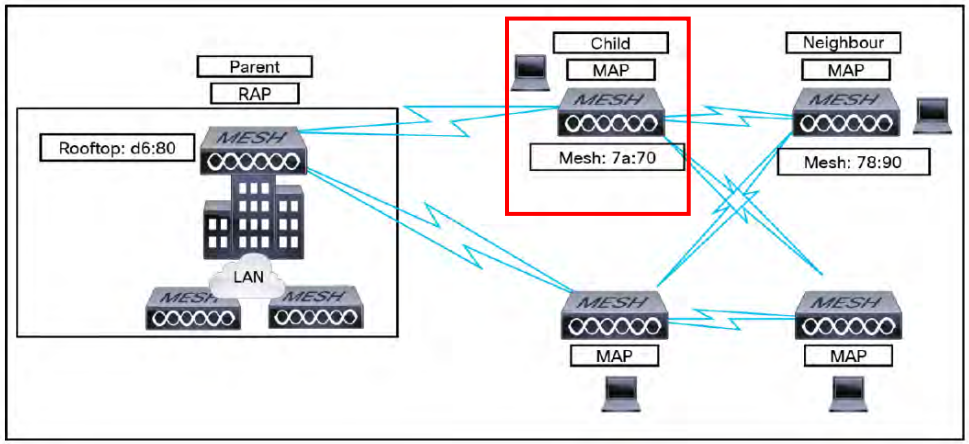
- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

	Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process uses criteria to select the best parent. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1d) wherein an AP node is in wireless communication with zero or more clients; and	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprise wherein an AP node (*e.g.*, MAP) is in wireless communication with zero or more clients (*e.g.*, laptops) as shown in Cisco's documentation:  Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
(1e) wherein an AP node includes a means for switching two-way data	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein an AP node includes a means for switching two-way data communication (*e.g.*, through an Adaptive Wireless Path Protocol (AWPP) process) from a first associated parent node to a second associated parent node based on the functioning parameters of

communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets;	the wireless mesh network and wherein an AP node contains one or more datasets (*e.g.*, a dataset comprising "the best path back to a RAP," using AWPP and soliciting neighbors to learn and determine the best path, and additional datasets identified *infra*). Catalyst 9100 and Aironet Access Points in Cisco Wireless Mesh Networks use AWPP to determine the best path through the mesh access points to the controller and switch the two-way data communications from an associated parent to another associated parent. This algorithm is run at start up and whenever a parent is lost or does not exist; and includes selection of parents based on functioning parameters of the network, including the number of hops and SNR. AWPP on the access points also continuously monitors conditions and changes routes to reflect changes in conditions. **Adaptive Wireless Path Protocol** The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

As confirmed in Cisco's documentation, and on information and belief as shown in Cisco's proprietary code, a list of potential parents/neighbors (*e.g.*, dataset) is obtained and maintained on each MAP, *e.g.*, from background scanning or "AWPP – Neighbor Delivery Request/Response (NDReq/NDResp)" messages.

Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Cisco's command line manual and mesh guide further show that the access points have a dataset showing what is surrounding the node, including neighbors and child node(s) identified with "Child" as the State field:

show wireless mesh neighbor

To see the neighbors of all connected mesh APs, use the **show wireless mesh neighbor** command.

show wireless mesh neighbor [**chassis** {*chassis-number* | **active** | **standby**} **R0**]

```
AP Name/Radio             Channel   Rate      Link-snr   Flags   State
-------------------------------------------------------------------------------

AP Name : Mesh-AP01

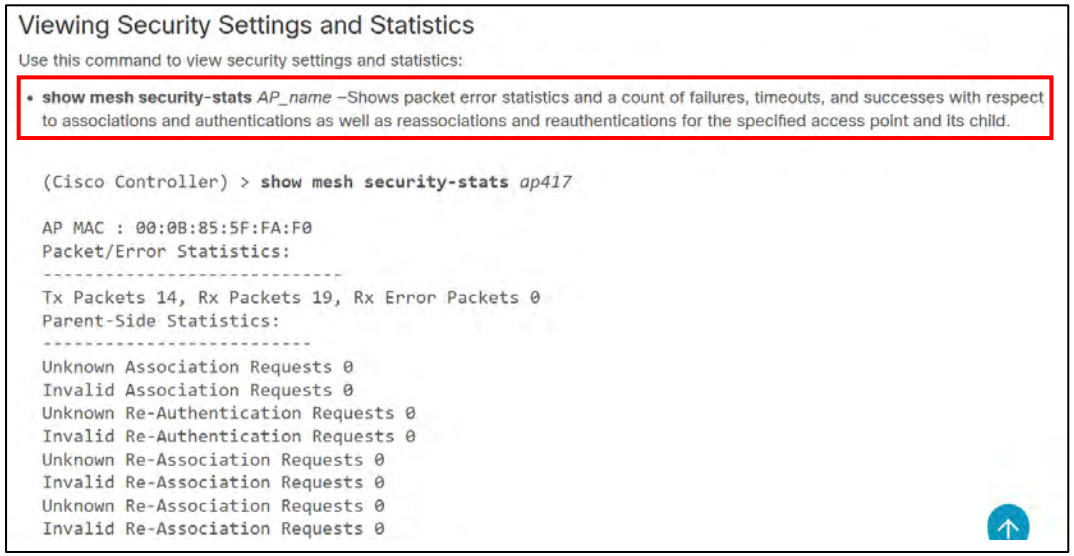
54:9f:c6:fa:5c:71         149       auto      0          40
b0:c5:3c:e5:d9:71         149       auto      22         49      UPDATED NEIGH
e8:eb:34:d5:88:d1         149       auto      0          40
e8:eb:34:d5:8d:d1         149       auto      18         49      UPDATED CHILD
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/cmd-ref/b_wl_16_10_cr/show-commands.html#wp3747685901

- **show mesh neighbor summary** –Displays summary information about mesh neighbors. Neighbor information includes MAC addresses, parent-child relationships, and uplink and downlink (SNRUp, SNRDown).

```
(Cisco Controller) > show mesh neighbor summary ap1500:62:39:70
AP Name/Radio Mac Channel Snr-Up Snr-Down Link-Snr Flags  State
mesh-45-rap1      165     15     18       16       0x86b  UPDATED NEIGH PARENT BEACON
00:0B:85:80:ED:D0 149     5      6        5        0x1a60 NEED UPDATE BEACON DEFAULT
00:17:94:FE:C3:5F 149     7      0        0        0x860  BEACON
```

Note After review of the **show mesh** commands above, you should be able to see the relationships between the nodes of your network and verify the RF connectivity by seeing the SNR values for every link.



Viewing Security Settings and Statistics

Use this command to view security settings and statistics:

- **show mesh security-stats** *AP_name* –Shows packet error statistics and a count of failures, timeouts, and successes with respect to associations and authentications as well as reassociations and reauthentications for the specified access point and its child.

```
(Cisco Controller) > show mesh security-stats ap417

AP MAC : 00:0B:85:5F:FA:F0
Packet/Error Statistics:
----------------------------
Tx Packets 14, Rx Packets 19, Rx Error Packets 0
Parent-Side Statistics:
-------------------------
Unknown Association Requests 0
Invalid Association Requests 0
Unknown Re-Authentication Requests 0
Invalid Re-Authentication Requests 0
Unknown Re-Association Requests 0
Invalid Re-Association Requests 0
Unknown Re-Association Requests 0
Invalid Re-Association Requests 0
```

```
Child-Side Statistics:
--------------------------
Association Failures 0
Association Timeouts 0
Association Successes 0
Authentication Failures 0
Authentication Timeouts 0
Authentication Successes 0
Re-Association Failures 0
Re-Association Timeouts 0
Re-Association Successes 0
Re-Authentication Failures 0
Re-Authentication Timeouts 0
Re-Authentication Successes 0
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-0/design/guide/mesh80/m_checking-the-health-of-the-network.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using parameters or criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and background scanning both assist the AP in maintaining a list of neighbors and switching to the best parent on-channel or off-channel.

	A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. To the extent this is considered an element under 35 U.S.C. §112 (¶ 6), in accordance with P.R. 3-1(c), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function include, *e.g.*, the software algorithms in Cisco's proprietary AWPP protocol and mesh background scanning for parent selection identified above. MeshDynamics's investigation into Cisco's proprietary algorithms remains ongoing and MeshDynamics will supplement or amend its P.R. 3-1(c) identification following discovery and/or claim construction. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1f) wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset;	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein the communication criteria (*e.g.*, AWPP ease calculation criteria) comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset (*e.g.*, a dataset comprising "the best path back to a RAP," which was received from a parent mesh access point). AWPP is designed specifically for wireless mesh networking and to dynamically find the best path back to the RAP. AWPP automatically defines the best path back to the RAP ("root node") by calculating the cost of each path in terms of the signal strength and number of hops. AWPP uses ease to determine the best path back to the RAP; thus, AWPP's communications criteria comprises instructions to select the parent node that is associated with a RAP or if the RAP is contained in the dataset comprising the best path back to the RAP.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

	Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh Cisco's Wireless Mesh Network can also be configured as a static mesh configuration and the criteria and route path dataset may be provided by manually assigning parent-child and RAP relationships. Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria dynamically and automatically find the best path back to the RAP. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1g) wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node;	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein one of the datasets contained in an AP node (*e.g.*, a dataset comprising "the best path back to a RAP") comprises a route path dataset comprising an identifier (*e.g.*, MAC address) for the associated parent node appended to the route path dataset for the associated parent node. AWPP determines the best path or route through the mesh access points back to the RAP and controller. A parent provides the best path back to the RAP. This includes selection of parents and establishing a parent-child relationship. The MAP learns of all the available neighbors back to a RAP and as discussed above, obtains and maintains a list of potential parents/neighbors (*e.g.*, dataset) on each MAP. Each RAP or MAP is identified by a MAC address. The routing information contains the MAC address of each hop to the RAP.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Loop Prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh access point can easily detect and discard routes that loop.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

AWPP is described as a distance-vector database. On information and belief, as shown in Cisco's proprietary code, one of the datasets maintained on the AP is a route path dataset, *i.e.*, the best path back to the RAP, with the MAC address for the associated parent.

Source: https://www.cisco.com/c/en/us/support/docs/wireless/embedded-wireless-controller-on-catalyst-9124ax-access-points/221557-configure-point-to-point-mesh-link-with.html

Cisco's command line manual further shows that the access points have a dataset with the route path.

show ap name mesh path

To see information about the mesh AP's path, use the **show ap name** *ap-name* **mesh path** command.

show ap name *ap-name* **mesh path** [**chassis** {*chassis-number* | **active** | **standby**} **R0**]

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/cmd-ref/b_wl_16_10_cr/show-commands.html#wp1177334405

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP, and the MAP learns all available neighbors and determines the route path to the RAP. The routing information

	contains the MAC address of each hop to the RAP. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1h) wherein one of the datasets contained in the AP node comprises an amount of AP data traffic wherein the amount of AP data traffic comprises the amount of data exchange between the AP node and the zero or more clients in communication with the node and wherein another of the datasets contained in the AP node comprises a maximum capacity amount of the AP node.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein one of the datasets contained in the AP node comprises an amount of AP data traffic wherein the amount of AP data traffic comprises the amount of data exchange between the AP node and the zero or more clients in communication with the node and wherein another of the datasets contained in the AP node comprises a maximum capacity amount of the AP node. Cisco's Catalyst 9100 and Aironet Access Points also comprise datasets with the amount of data exchanged between the AP node and the clients as shown by the "config ap client-trace" command. The APs keep a log of client traffic, and the data exchanged between the node and the clients in communication with the node. **config ap client-trace** To configure client trace on the access point, use the **config ap client-trace** command. **config ap client-trace {address {add | clear-all | delete} | all-clients {enable | disable} | filter {all {enable | disable} | arp {enable | disable} | assoc {enable | disable} | auth {enable | disable} | dhcp {enable | disable} | eap {enable | disable} | icmp {enable | disable} | ndp {enable | disable} | probe {enable | disable}} | inline-mon {enable | disable} | output *console-log* | start | stop}** Source: https://www.cisco.com/c/en/us/td/docs/wireless/access_point/ios-xe/command-reference/b-cisco-cat-ap-iosxe-cr/config_commands.pdf#page=1.30; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/access_point/wave2-ap/command-reference/8-10/b-cisco-wave2-ap-cr-810/config_commands.html#config-ap-client-trace The existence of the claimed dataset comprising the amount of data exchanged between the AP node and the zero or more clients of the access points is further evidenced by at least the output of the "show ap <ap-name> mesh cac access" command, which displays "the number of active calls on access radio for a given AP and its Mesh tree."

```
Device# show ap <ap-name> mesh cac access

Depth   AP Name                          Slot  Radio      BW Used  Call
------  -------------------------------  ----  ---------  -------  ----
0       AALUKKAL-1562-RAP                0     802.11b/g  0        0
                                         1     802.11a    0        0
1       AP380E.4DBF.C6B0                 0     802.11b/g  1072     2
                                         1     802.11a    0        0
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-3/cmd-ref/b_wl_17_3_cr/show-commands.html#wp1976338546

The existence of the claimed dataset comprising a maximum capacity amount of the AP node is evidenced by Cisco's Mesh Call Admission Control (CAC). Each access point determines whether it is capable of accommodating a particular call by looking at the bandwidth available.

Mesh CAC

The Call Admission Control (CAC) enables a mesh access point to maintain controlled quality of service (QoS) on the controller to manage voice quality on the mesh network. Bandwidth-based, or static CAC enables the client to specify how much bandwidth or shared medium time is required to accept a new call. Each access point determines whether it is capable of accommodating a particular call by looking at the bandwidth available and compares it against the bandwidth required for the call. If there is not enough bandwidth available to maintain the maximum allowed number of calls with acceptable quality, the mesh access point rejects the call.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html#mesh-cac

This evidence supports that, on information and belief, the Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprise datasets with an amount of data exchange between the AP node and the zero or more clients in communication with the node and a maximum capacity amount of the AP node.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP, and the MAP learns all available neighbors and determines the route path to the RAP. AWPP's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change.

	A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 2	
2. The wireless mesh network of claim 1 wherein one of the datasets contained in the AP node comprises a cost to connect and said cost to connect is calculated using the level of network congestion experienced by the available parent node.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein one of the datasets contained in the AP node comprises a cost to connect and said cost to connect is calculated using the level of network congestion experienced by the available parent node. As shown in Cisco's documentation, AWPP uses ease to determine the preferred path, with ease considered the opposite of cost, such that the preferred path is the path with the highest ease/lowest cost. Cost is calculated at the AP in terms of signal strength and the hops to get back to the controller. The child node or MAPs continuously send NDReq messages to the parent to determine the availability of the parent node to service them adequately. **Criteria to Choose the Best Parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: • A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels. • The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor. • The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state. • Parent maintenance and optimization occurs in the *maintain* state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

	The mesh access point can also operate as a relay node for other access points not directly connected to a wired network. Intelligent wireless routing is provided by the Adaptive Wireless Path Protocol (AWPP). This Cisco protocol enables each mesh access point to identify its neighbors and intelligently choose the optimal path to the wired network by calculating the cost of each path in terms9 of the signal strength and the number of hops required to get to a controller. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh "Load value" (*e.g.*, network or parent capacity) is also considered in calculating the "ease" value and cost to connect. This feature tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink. This 'switching' parent procedure need not be a triggered event like parent loss detection, but also on identifying a better parent using 'Auto parent selection algorithm' when the child MAP still has its current parent uplink active. The " Auto Parent selection algorithm" is based on the new "ease" values. For better convergence calculation in rel 8.3 a new "Ease" value introduced for smoother and faster parent or neighbor finding and auto parent connection algorithm. The Ease value is based on SNR, number of Hops, Timers and load values. For the off-channel neighbors the AdjustedEase value will be used and a best neighbor per off-channel shall be identified based on its highest AdjustedEase value. StickyEase shall be applicable only for on-channel parent. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf AWPP's and the AP's focus on factors like link quality/signal strength, reducing the number of hops and load value helps to minimize and avoid congestion in the mesh network. For example, lower throughput is an indication of congestion. By choosing the "best" path to the wired network based on these criteria, AWPP's route path dataset addresses the level of network congestion experienced by an available parent node. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP, and the MAP learns all available neighbors and determines the route path to the RAP. AWPP's calculations are

	continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. AWPP thus inherently addresses network congestion by selecting parents and paths that are less utilized and more reliable. This reduces the likelihood of bottlenecks and improves overall network performance. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 3	
3. The wireless mesh network of claim 1 wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein the cost to connect of the associated parent is the lowest of all available parent nodes.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprise wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein the cost to connect of the associated parent is the lowest of all available parent nodes. As shown in Cisco's documentation, AWPP uses ease (*e.g.*, communication criteria) to determine the preferred path, with ease considered the opposite of cost, such that the preferred path (and parent) is the path with the highest ease/lowest cost (*e.g.*, "the cost to connect of the associated parent is the lowest of all available parent nodes"). **Criteria to Choose the Best Parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: • A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels. • The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor. • The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state. • Parent maintenance and optimization occurs in the *maintain* state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

	The mesh access point can also operate as a relay node for other access points not directly connected to a wired network. Intelligent wireless routing is provided by the Adaptive Wireless Path Protocol (AWPP). This Cisco protocol enables each mesh access point to identify its neighbors and intelligently choose the optimal path to the wired network by calculating the cost of each path in terms9 of the signal strength and the number of hops required to get to a controller. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP, and the MAP learns all available neighbors and determines the route path to the RAP. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 4	
4. The wireless mesh network of claim 1 wherein the communication criteria comprises instructions for the AP node to associate with an available parent node wherein the amount of AP data traffic is less than the sum of the maximum capacity amount of the available parent node and the	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein the communication criteria comprises instructions for the AP node to associate with an available parent node wherein the amount of AP data traffic is less than the sum of the maximum capacity amount of the available parent node and the available parent amount of AP data traffic. As discussed above, a parent offers the best route back to the RAP based on its ease values (*e.g.*, "communication criteria"). AWPP will calculate the best path back to the RAP.

available parent amount of AP data traffic.

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
  - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

"Load value" (*e.g.*, network or parent capacity) is also considered in calculating the "ease" value and cost to connect.

This feature tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink. This 'switching' parent procedure need not be a triggered event like parent loss detection, but also on identifying a better parent using 'Auto parent selection algorithm' when the child MAP still has its current parent uplink active. The " Auto Parent selection algorithm" is based on the new "ease" values. For better convergence calculation in rel 8.3 a new "Ease" value introduced for smoother and faster parent or neighbor finding and auto parent connection algorithm. The Ease value is based on SNR, number of Hops, Timers and load values. For the off-channel neighbors the AdjustedEase value will be used and a best neighbor per off-channel shall be identified based on its highest AdjustedEase value. StickyEase shall be applicable only for on-channel parent.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP, and the MAP learns all available neighbors and determines the route path to the RAP. AWPP's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. AWPP contributes to load balancing by selecting the path with the best quality and minimizing the number of hops. This

	helps to distribute traffic more effectively across the mesh network and potentially reduces congestion. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 5	
5. The wireless mesh network of claim 4 wherein the communication criteria comprises instructions for the AP node to associate with a single suitable parent node wherein a parent node is suitable if the throughput capacity of the parent node is the highest of all available parent nodes.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 4 as shown above, comprises wherein the communication criteria comprises instructions for the AP node to associate with a single suitable parent node wherein a parent node is suitable if the throughput capacity of the parent node is the highest of all available parent nodes. AWPP determines the parent that offers the best route back to the RAP based on ease values (*e.g.*, "communication criteria"). Ease is calculated by using, *inter alia*, the link hop value. The parent with a higher ease value is selected. The parent with the lower number of hops will be chosen as the highest throughput capacity. If defined as "throughput," signal strength can also be used and the parent with the highest signal strength is chosen. **Mesh neighbors, parents, and children** Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73). • A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP. ◦ Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected. • A child AP selects the parent AP as its best route back to the RAP. • A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Note With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

- Number of backhaul hops is limited to eight but we recommend three to four hops.

  The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.

- Number of MAPs per RAP.

  There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf;

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP and chooses the parent with the least number of hops available. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 6	
6. The wireless mesh network of claim 1 wherein one of the datasets contained in an AP node comprises a throughput capacity of the AP node.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein one of the datasets contained in an AP node comprises a throughput capacity of the AP node. AWPP is run on an access point and automatically determines the best path back to the RAP for that access point by calculating the cost of the path in terms of signal strength and number of hops ("dataset contained in an AP node comprises a throughput capacity of the AP node"). AWPP continuously monitors conditions and changes routes, if necessary, based on the changes in the conditions. On information and belief, a MAP node contains a dataset with its throughput capacity based on its number of hops or signal strength/distance from the RAP so it can determine whether a better path exists. **Adaptive Wireless Path Protocol** The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP and chooses the parent with the least number of hops available.

	A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 7	
7. The wireless mesh network of claim 1 wherein the means for switching two-way data communications selects the second associated parent node when the first associated parent node amount of AP data traffic approaches the parent node maximum capacity amount within a congestion value set by the access server.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein the means for switching two-way data communications selects the second associated parent node when the first associated parent node amount of AP data traffic approaches the parent node maximum capacity amount within a congestion value set by the access server. As discussed above, AWPP's and the AP's focus on factors like link quality/signal strength and reducing the number of hops and load value helps to minimize and avoid congestion in the mesh network. As also discussed above, parameters set at the access server, including number of hops, link SNR, global mesh parameters, and data rate, address congestion in the mesh network. By choosing the "best" path to the wired network based on these criteria, AWPP's route path dataset addresses the level of congestion experienced by an available parent node. **Adaptive Wireless Path Protocol** The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's mesh access points will monitor throughput and will steer to another node as a part of load balancing, including client connection limits set at the controller:

show wireless load-balancing

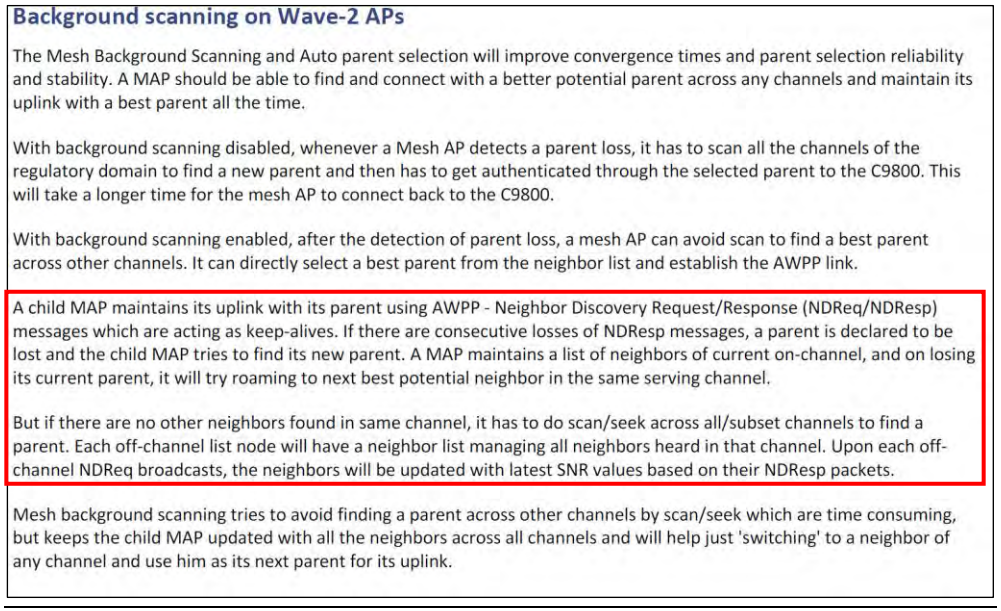
To display the status of the load-balancing feature, use the **show wireless load-balancing** command in privileged EXEC mode.

```
 > show wireless load-balancing
Aggressive Load Balancing.............................: per WLAN enabling
Aggressive Load Balancing Window (clients)..................:: 5
Aggressive Load Balancing Denial Count.....................:: 3

Statistics
Total Denied Count (clients)................................:: 0
Total Denial Sent (messages)................................:: 0
Exceeded Denial Max Limit Count (times).....................:: 0
None 5G Candidate Count (times).............................:: 0
None 2.4G Candidate Count (times)...........................:: 0
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/cmd-ref/b_wl_16_10_cr/show-commands.html#wp3747685901l; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_16_10_cg/client-load-balancing.pdf

As confirmed in Cisco's documentation, and on information and belief as shown in Cisco's proprietary code, a list of maximum capacity amounts and congestion threshold values are obtained and maintained on the access points.



Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

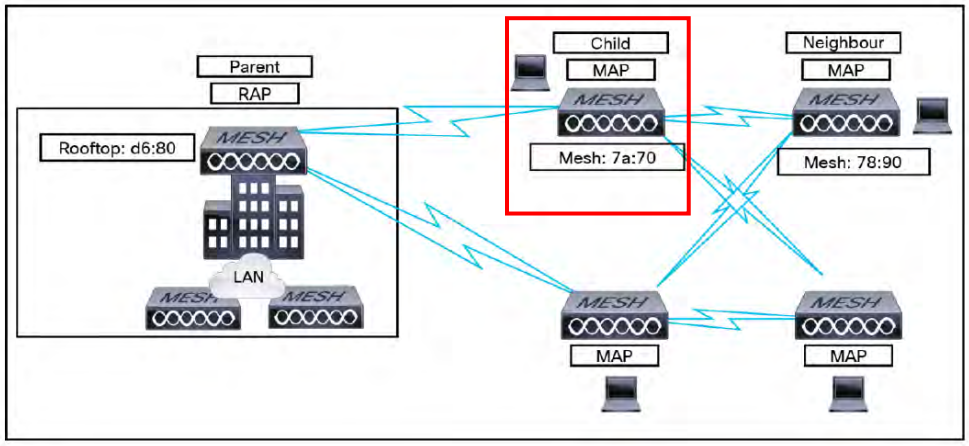
But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by using parameters or criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AWPP process and criteria finds the best path back to the RAP, and the MAP learns all available neighbors and determines the route path to the RAP. AWPP's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. AWPP thus inherently addresses congestion by selecting parents and paths that are less utilized and more reliable. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

To the extent this is considered an element under 35 U.S.C. §112 (¶ 6), in accordance with P.R. 3-1(c), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function include, *e.g.*, the software algorithms in Cisco's proprietary AWPP protocol and mesh background scanning for parent selection identified above. MeshDynamics's investigation into Cisco's proprietary algorithms remains ongoing and MeshDynamics will supplement or amend its P.R. 3-1(c) identification following discovery.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 8	
8. The wireless mesh network of claim 1 wherein the wireless communication with client includes a means to adjust a signal strength of the AP node wireless communication with clients and wherein the means for adjusting the signal strength of the AP node decreases the signal strength of the AP node when the amount of data sent from the zero or more clients exceeds a data congestion threshold value.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein the wireless communication with client includes a means to adjust a signal strength of the AP node wireless communication with clients and wherein the means for adjusting the signal strength of the AP node decreases the signal strength of the AP node when the amount of data sent from the zero or more clients exceeds a data congestion threshold value (*e.g.*, wireless load-balancing or aggressive client load balancing). The AP node (*e.g.*, MAP) is in wireless communication with zero or more clients (*e.g.*, laptops) as shown in Cisco's documentation:  Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh Cisco utilizes aggressive client load balancing as a means to balance clients and optimize overall network performance and prevent overloading an AP node:

Aggressive client load balancing

Aggressive client load balancing is a wireless network feature that

- distributes wireless clients across lightweight APs managed by the same controller
- redirects client association attempts from overutilized APs to less busy APs using standardized response codes, and
- helps optimize overall network performance by preventing AP overload.

Guidelines of aggressive client load balancing

When a wireless client attempts to associate to a lightweight AP, the associated response packets are sent to a client with an 802.11 response packet including status code 17. This code 17 indicates that the corresponding AP is busy. The AP does not respond with 'success' if the AP threshold is not met. If the AP utilization threshold is exceeded and another less busy AP hears the client request, the AP responds with code 17 (AP busy).

If AP1 has more clients than AP2 during the load-balancing window, AP1 is considered busier. When a client attempts to associate to AP1, the client receives an 802.11 response packet with status code 17, indicating that the AP is busy, and the client attempts to associate to a different AP.

Configure the controller to deny client associations up to 10 times. If a client tries to connect 11 times, it will connect on the 11th attempt. Enable or disable load balancing on a particular WLAN. This option is helpful if you need to exclude time-sensitive voice clients from load balancing.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_16_10_cg/client-load-balancing.html

Cisco's mesh access points will monitor throughput and will steer to another node as a part of load balancing, including client connection limits set at the controller:

show wireless load-balancing

To display the status of the load-balancing feature, use the **show wireless load-balancing** command in privileged EXEC mode.

```
 > show wireless load-balancing
Aggressive Load Balancing.............................: per WLAN enabling
Aggressive Load Balancing Window (clients)..................:: 5
Aggressive Load Balancing Denial Count.....................:: 3

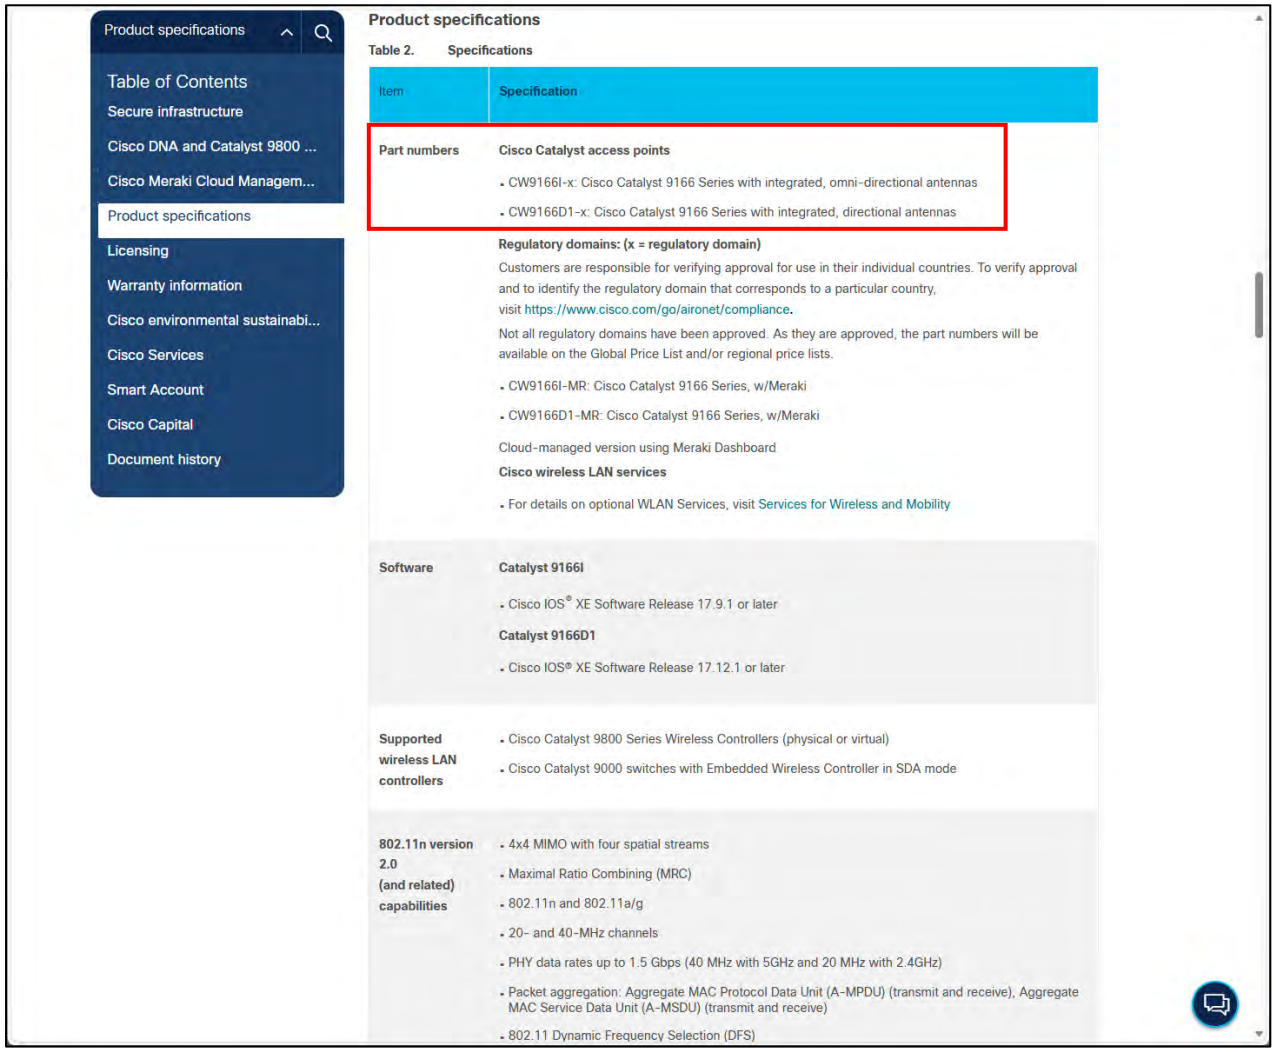
Statistics
Total Denied Count (clients)................................:: 0
Total Denial Sent (messages)................................:: 0
Exceeded Denial Max Limit Count (times).....................:: 0
None 5G Candidate Count (times).............................:: 0
None 2.4G Candidate Count (times)...........................:: 0
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/cmd-ref/b_wl_16_10_cr/show-commands.html#wp3747685901l; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/config-guide/b_wl_16_10_cg/client-load-balancing.pdf

Client steering	Enhanced to help clients that are 6GHz capable to leave the 5GHz radio and connect to the 6GHz radio. Wi-Fi 6E clients are automatically directed to connect to the 6GHz radio to take advantage of the benefits it offers and free up the 2.4GHz and 5GHz radios for legacy clients.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9166-series-access-points/catalyst-9166-series-access-points-ds.html#Productspecifications.

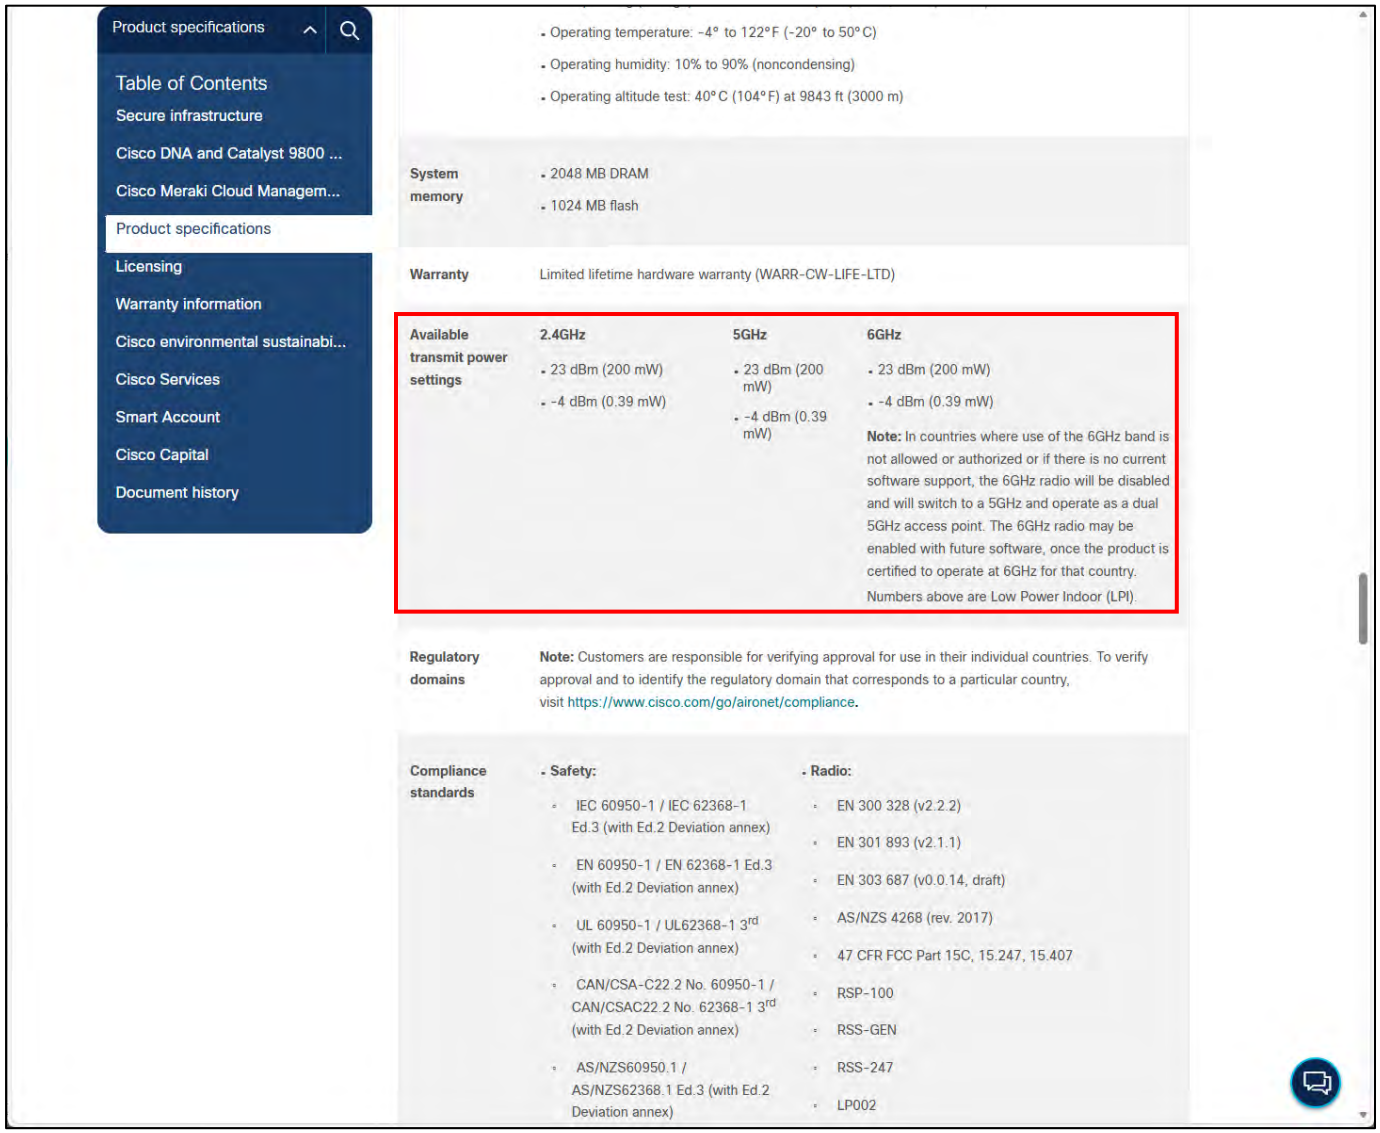
On information and belief, Cisco Catalyst access points have a means for adjusting the signal strength of the AP node that decreases the signal strength of the AP node transmitted outward from the omni-directional antennas or directional antennas:



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9166-series-access-points/catalyst-9166-series-access-points-ds.html#Productspecifications.

Cisco effectively changes the strength of signal to the clients with the available transmit power settings to steer the beam via directional antennas, or by decreasing the transmit power settings strength of the signal to dissuade clients from connecting to the AP node. For example, when traffic from one or more clients is high and the directional antennas are steering the RF beam to those clients, other clients see a reduced signal strength which may cause those clients to seek a

different AP node. Additionally, the AP node connected to the client may decrease the output signal strength to encourage clients to seek a different AP node to select:



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9166-series-access-points/catalyst-9166-series-access-points-ds.html#Productspecifications

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting an appropriate AP for a client in the same way by maintaining a list of neighboring AP's and using parameters or criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The aggressive client load balancing process distributes wireless clients across lightweight APs managed by the same controller and helps optimize overall network performance by preventing AP overload. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. To the extent this is considered an element under 35 U.S.C. §112 (¶ 6), in accordance with P.R. 3-1(c), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function include, *e.g.*, the software algorithms in Cisco's proprietary AWPP protocol, aggressive client load balancing, and mesh background scanning for parent selection identified above. MeshDynamics's investigation into Cisco's proprietary algorithms remains ongoing and MeshDynamics will supplement or amend its P.R. 3-1(c) identification following discovery and/or claim construction. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 9	
A wireless mesh network comprising:	*See* (1pre) above, the contents of which are incorporated herein by reference.
an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	*See* (1a) above, the contents of which are incorporated herein by reference.
one or more root nodes connected to said access server and an external network;	*See* (1b) above, the contents of which are incorporated herein by reference.
one or more AP nodes wherein each AP node is in wireless two-way data communication with an	*See* (1c) above, the contents of which are incorporated herein by reference.

associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria or the associated parent node is a root node within wireless communication range;	
wherein an AP node is in wireless communication with zero or more clients; and	*See* (1d) above, the contents of which are incorporated herein by reference.
wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets	*See* (1e) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an	*See* (1g) above, the contents of which are incorporated herein by reference.

AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node	
wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset; and	*See* (1f) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an AP node comprises a dataset of child node identifiers wherein the dataset of child node identifiers is a dataset identifying each AP node in wireless communication with the AP.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein one of the datasets contained in an AP node comprises a dataset of child node identifiers wherein the dataset of child node identifiers is a dataset identifying each AP node in wireless communication with the AP. As confirmed in Cisco's documentation, and on information and belief as shown in Cisco's proprietary code, a list of nodes (*e.g.*, dataset) identifying each AP node in RF range, including those in wireless communication, is obtained and maintained on each MAP, including those off-channel that will have a neighbor list as well, *e.g.*, from background scanning or "AWPP – Neighbor Delivery Request/Response (NDReq/NDResp)" messages.

Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
  - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco uses Address Resolution Protocol (ARP) to find hardware addresses of a host from a known IP address. This will show what nodes that the APs have communicated with.

show arp

To view the ARP table, use the **show arp** command.

show arp

Syntax Description

arp	Shows ARP table

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
8.1.111.0	This command was introduced.

The following example shows a sample output of the command:

```
cisco-ap# show arp

Address Age (min)     Hardware Addr
   9.11.8.1         0 84:80:2D:A0:D2:E6
9.11.32.111         0 3C:77:E6:02:33:3F
```

Source: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr_arp/configuration/xe-3s/arp-xe-3s-book/arp-config-arp.pdf

Cisco's command line manual further shows that the access points have a dataset showing what is surrounding the node, including child node identifies with "Child" as the State field:

show wireless mesh neighbor

To see the neighbors of all connected mesh APs, use the **show wireless mesh neighbor** command.

show wireless mesh neighbor [**chassis** {*chassis-number* | **active** | **standby**} **R0**]

```
AP Name/Radio             Channel   Rate      Link-snr   Flags   State
-------------------------------------------------------------------------------

AP Name : Mesh-AP01

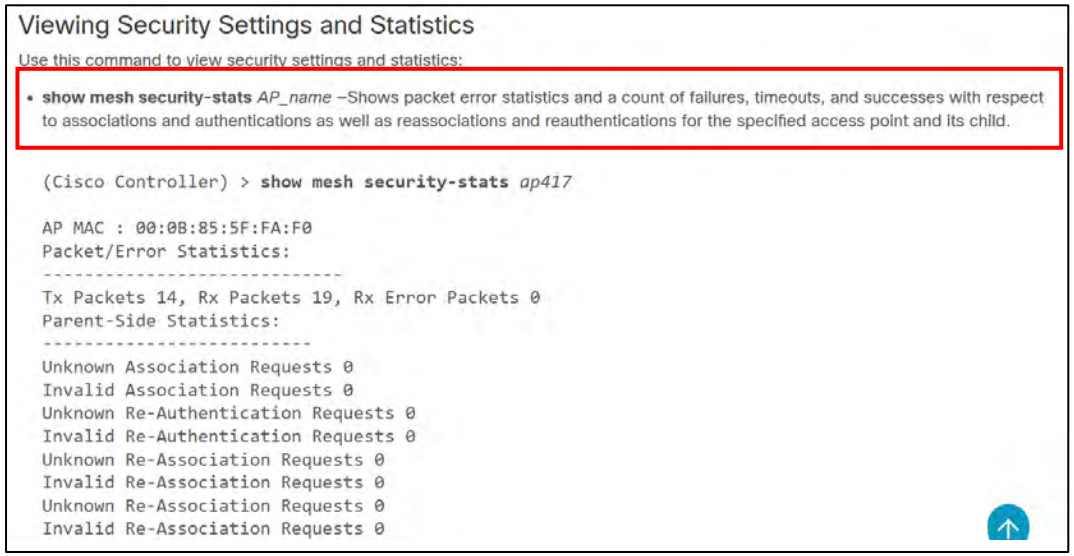
54:9f:c6:fa:5c:71         149       auto      0          40
b0:c5:3c:e5:d9:71         149       auto      22         49      UPDATED NEIGH
e8:eb:34:d5:88:d1         149       auto      0          40
e8:eb:34:d5:8d:d1         149       auto      18         49      UPDATED CHILD
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/cmd-ref/b_wl_16_10_cr/show-commands.html#wp3747685901

- **show mesh neighbor summary** –Displays summary information about mesh neighbors. Neighbor information includes MAC addresses, parent-child relationships, and uplink and downlink (SNRUp, SNRDown).

```
(Cisco Controller) > show mesh neighbor summary ap1500:62:39:70
AP Name/Radio Mac Channel Snr-Up Snr-Down Link-Snr Flags  State
mesh-45-rap1      165     15     18       16       0x86b  UPDATED NEIGH PARENT BEACON
00:0B:85:80:ED:D0 149     5      6        5        0x1a60 NEED UPDATE BEACON DEFAULT
00:17:94:FE:C3:5F 149     7      0        0        0x860  BEACON
```

Note After review of the **show mesh** commands above, you should be able to see the relationships between the nodes of your network and verify the RF connectivity by seeing the SNR values for every link.



Viewing Security Settings and Statistics

Use this command to view security settings and statistics:

- **show mesh security-stats** *AP_name* –Shows packet error statistics and a count of failures, timeouts, and successes with respect to associations and authentications as well as reassociations and reauthentications for the specified access point and its child.

```
(Cisco Controller) > show mesh security-stats ap417

AP MAC : 00:0B:85:5F:FA:F0
Packet/Error Statistics:
----------------------------
Tx Packets 14, Rx Packets 19, Rx Error Packets 0
Parent-Side Statistics:
-------------------------
Unknown Association Requests 0
Invalid Association Requests 0
Unknown Re-Authentication Requests 0
Invalid Re-Authentication Requests 0
Unknown Re-Association Requests 0
Invalid Re-Association Requests 0
Unknown Re-Association Requests 0
Invalid Re-Association Requests 0
```

```
Child-Side Statistics:
-------------------------
Association Failures 0
Association Timeouts 0
Association Successes 0
Authentication Failures 0
Authentication Timeouts 0
Authentication Successes 0
Re-Association Failures 0
Re-Association Timeouts 0
Re-Association Successes 0
Re-Authentication Failures 0
Re-Authentication Timeouts 0
Re-Authentication Successes 0
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-0/design/guide/mesh80/m_checking-the-health-of-the-network.html

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Each MAP in the wireless mesh network, on the channel and off-channel, is maintaining a list of nodes within the RF range of each other

	and thus includes all child nodes in the network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 10	
10. The wireless mesh network of claim 9 wherein the dataset of child node identifiers is accessible to each child node.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 9 as shown above, comprises wherein the dataset of child node identifiers is accessible to each child node. As discussed above, each MAP in the wireless mesh network, on the channel and off-channel, is maintaining a list of nodes within the RF range of each other, and on information and belief, in wireless communication, and thus includes all child nodes in the network. MAPs communicate among themselves, and on information and belief, the list (*e.g.*, dataset) of child node identifiers that identifies each AP node in wireless communication with the AP is accessible to each child node. • **show mesh neighbor summary** –Displays summary information about mesh neighbors. Neighbor information includes MAC addresses, parent-child relationships, and uplink and downlink (SNRUp, SNRDown). (Cisco Controller) > **show mesh neighbor summary** *ap1500:62:39:70* AP Name/Radio Mac Channel Snr-Up Snr-Down Link-Snr Flags State mesh-45-rap1 165 15 18 16 0x86b UPDATED NEIGH PARENT BEACON 00:0B:85:80:ED:D0 149 5 6 5 0x1a60 NEED UPDATE BEACON DEFAULT 00:17:94:FE:C3:5F 149 7 0 0 0x860 BEACON **Note** After review of the **show mesh** commands above, you should be able to see the relationships between the nodes of your network and verify the RF connectivity by seeing the SNR values for every link.

Viewing Security Settings and Statistics

Use this command to view security settings and statistics:

- **show mesh security-stats** *AP_name* –Shows packet error statistics and a count of failures, timeouts, and successes with respect to associations and authentications as well as reassociations and reauthentications for the specified access point and its child.

```
(Cisco Controller) > show mesh security-stats ap417

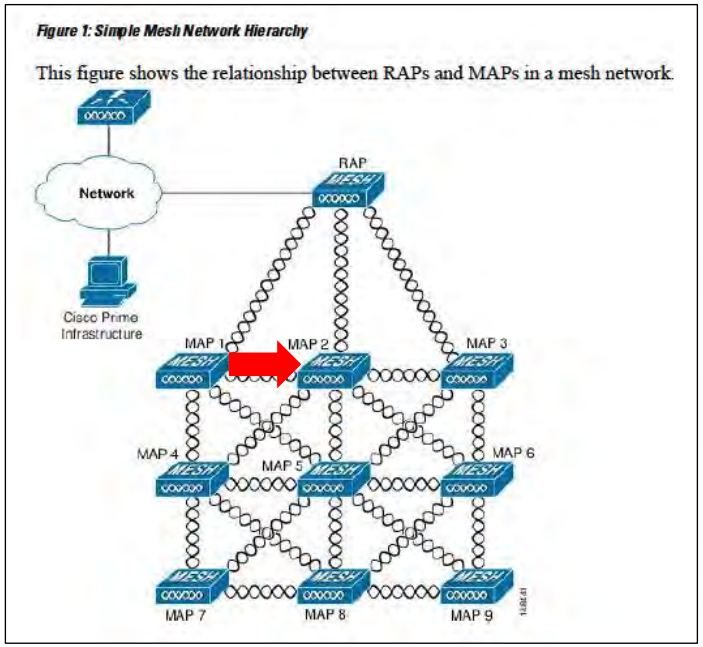
AP MAC : 00:0B:85:5F:FA:F0
Packet/Error Statistics:
----------------------------
Tx Packets 14, Rx Packets 19, Rx Error Packets 0
Parent-Side Statistics:
-------------------------
Unknown Association Requests 0
Invalid Association Requests 0
Unknown Re-Authentication Requests 0
Invalid Re-Authentication Requests 0
Unknown Re-Association Requests 0
Invalid Re-Association Requests 0
Unknown Re-Association Requests 0
Invalid Re-Association Requests 0
```

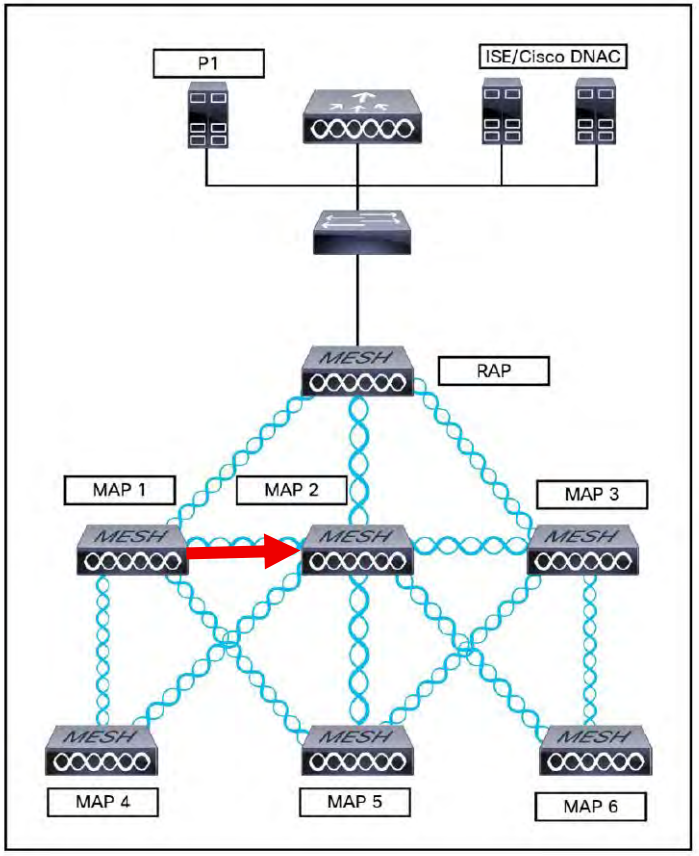
```
Child-Side Statistics:
-------------------------
Association Failures 0
Association Timeouts 0
Association Successes 0
Authentication Failures 0
Authentication Timeouts 0
Authentication Successes 0
Re-Association Failures 0
Re-Association Timeouts 0
Re-Association Successes 0
Re-Authentication Failures 0
Re-Authentication Timeouts 0
Re-Authentication Successes 0
```

Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-0/design/guide/mesh80/m_checking-the-health-of-the-network.html

	For Mesh mode, an AP should be configured with Bridge mode. Access points within a mesh network operate in one of the following two ways: • Root AP (RAP) • Mesh AP (MAP) **Note:** All APs are configured and shipped as mesh APs. To use an APs as a RAP, you must reconfigure the mesh AP as a RAP. In all mesh networks, ensure that there is at least one RAP. While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Each MAP in the wireless mesh network, on the channel and off-channel, is maintaining a list of nodes within the RF range of each other and thus includes all child nodes in the network. MAPs communicate among themselves. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 11	
11. The wireless mesh network of claim 10 wherein the AP node dataset of child nodes contains two or more child nodes wherein the zero or more clients in communication with a first child node sends data wherein a	Cisco Catalyst 9100 Access Points in Cisco's Wireless Mesh Network, in addition to the elements of claim 10 as shown above, comprises wherein the AP node dataset of child nodes contains two or more child nodes wherein the zero or more clients in communication with a first child node sends data wherein a destination of the data is a second child node, the first child node sends the data directly to the second child node. As discussed above, MAPs are aware of neighbors and thus can communicate directly over wireless. In the figures below, the client (*e.g.* Laptop) within the first child node (MAP 1) will send data to a destination within a second child node (MAP 2) and that data will go directly from MAP 1 to MAP 2.

destination of the data is a second child node, the first child node sends the data directly to the second child node.



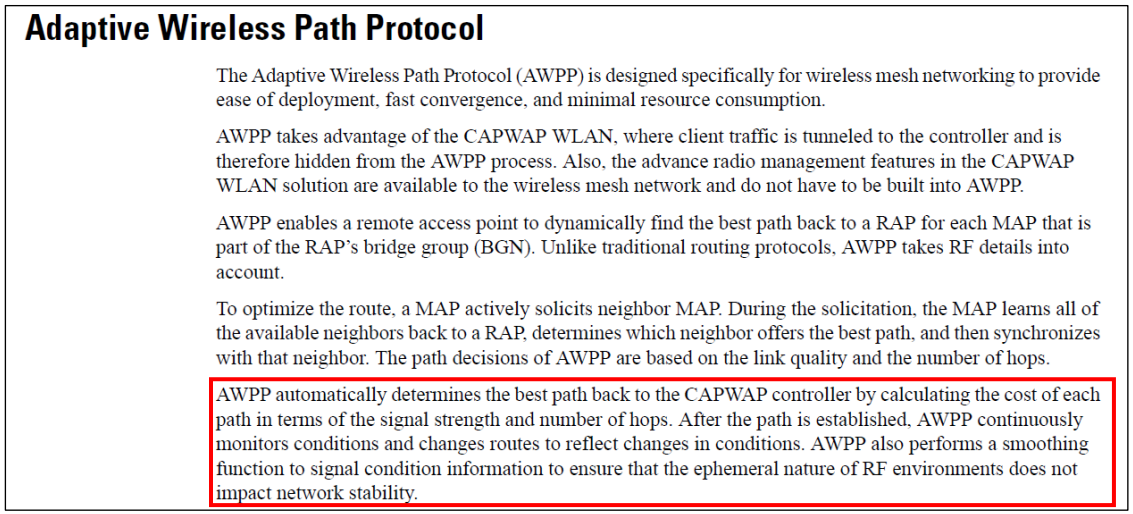


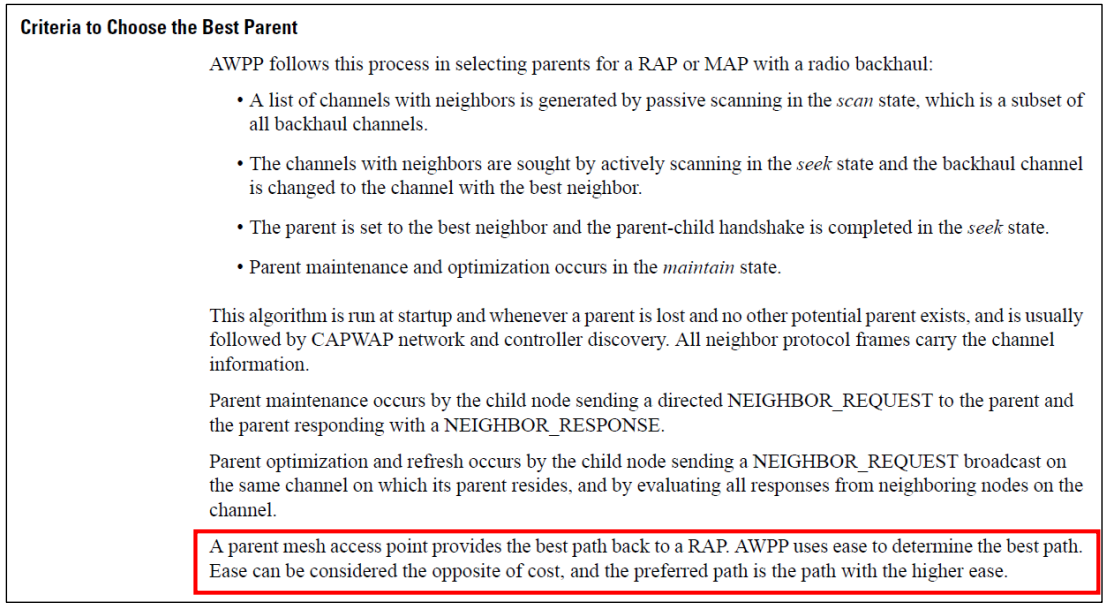
Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of sending data traffic through mesh network in same way by maintaining a list of neighbors and finding the best path to the data traffic destination and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Rather than forcing data back up the hierarchical chain to a root node, and then back down the chain to a destination child node, the mesh is cooperative and since it is aware of the destination node will take the quickest path there (*i.e.*, directly from MAP1 to MAP2). A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 12	
A wireless mesh network comprising:	*See* (1pre) above, the contents of which are incorporated herein by reference.
an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	*See* (1a) above, the contents of which are incorporated herein by reference.
one or more root nodes connected to said access server and an external network	*See* (1b) above, the contents of which are incorporated herein by reference.
one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria or the associated parent node is a root node	*See* (1c) above, the contents of which are incorporated herein by reference.

within wireless communication range;	
wherein an AP node is in wireless communication with zero or more clients; and	*See* (1d) above, the contents of which are incorporated herein by reference.
wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets	*See* (1e) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node	*See* (1g) above, the contents of which are incorporated herein by reference.
wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to	*See* (1f) above, the contents of which are incorporated herein by reference.

become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset; and	
wherein the communication criteria further comprises instructions for the AP node to associate with a single suitable parent node wherein the route path dataset of the parent node is the shortest route path dataset of all available parent nodes.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein the communication criteria further comprises instructions for the AP node to associate with a single suitable parent node wherein the route path dataset of the parent node (*e.g.*, a dataset comprising "the best path back to a RAP") is the shortest route path dataset (*e.g.*, the route path dataset with the lowest number of hops back to the RAP) of all available parent nodes. The AWPP criteria to choose the best parent includes instructions for the MAP or child node to associate with a RAP or parent, which on information and belief, as shown in Cisco's proprietary code, is the shortest route or path, *e.g.*, the best route with the least number of hops back to the RAP.  **Adaptive Wireless Path Protocol** The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh



Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the *scan* state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the *seek* state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the *seek* state.
- Parent maintenance and optimization occurs in the *maintain* state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's Wireless Mesh Network can also be configured as a static mesh configuration and the criteria and route path dataset may be provided by manually assigning parent-child and RAP relationships.

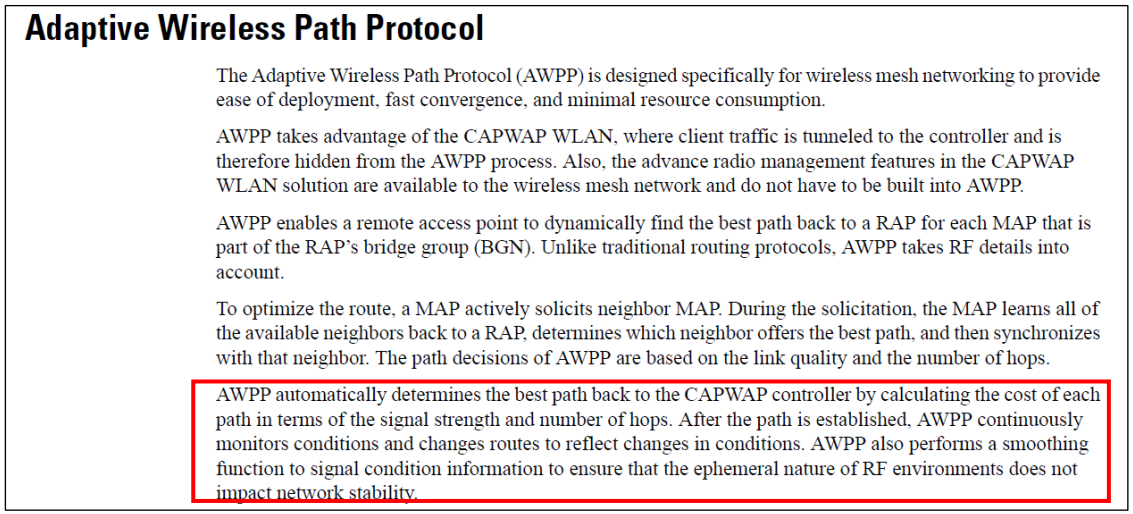
Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AP node in a Cisco Wireless Network is associated with a single suitable parent node when the route path of the parent node is the best path of all available parent nodes. It is the best path and shortest when it is the least number of hops, for example. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 13	
A wireless mesh network comprising:	*See* (1pre) above, the contents of which are incorporated herein by reference.
an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	*See* (1a) above, the contents of which are incorporated herein by reference.
one or more root nodes connected to said access server and an external network;	*See* (1b) above, the contents of which are incorporated herein by reference.
one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication	*See* (1c) above, the contents of which are incorporated herein by reference.

criteria or the associated parent node is a root node within wireless communication range;	
wherein an AP node is in wireless communication with zero or more clients; and	*See* (1d) above, the contents of which are incorporated herein by reference.
wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets	*See* (1e) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node	*See* (1g) above, the contents of which are incorporated herein by reference.
wherein the communication criteria comprises instructions for the AP node to select the associated parent node	*See* (1f) above, the contents of which are incorporated herein by reference.

wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset; and	
wherein the access server functioning parameters includes a latency modifier wherein the AP node means for switching from the first associated parent node to a second associated parent node result in selection of the second associated parent wherein the route path of the second associated parent node is shorter than the first associated route path by a value related to the latency modifier.	Cisco Catalyst 9100 and Aironet Access Points in Cisco's Wireless Mesh Network comprises wherein the access server functioning parameters includes a latency modifier wherein the AP node means for switching from the first associated parent node to a second associated parent node result in selection of the second associated parent wherein the route path of the second associated parent node is shorter than the first associated route path by a value related to the latency modifier. As discussed above, functioning parameters set at the access server includes the signal strength and number of hops ("latency modifier"). The AWPP algorithm results in the selection of a parent with a best or shorter route path shorter by signal strength and number of hops. AWPP determines the lowest cost in terms of signal strength and number of hops. **Mesh neighbors, parents, and children** Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73). • A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP. ◦ Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected. • A child AP selects the parent AP as its best route back to the RAP. • A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh



Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's Wireless Mesh Network can also be configured as a static mesh configuration and the criteria and route path dataset may be provided by manually assigning parent-child and RAP relationships.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AP node in a Cisco Wireless Network is associated with a single suitable parent node when the route path of the parent node is the best path of all available parent nodes. It is the best path and shortest when it is the least number of hops such as when a MAP can connect directly to a RAP, for example. This path would be the shortest number of hops and the least latency. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Exhibit 6 – U.S. Patent No. 7,885,243 (Cisco Wireless Mesh Networks with Meraki and Catalyst CW Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Meraki Wireless Mesh Network, including but not limited to Cisco's Meraki and Catalyst CW Access Points with related cloud architecture and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1-13 (the "Asserted Claims") of U.S. Patent No. 7,885,243 (the "'243 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '243 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale, and importing the Cisco Accused Products. Cisco further directly infringed by testing, setting up, and using Cisco Meraki Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Meraki Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Meraki Wireless Mesh Networks. As indicated in the Meraki Mesh Networking Guide, each Meraki access point has unique mesh characteristics independent of general technology specifications listed on the datasheet. Cisco recommends a Meraki Systems Engineer as a resource to ensure that an access point's capabilities match the desired design targets.[3] The Meraki Trial Program allows customers to run a real-world proof of concept to test the design, which is arranged by a Cisco Meraki account representative. Meraki is also cloud-based architecture, and the infringing mesh networking methods and systems were controlled and put into service by Cisco.[4] Cisco knew how its customers were configuring the specific access points sold, as customers must register to the Meraki cloud and use the invoice or serial number of each purchased access point.[5]

---

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g*., https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

[3] *See, e.g*., https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

[4] *See Centillion Data Systems, LLC v. Qwest Communications International, Inc.*, 631 F.3d 1279 (Fed. Cir. 2011); *NTP, Inc. v. Research in Motion, Ltd*., 418 F. 3d 1282, 1317-1318 (Fed. Cir. 2005); *see also SiRF Tech. Inc. v. ITC*, No. 2009-1262 (Fed. Cir. Apr. 12, 2010) (direct infringement where the claims do not specifically require action by its customers or end users).

[5] *See, e.g*., https://documentation.meraki.com/MR/MR_Installation_Guides/CW9163E-MR_Installation_Guide.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the ’243 Patent and make and use the infringing Cisco Meraki Wireless Mesh Network. Specifically, Cisco’s Meraki Mesh Deployment Guides, Meraki Mesh Networking Guide, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco’s website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Meraki Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-13 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Meraki Wireless Mesh Network, and the infringing configuration and operation of the cloud, access points, and software that comprise the infringing Cisco Meraki Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Meraki Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the ’243 Patent since 2009 when Cisco executives met with MeshDynamics to discuss partnering opportunities and the application that would result in the ’243 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale, and importing the Cisco Accused Products that comprise the Cisco Meraki Wireless Mesh Network in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’243 Patent, or in the alternative was willfully blind to the ’243 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Meraki Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’243 Patent. Cisco Accused Products come enabled as a mesh access point and automatically connect to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particularly suited for Cisco’s Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Meraki Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’243 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the '243 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics's preliminary infringement analysis of the Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco's Meraki Wireless Mesh Network with Meraki and Catalyst CW Access Points
(1pre) A wireless mesh network comprising:	Cisco Meraki Wireless Mesh Network comprises Wi-Fi relay or mesh nodes (*e.g.*, Cisco Meraki and Catalyst CW Access Points) and additional components, including the cloud and software. The Meraki cloud-managed wireless series is a line of enterprise-grade WLAN access points ("APs") that leverage the Cisco Meraki cloud-managed architecture. Both Cisco Meraki and Catalyst CW Access Points support Wi-Fi capabilities and can be paired with the Meraki cloud platform. Meraki AP models start with MR, *e.g.* MR33, MR36,

MR46E, and the Catalyst APs with CW, *e.g.* CW9162, CW9164, CW9166, are dual-mode APs which can be used with the Catalyst 9800 controller or the Meraki cloud architecture.

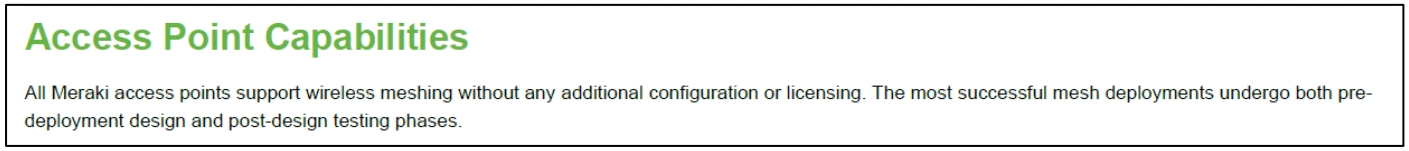
Source: Meraki Family Datasheet; Meraki Wi-Fi Models

The Meraki and Catalyst CW Access Points contain the same mesh networking capabilities, as shown Cisco's documentation and the Meraki Family Datasheet, and the Meraki Wireless Mesh Networking and Mesh Deployment Guide apply to both product lines for setting up a Cisco Meraki Wireless Mesh Network. Multiple access points communicate over wireless interfaces to form a single network. Meraki's mesh networking functionality is automatic, self-healing and available on all Meraki (including Catalyst CW) Access Points.

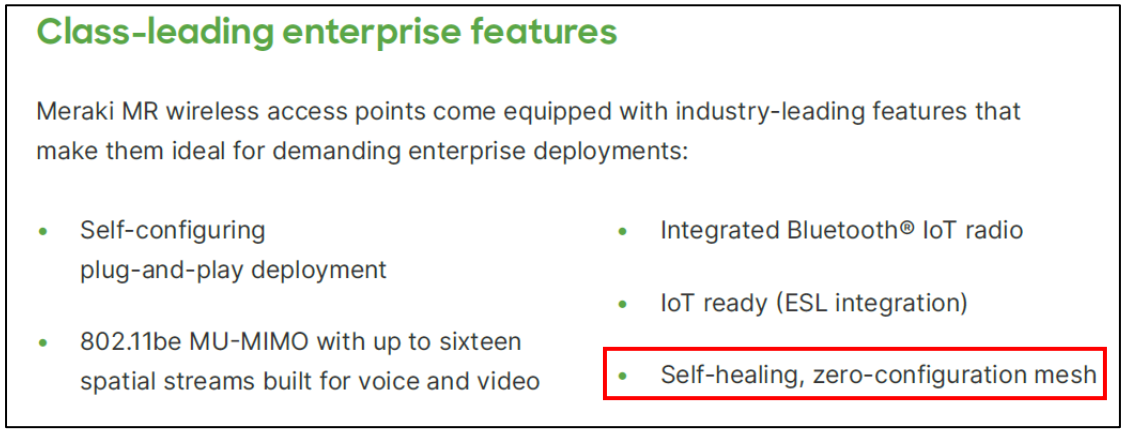
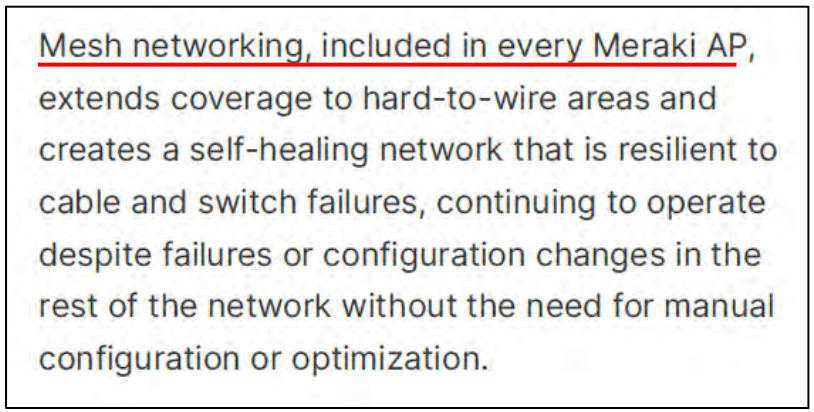


Source: Wireless Mesh Networking; *see also* Meraki Best Practices - Wireless Mesh Networking

Mesh functionality is enabled by default in every Meraki AP. All Meraki APs support wireless meshing without any additional configuration or licensing.



Source: Mesh Deployment Guide

	 Source: Meraki Family Datasheet  Source: Meraki Family Datasheet
(1a) an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	Cisco's Meraki Wireless Mesh Network comprises an access server (*e.g.*, Meraki's cloud centralized management service) wherein the access server sets one or more functioning parameters of the wireless mesh network (*e.g.*, number of hops and link quality). Cisco's Meraki Wireless Mesh Network is a cloud-based system. Users are able to deploy, monitor, and configure the access points in the Meraki wireless mesh network using the Meraki dashboard. Once a user makes a configuration change, the change request is sent to the Meraki cloud and then pushed to the relevant devices.

Meraki Cloud Architecture

The Meraki cloud solution is a centralized management service that allows users to manage all of their Meraki network devices via a single, simple and secure platform.



Users are able to deploy, monitor, and configure their Meraki devices via the Meraki dashboard web interface or via APIs. Once a user makes a configuration change, the change request is sent to the Meraki cloud and is then pushed to the relevant device(s).

Source: Meraki Cloud Architecture

Cisco provides link quality and number of hops recommendations for the Meraki Wireless Mesh Networks. On information and belief, the number of hops and link quality can be set at the cloud-based servers consistent with Cisco's recommendations.

Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

Impact of Meshing on Throughput

Due to the half-duplex nature of wireless communication and that signals being passed through a repeater access point must be retransmitted to the next hop, throughput is greatly reduced when using a repeater. While many factors impact wireless throughput, it is safe to assume that the addition of meshing can reduce throughput by approximately 50%, with that reduction being applied for each subsequent repeater that must be traversed to reach a gateway. Therefore, it is advised to minimize the number of hops between a client and gateway.

Source: Wireless Mesh Networking

Monitoring tools are also available at the access server, including how to ensure each access points has strong neighbors.

Maximum Mesh Hops

There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device.

Multi-Path Mesh

It is recommended that each mesh access point has at least three strong "neighbors," or other access points in the mesh that the access point can "see," so that multiple failover paths exist. Wireless Mesh Networking gives an overview of the monitoring tools available on dashboard, including information about the mesh neighbors table found on individual access point detail pages.

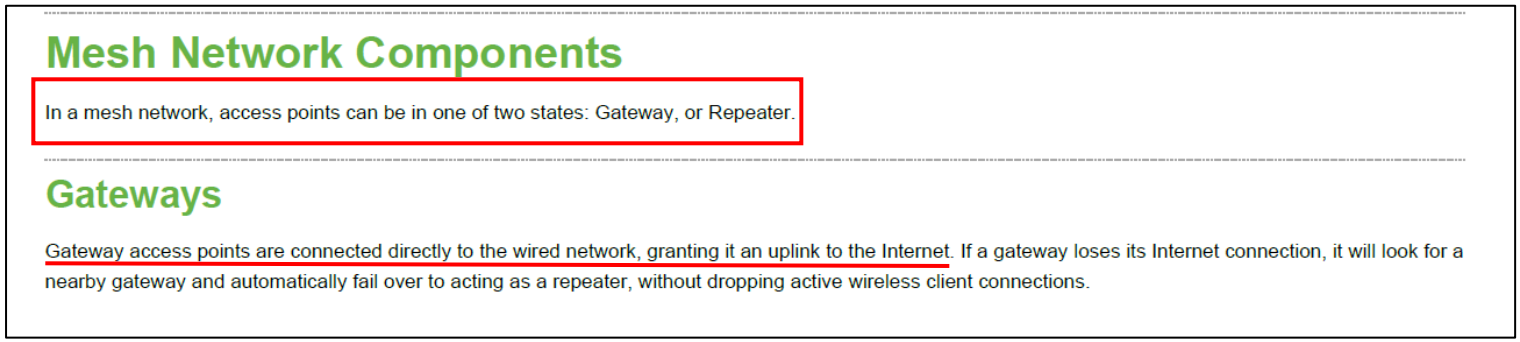
Source: Mesh Deployment Guide

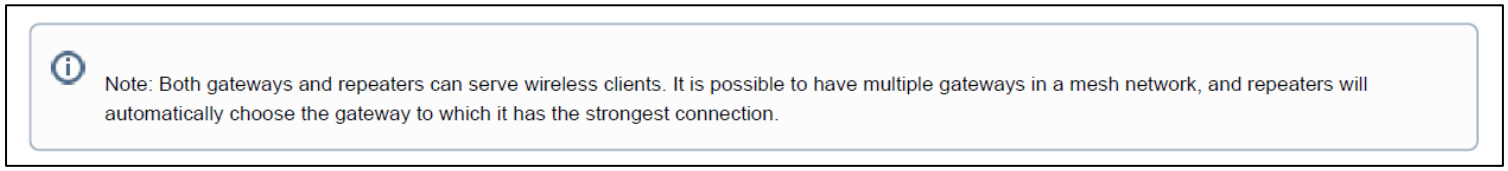
Manual channel selection can be configured using the dashboard at the access server.

Manually Changing Mesh Channels

Allowing Cisco Meraki access points to automatically converge using our proprietary routing protocols and mesh network software mitigates the need for additional configurations, RF Planning and cumbersome wireless site surveys. It is not recommended to manually change channels in a mesh network unless absolutely necessary. In the event that it is deemed necessary to manually select a channel for the mesh network to service clients from, it is recommended that the process illustrated below is used to do so.

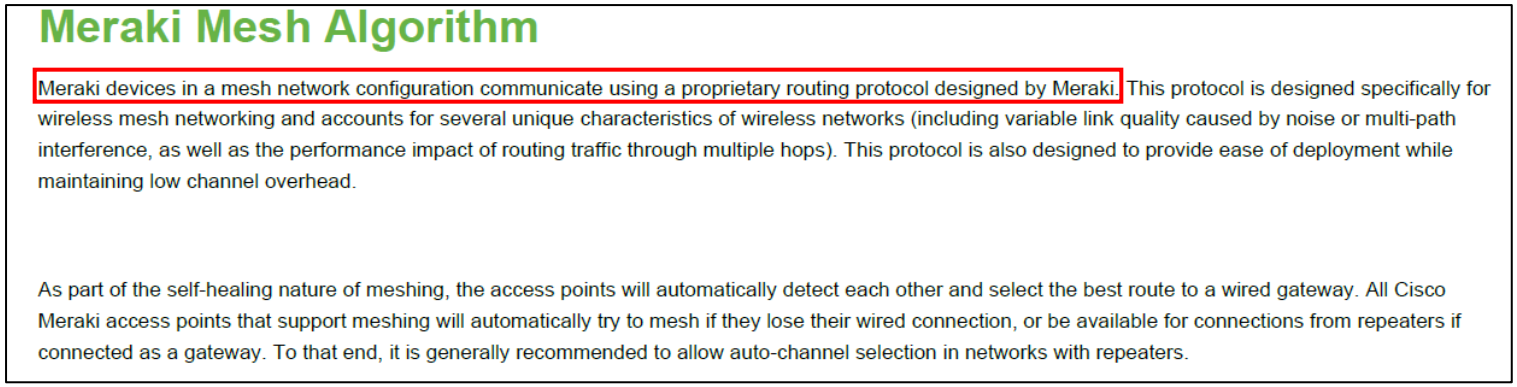
Source: Manually Changing Channels

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of establishing one or more functioning parameters of the wireless mesh network in the same way by using a cloud server to set recommended or global parameters for the wireless mesh network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. To the extent Cisco argues the cloud is not the claimed "access server," for example, a person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1b) one or more root nodes connected to said access server and an external network;	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprise one or more root nodes (*e.g.*, gateway) connected to said access server and an external network. APs within the Cisco Meraki Wireless Mesh Network operate as either a gateway or repeater and are capable of relaying traffic from a child node (*e.g.*, repeater) associated with the parent node (*e.g.*, gateway). In Cisco's Meraki Wireless Mesh Network, APs can be in one of two states: gateway or repeater. Gateway APs are connected directly to the wired network, granting it an uplink to the Internet and connection to the Meraki cloud architecture and access server.  Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide Both gateways and repeaters can serve clients, and it is possible to have multiple gateways in a mesh network.



Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

APs in Cisco's Meraki Wireless Mesh Network communicate traffic among themselves and back to the gateway using a proprietary protocol.

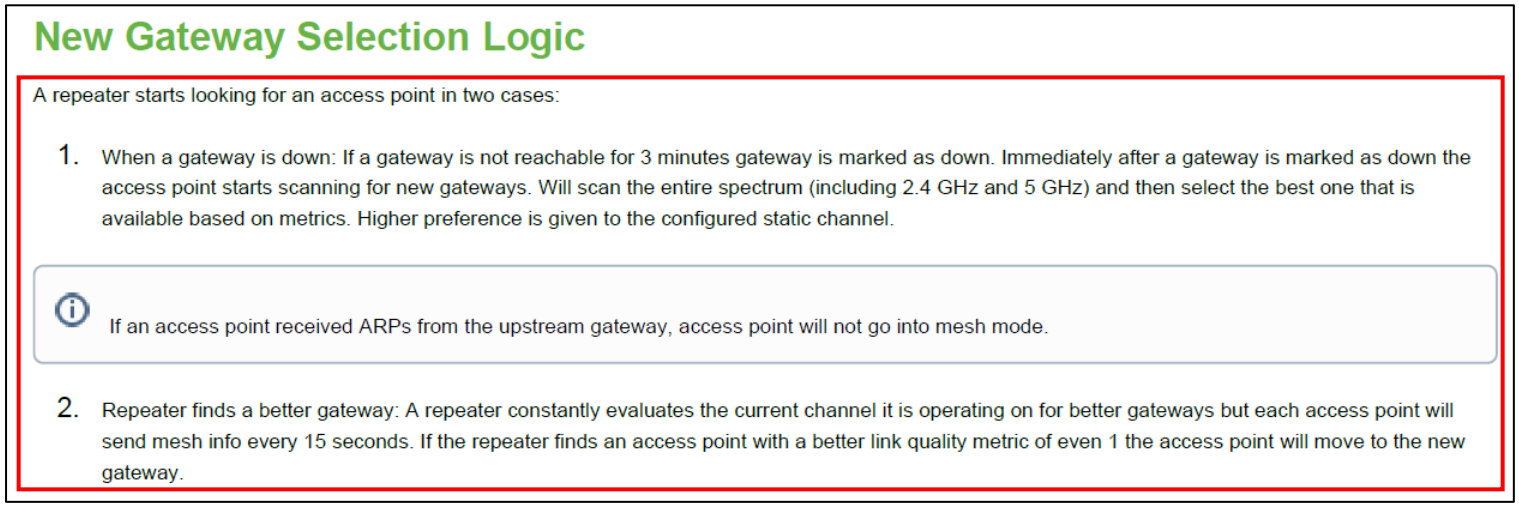


Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of connecting the mesh network to the external network in the same way by using a root access point and achieves the same result of creating a wireless mesh network. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The gateway connects the wireless mesh network to the wired network, providing access to the external network, such as the Internet. The gateway is directly connected to the wired network and all mesh access points are connected to the cloud servers. A person skilled in the relevant art would consider any difference between the claimed invention and the accused product to be minor.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

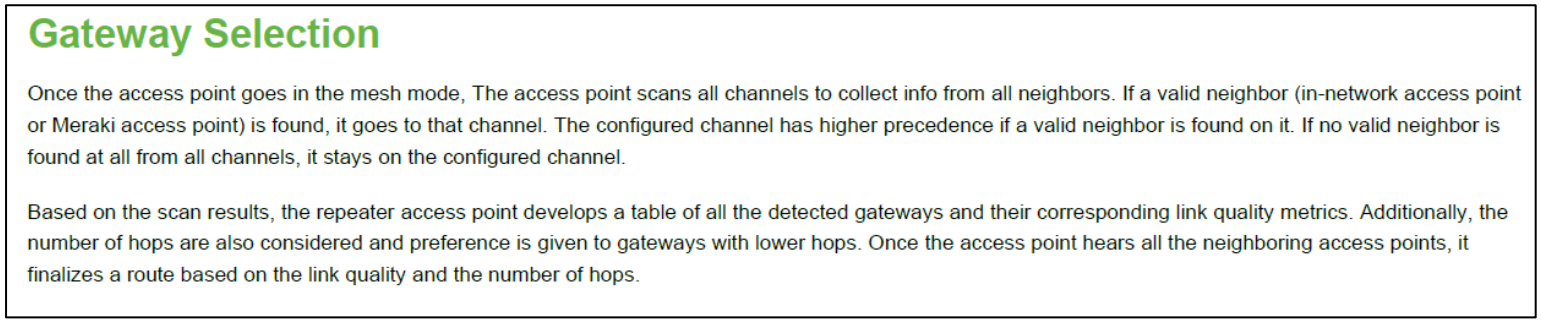
(1c) one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria or the associated parent node is a root node within wireless communication range;	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprise one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node (*e.g.*, gateway) wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria (*e.g.*, Meraki Mesh Algorithm's criteria, including the gateway selection logic and routing metric calculations, and channel selection algorithm criteria) or the associated parent node is a root node (*e.g.*, gateway) within wireless communication range. APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference possible using channel selection algorithms. Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels A repeater can also determine (*e.g.*, for parent optimization and upon the loss of a parent connection) that it should change channels and connect with a better potential parent across any channel to maintain its uplink with the best parent (*i.e.*, the best path back to the gateway determined using criteria based on parameters including, for example, signal strength and number of hops). The channel selection change will choose the best parent. Repeaters in a Cisco Meraki Wireless Mesh Network look for an AP in two cases: when the gateway is down or it finds a better gateway. AP will scan the entire spectrum in the first case. In the second, a repeater is constantly evaluating the current channel for better gateways. If a repeater finds an AP with a better link quality metric, it will move to the new gateway.



Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network scan all channels to collect information from its neighbors. If a valid neighbor is found, the AP switches to the channel.

Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. The number of hops is considered as well, with preference given to gateways with the lower number of hops. An AP finalizes its route based on link quality and number of hops.



Source: Wireless Mesh Networking

Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead.

APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

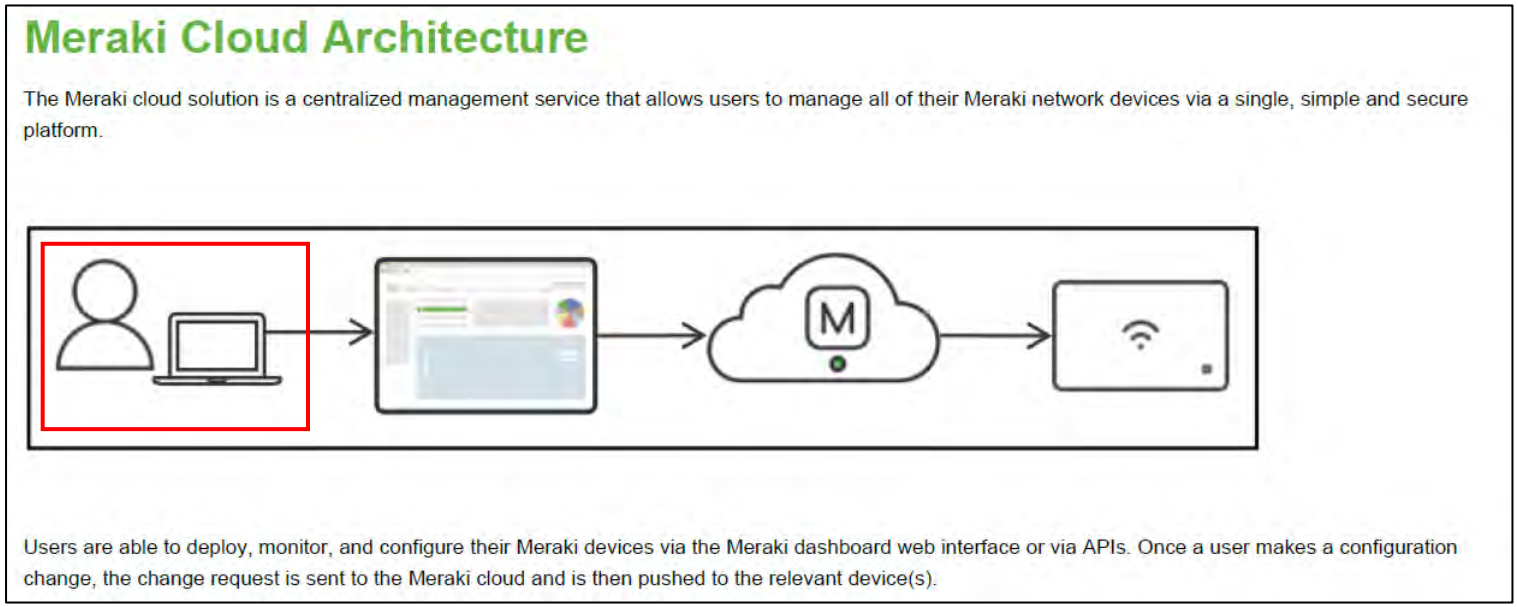
APs in Cisco's Meraki Wireless Mesh Network constantly update their routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the

	claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm uses criteria to select the best parent. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1d) wherein an AP node is in wireless communication with zero or more clients; and	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprise wherein an AP node (*e.g.*, Meraki access point) is in wireless communication with zero or more clients (*e.g.*, laptops) as shown in Cisco's documentation:  Source: Meraki Cloud Architecture
(1e) wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprises wherein an AP node includes a means for switching two-way data communication (*e.g.*, Meraki Mesh Algorithm, gateway selection, and channel selection algorithms) from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets (*e.g.*, routing table with optimal path to gateway and additional datasets identified *infra*).

based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets;	The Cisco Meraki and Catalyst CW Access Points (*e.g.*, repeaters) are in wireless two-way data communication with a parent node by forming a mesh link. A parent can be either the gateway itself or another repeater. Repeaters communicate with each other and the gateway using wireless 802.11 (WiFi) backhaul channels, and 802.11 request/response protocol. Source: Wireless Mesh Networking; *see also* Wi-Fi Basics and Best Practices APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference possible using channel selection algorithms. > Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels A repeater can also determine (*e.g.*, for parent optimization and upon the loss of a parent connection) that it should change channels and connect with a better potential parent across any channel to maintain its uplink with the best parent (*i.e.*, the best path back to the gateway based on signal strength and number of hops/ease criteria). The channel selection change will choose the best parent. Repeaters in a Cisco Wireless Mesh Network look for an AP in two cases: when the gateway is down or it finds a better gateway. AP will scan the entire spectrum in the first case. In the second, a repeater is constantly evaluating the current channel for better gateways. If a repeater finds an AP with a better link quality metric, it will move to the new gateway.

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

> ⓘ If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network scan all channels to collect information from its neighbors. If a valid neighbor is found, the AP switches to the channel.

Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. The number of hops is considered as well, with preference given to gateways with the lower number of hops. An AP finalizes its route based on link quality and number of hops.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead.

APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway and will switch two-way data communication from a parent node to a second parent node based on the functioning parameters of the wireless mesh network, including variable link quality and number of hops.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using parameters (or criteria based on those parameters) to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm process and background scanning both assist the AP in maintaining a list of neighbors and switching to the best parent on-channel or off-channel. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. To the extent this is considered an element under 35 U.S.C. §112 (¶ 6), in accordance with P.R. 3-1(c), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function include, *e.g.*, the software algorithms in Cisco's proprietary channel selection algorithms, Meraki Mesh Algorithm and gateway selection identified above. MeshDynamics's investigation into Cisco's proprietary algorithms remains ongoing and MeshDynamics will supplement or amend its P.R. 3-1(c) identification following discovery. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1f) wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset;	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprises wherein the communication criteria (*e.g.,* Meraki Mesh Algorithm's criteria, including the gateway selection logic and routing metric calculations, and channel selection algorithm criteria) comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset (*e.g.*, a dataset comprising "the best-known path" to the gateway which was received from a parent mesh access point). Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway ("root node").

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

	7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**: **Mesh routes** **Route** | **Avg Mbps** | **Avg metric** | **Usage** Repeater > Gateway | 186.6 | 1289 | 100% Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria dynamically and automatically find the best path back to the gateway or RAP. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
(1g) wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node;	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprises wherein one of the datasets contained in an AP node (*e.g.*, a dataset comprising "the best-known path") comprises a route path dataset comprising an identifier (*e.g.*, name and related information) for the associated parent node appended to the route path dataset for the associated parent node. APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:

Mesh routes

Route	Avg Mbps	Avg metric	Usage
Repeater > Gateway	186.6	1289	100%

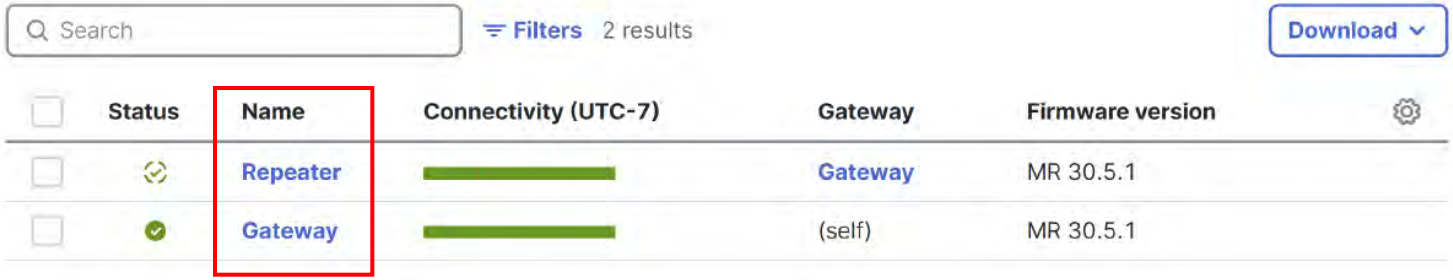
Source: Wireless Mesh Networking

Repeaters will identify the other AP in the network that it is using as its gateway, and its route back to the gateway, as well identify and track neighbors. Repeaters know the distance to their neighbors as well as the link quality. This is shown at the access server using identifier values to identify the repeaters, *e.g.*, "Name."

Identifying a Repeater access point vs Gateway access point

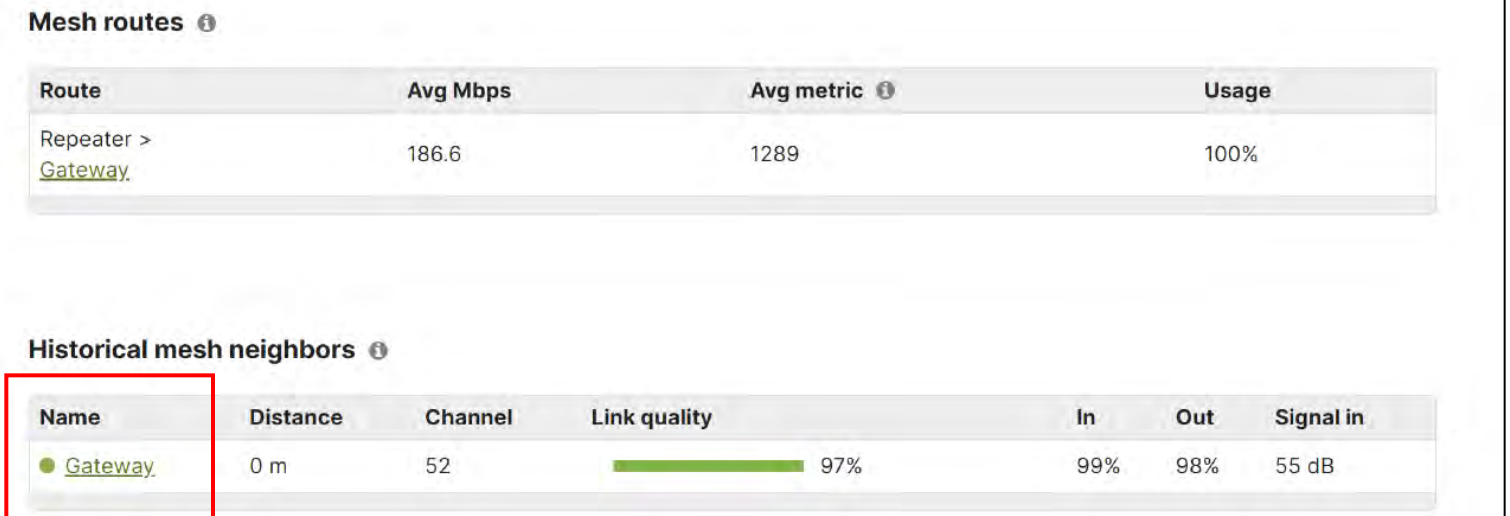
1. Navigate to **Wireless > Monitor > Access points**
2. Click the Gear icon on the right-hand side and make sure the **Gateway** option is selected:

3. A gateway access point will be listed as "(self)" under the **Gateway** column, while a repeater access point will list some other access point in the network (denoting that it is using that access point as its gateway):



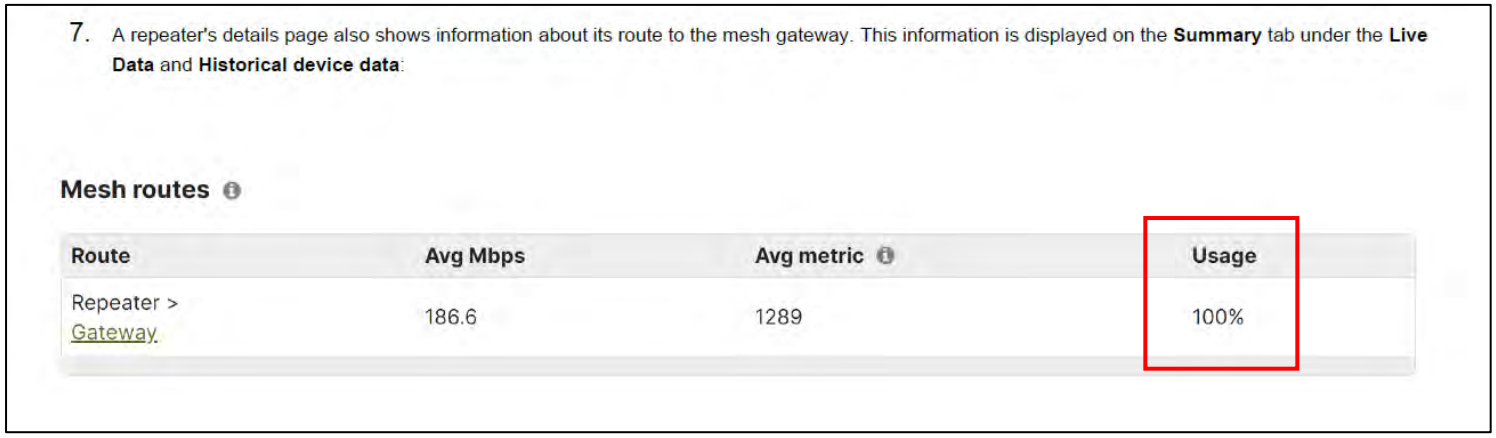
Source: Wireless Mesh Networking

7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:



	Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria find the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. The repeater's detail pages show information about its route to the mesh gateway and contains the name and identity of access points within the network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to..
(1h) wherein one of the datasets contained in the AP node comprises an amount of AP data traffic wherein the amount of AP data traffic comprises the amount of data exchange between the AP node and the zero or more clients in communication with the node and wherein another of the datasets contained in the AP node comprises a maximum capacity amount of the AP node.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprises wherein one of the datasets contained in the AP node comprises an amount of AP data traffic wherein the amount of AP data traffic comprises the amount of data exchange between the AP node and the zero or more clients in communication with the node and wherein another of the datasets contained in the AP node comprises a maximum capacity amount of the AP node. Cisco Meraki and Catalyst CW Access Points also comprise datasets with the amount of data exchanged between the AP node and the clients. On information and belief, the APs keep a log of client traffic, and the data exchanged between the node and the clients in communication with the node. APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route. To find the best path, the APs should know the capacity of the nodes in that path. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking



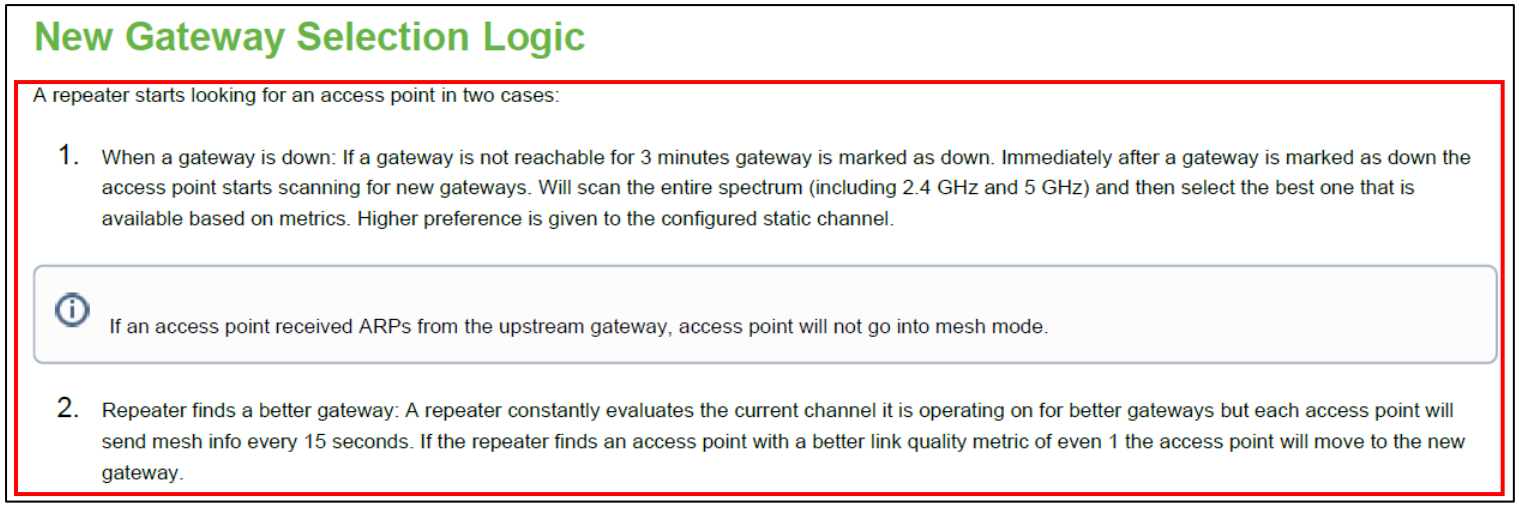
7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:

Mesh routes

Route	Avg Mbps	Avg metric	Usage
Repeater > Gateway	186.6	1289	100%

Source: Wireless Mesh Networking

Repeaters in a Cisco Wireless Mesh Network look for an AP in two cases: when the gateway is down or it finds a better gateway. AP will scan the entire spectrum in the first case. In the second, a repeater is constantly evaluating the current channel for better gateways. If a repeater finds an AP with a better link quality metric, it will move to the new gateway.



New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network scan all channels to collect information from its neighbors. If a valid neighbor is found, the AP switches to the channel.

Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. The number of hops is considered as well, with preference given to gateways with the lower number of hops. An AP finalizes its route based on link quality and number of hops.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

APs in the Cisco Meraki Wireless Mesh Network exchange mesh route information to allow repeaters to discover the best metric route to the gateway. On information and belief, this includes traffic and capacity data maintained by each AP.

Check for Mesh Neighbors

If two APs operate on the same channel and are within range of each other, they can become mesh neighbors. Mesh neighbors exchange mesh route information allowing repeaters to discover the best metric route to the gateway. If two APs are in the same Dashboard network and mesh neighbors, they may establish a mesh routing relationship where they send and receive client traffic of the mesh link.

This information can be immensely helpful in narrowing down bad mesh links and determining next steps.

Source: Troubleshooting Mesh Communication between Access Points

APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

The amounts of data traffic per route is kept on the AP and can be accessed at the cloud server, as shown in Cisco's documentation. This includes, on information and belief, the amount of data traffic and capacity at the AP.

Monitoring Mesh

Mesh monitoring tools are located at the bottom of every access point detail page, which can be accessed by navigating to **Wireless > Monitor > Access Points**, then clicking on an Access Point.

The image below shows an example access point acting as a repeater. The time selector at the top right-hand corner will adjust the timeframe of all of the UI components in the mesh monitoring section of the UI.

The time selector may select data from:

- 2 Hours
- 1 day
- 1 week
- 1 month



Mesh Routes

The Routes Table shows the routes used by different flows over time. As new routes are selected, they are added to the routes table. The overall amount of traffic per-route over the time period selected is shown in the Usage column. The Metric is also displayed in this table, representing a combination of loss and packet delivery times. Avg. Mbps throughput values are also provided for customers to gauge the capacity of that particular mesh route.

Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

	Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using parameters or criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria find the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. Meraki Mesh Algorithm's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 2	
2. The wireless mesh network of claim 1 wherein one of the datasets contained in the AP node comprises a cost to connect and said cost to connect is calculated using the level of network congestion experienced by the available parent node.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein one of the datasets contained in the AP node comprises a cost to connect and said cost to connect is calculated using the level of network congestion experienced by the available parent node. As shown in Cisco's documentation, the Meraki Mesh Algorithm uses a calculation of link quality caused by noise or multi-path interference, as well as performance impact of routing traffic through multiple hops to determine the preferred path. Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

Meraki Mesh Algorithm's and the AP's focus on factors like link quality/signal strength, reducing the number of hops and load value helps to minimize and avoid congestion in the mesh network. For example, lower throughput is an indication of congestion. By choosing the "best" path to the wired network based on these criteria, Meraki Mesh Algorithm's route path dataset addresses the level of network congestion experienced by an available parent node.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm

	process and criteria finds the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. Meraki Mesh Algorithm's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. Meraki Mesh Algorithm thus inherently addresses network congestion by selecting parents and paths that are less utilized and more reliable. This reduces the likelihood of bottlenecks and improves overall network performance. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 3	
3. The wireless mesh network of claim 1 wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein the cost to connect of the associated parent is the lowest of all available parent nodes.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprise wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein the cost to connect of the associated parent is the lowest of all available parent nodes. As shown in Cisco's documentation, the Meraki Mesh Algorithm uses a calculation of variable link quality caused by noise or multi-path interference, as well as performance impact of routing traffic through multiple hops to determine the preferred path, such that the preferred path (and parent) is the path with the best route to a wired gateway. Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, with the instructions that the lower metric ("cost to connect") equals the better route to select the parent.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

	7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**: **Mesh routes** **Route** | **Avg Mbps** | **Avg metric** | **Usage** Repeater > Gateway | 186.6 | 1289 | 100% Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm process and criteria finds the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 4	
4. The wireless mesh network of claim 1 wherein the communication criteria comprises instructions for the AP node to associate with an available parent node wherein the amount of AP data traffic is less	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein the communication criteria comprises instructions for the AP node to associate with an available parent node wherein the amount of AP data traffic is less than the sum of the maximum capacity amount of the available parent node and the available parent amount of AP data traffic. As discussed above, a parent offers the best route back to the gateway based on its calculation of variable link quality caused by noise or multi-path interference, as well as performance impact of routing traffic through multiple hops. Meraki Mesh Algorithm will calculate the best path back to the gateway.

than the sum of the maximum capacity amount of the available parent node and the available parent amount of AP data traffic.	Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: Wireless Mesh Networking APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, with the instructions that the lower metric equals the better route. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. Source: Wireless Mesh Networking APs in the Cisco Meraki Wireless Mesh Network exchange mesh route information to allow repeaters to discover the best metric route to the gateway. On information and belief, this includes traffic and capacity data maintained by each AP.

	Check for Mesh Neighbors If two APs operate on the same channel and are within range of each other, they can become mesh neighbors. Mesh neighbors exchange mesh route information allowing repeaters to discover the best metric route to the gateway. If two APs are in the same Dashboard network and mesh neighbors, they may establish a mesh routing relationship where they send and receive client traffic of the mesh link. This information can be immensely helpful in narrowing down bad mesh links and determining next steps. Source: Troubleshooting Mesh Communication between Access Points If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria find the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. Meraki Mesh Algorithm's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. Meraki Mesh Algorithm contributes to load balancing by selecting the path with the best quality and minimizing the number of hops. This helps to distribute traffic more effectively across the mesh network and potentially reduces congestion. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 5	
5. The wireless mesh network of claim 4 wherein the communication criteria comprises instructions for the AP node to associate with a single suitable parent node wherein a parent node is suitable if the throughput capacity	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 4 as shown above, comprises wherein the communication criteria comprises instructions for the AP node to associate with a single suitable parent node wherein a parent node is suitable if the throughput capacity of the parent node is the highest of all available parent nodes. Meraki Mesh Algorithm determines the parent that offers the best route back to the gateway based on a calculation of variable link quality caused by noise or multi-path interference, as well as performance impact of routing traffic through multiple hops. The parent with the lower number of hops will be chosen as the highest throughput capacity. If defined as "throughput," signal strength can also be used and the parent with the highest signal strength is chosen.

of the parent node is the highest of all available parent nodes.	Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: Wireless Mesh Networking APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, with the instructions that lower metric equals the better route. As shown above, highest throughput is part of the communication criteria's lower metric and instructions for suitable parent selection. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. Source: Wireless Mesh Networking

	7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**: **Mesh routes** **Route** | **Avg Mbps** | **Avg metric** | **Usage** Repeater > Gateway | 186.6 | 1289 | 100% Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria find the best path back to the gateway and chooses the parent with the least number of hops available or has the highest throughput capacity. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 6	
6. The wireless mesh network of claim 1 wherein one of the datasets contained in an AP node comprises a throughput capacity of the AP node.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein one of the datasets contained in an AP node comprises a throughput capacity of the AP node. Meraki Mesh Algorithm is run on an access point and automatically determines the best path back to the gateway for that access point by calculating the cost of the path in terms of signal strength and number of hops. Meraki Mesh Algorithm continuously monitors conditions and changes routes if necessary based on the changes in the conditions.

Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead.

APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

> **Meraki Mesh Algorithm**
>
> Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.
>
> As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric equals the better route. On information and belief, a repeater node contains a dataset with its throughput capacity based on its number of hops or signal strength/distance from the gateway so it can determine whether a better path exists.

> Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.
>
> In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

	7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**: **Mesh routes**  Route / Avg Mbps / Avg metric / Usage Repeater > Gateway / 186.6 / 1289 / 100% Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm process and criteria finds the best path back to the gateway and chooses the parent with the least number of hops available or has the highest throughput capacity. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 7	
7. The wireless mesh network of claim 1 wherein the means for switching two-way data communications selects the second associated parent node when the first associated parent node amount of AP data	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein the means for switching two-way data communications selects the second associated parent node when the first associated parent node amount of AP data traffic approaches the parent node maximum capacity amount within a congestion value set by the access server. As discussed above, Meraki Mesh Algorithm and the AP's focus on factors like link quality/signal strength and reducing the number of hops and load value helps to minimize and avoid congestion in the mesh network. By choosing the "best" path to the wired network based on these criteria, Meraki Mesh Algorithm's route path dataset addresses the level of congestion experienced by an available parent node.

traffic approaches the parent node maximum capacity amount within a congestion value set by the access server.	APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: Wireless Mesh Networking APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference possible using channel selection algorithms. Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels A repeater can also determine (*e.g.*, for parent optimization and upon the loss of a parent connection) that it should change channels and connect with a better potential parent across any channel to maintain its uplink with the best parent (*i.e.*, the best path back to the gateway based on signal strength and number of hops/ease criteria). The channel selection change will choose the best parent. Repeaters in a Cisco Wireless Mesh Network look for an AP in two cases: when the gateway is down or it finds a better gateway. AP will scan the entire spectrum in the first case. In the second, a repeater is constantly evaluating the current channel for better gateways. If a repeater finds an AP with a better link quality metric, it will move to the new gateway.

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

> ⓘ If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network scan all channels to collect information from its neighbors. If a valid neighbor is found, the AP switches to the channel.

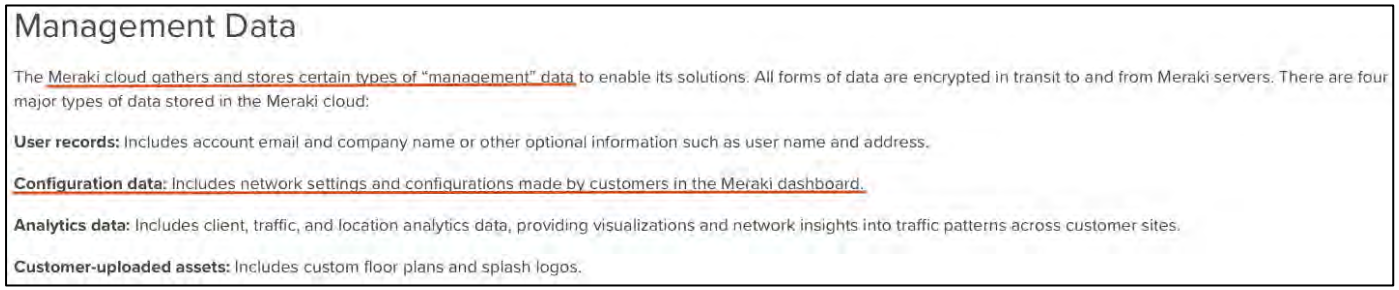
Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. The number of hops is considered as well, with preference given to gateways with the lower number of hops. An AP finalizes its route based on link quality and number of hops.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

	The Cisco Meraki Wireless Mesh Network includes a Meraki cloud server, which is in data communication with the access points/nodes, *e.g.*, the MR and CW Access Points, and is used to manage, configure, and monitor the access points/nodes:  Management Data The Meraki cloud gathers and stores certain types of "management" data to enable its solutions. All forms of data are encrypted in transit to and from Meraki servers. There are four major types of data stored in the Meraki cloud: **User records:** Includes account email and company name or other optional information such as user name and address. **Configuration data:** Includes network settings and configurations made by customers in the Meraki dashboard. **Analytics data:** Includes client, traffic, and location analytics data, providing visualizations and network insights into traffic patterns across customer sites. **Customer-uploaded assets:** Includes custom floor plans and splash logos. A maximum capacity amount within a congestion value is set by the access server by monitoring the mesh traffic and setting deployment considerations to address congestion and traffic patterns. This is used by the Meraki Mesh Algorithm to select the new parent mode when the AP data traffic approaches the parent node maximum capacity.

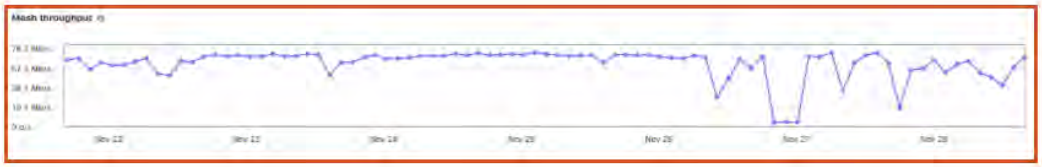
Monitoring Mesh

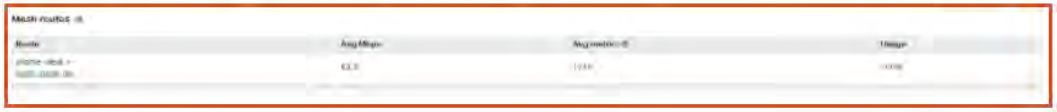
Mesh monitoring tools are located at the bottom of every access point detail page, which can be accessed by navigating to **Wireless > Monitor > Access Points**, then clicking on an Access Point.

The image below shows an example access point acting as a repeater. The time selector at the top right-hand corner will adjust the timeframe of all of the UI components in the mesh monitoring section of the UI.

The time selector may select data from:

- 2 Hours
- 1 day
- 1 week
- 1 month





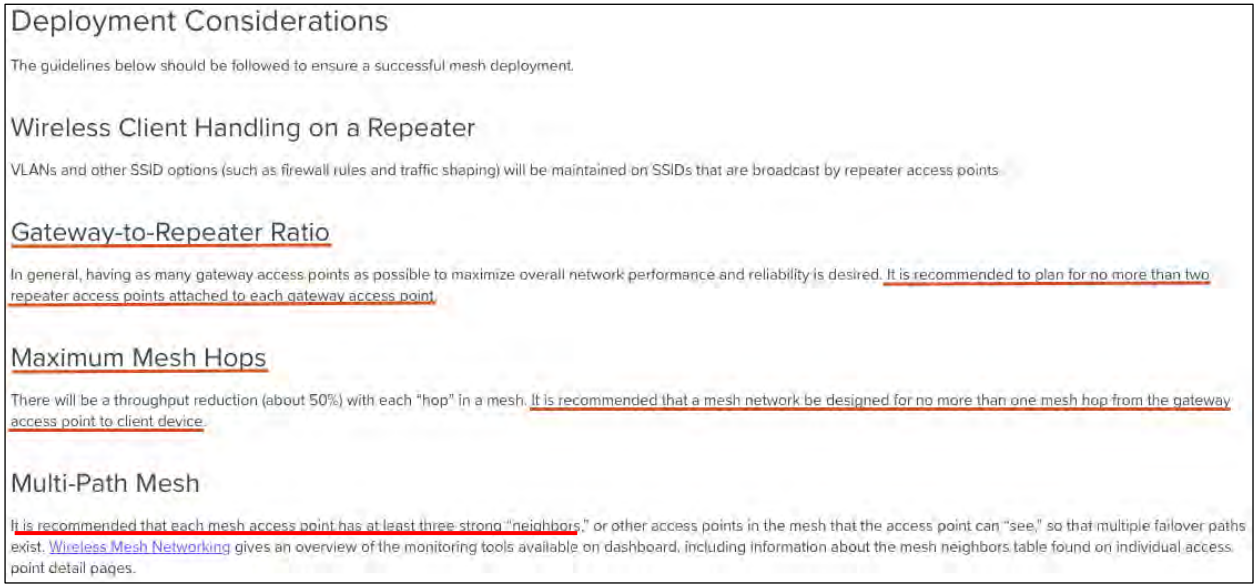
Mesh Routes

The Routes Table shows the routes used by different flows over time. As new routes are selected, they are added to the routes table. The overall amount of traffic per-route over the time period selected is shown in the Usage column. The Metric is also displayed in this table, representing a combination of loss and packet delivery times. Avg. Mbps throughput values are also provided for customers to gauge the capacity of that particular mesh route.

Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking



Deployment Considerations

The guidelines below should be followed to ensure a successful mesh deployment.

Wireless Client Handling on a Repeater

VLANs and other SSID options (such as firewall rules and traffic shaping) will be maintained on SSIDs that are broadcast by repeater access points.

Gateway-to-Repeater Ratio

In general, having as many gateway access points as possible to maximize overall network performance and reliability is desired. It is recommended to plan for no more than two repeater access points attached to each gateway access point.

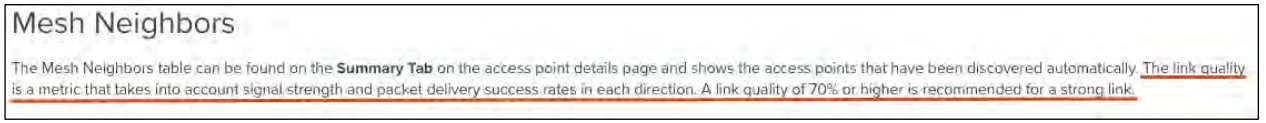
Maximum Mesh Hops

There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device.

Multi-Path Mesh

It is recommended that each mesh access point has at least three strong "neighbors," or other access points in the mesh that the access point can "see," so that multiple failover paths exist. Wireless Mesh Networking gives an overview of the monitoring tools available on dashboard, including information about the mesh neighbors table found on individual access point detail pages.

Source: Mesh Deployment Guide

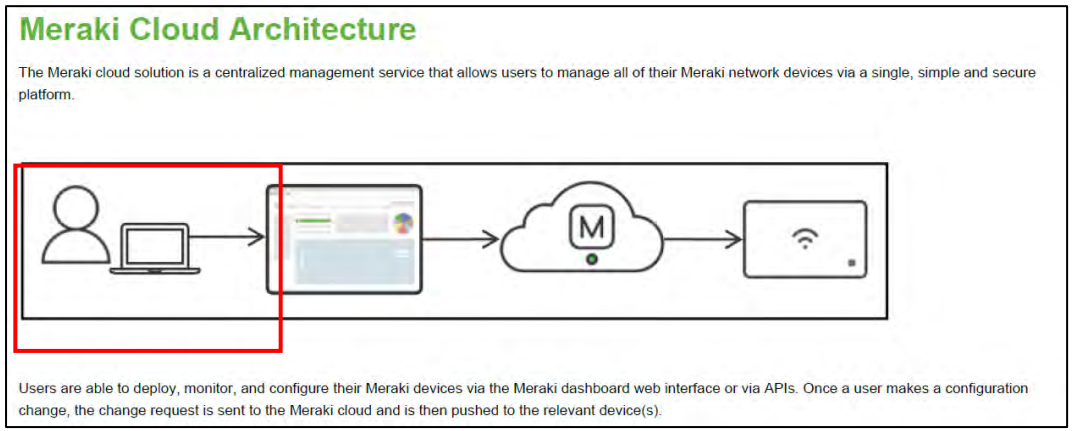


Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by using criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria find the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. Meraki Mesh Algorithm's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. Meraki Mesh Algorithm thus inherently addresses congestion by selecting parents and paths that are less utilized and more reliable. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	To the extent this is considered an element under 35 U.S.C. §112 (¶ 6), in accordance with P.R. 3-1(c), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function include, *e.g.*, the software algorithms in Cisco's proprietary Meraki Mesh Algorithm and gateway selection logic for parent selection identified above. MeshDynamics's investigation into Cisco's proprietary algorithms remains ongoing and MeshDynamics will supplement or amend its P.R. 3-1(c) identification following discovery. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 8	
8. The wireless mesh network of claim 1 wherein the wireless communication with client includes a means to adjust a signal strength of the AP node wireless communication with clients and wherein the means for adjusting the signal strength of the AP node decreases the signal strength of the AP node when the amount of data sent from the zero or more clients exceeds a data congestion threshold value.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises wherein the wireless communication with client includes a means to adjust a signal strength of the AP node wireless communication with clients and wherein the means for adjusting the signal strength of the AP node decreases the signal strength of the AP node when the amount of data sent from the zero or more clients exceeds a data congestion threshold value (*e.g.*, client balancing). The AP node (*e.g.*, Meraki access point) is in wireless communication with zero or more clients (*e.g.*, laptops) as shown in Cisco's documentation:  Source: Meraki Cloud Architecture Cisco's Meraki Wireless Mesh Network uses Client Balancing as a means to balance clients and optimize overall network performance and prevent overloading an AP node:

Client Balancing

Last updated: Jul 16, 2025

Click 日本語 for Japanese

+ TABLE OF CONTENTS

Overview

Client balancing aims to steer devices to associate with the best available access point (AP). When a client attempts to connect to a busier AP, the AP rejects the connection, prompting the client to connect to a different AP until it finds the one with the least usage.

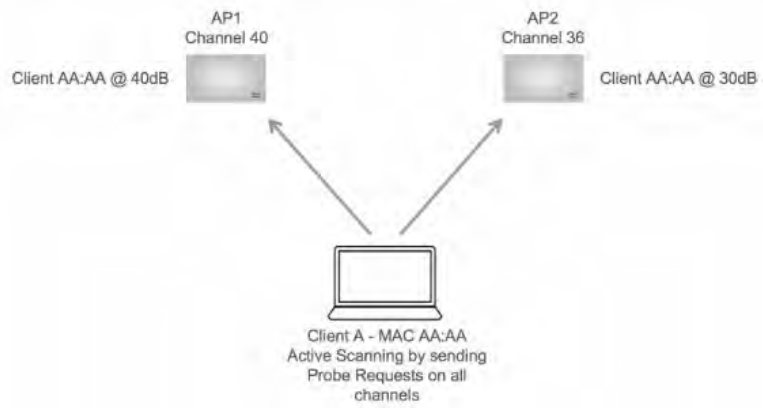
Client Balancing is not enabled by default. Some clients and applications may not handle rejections gracefully when Client Balancing is active. Therefore, it is recommended to enable Client Balancing only after thorough evaluation and validation to ensure compatibility with your specific clients and applications.

Depending on the firmware version and client type, the AP rejecting the association might provide specific information about the least loaded AP, giving the client a higher chance of connecting to it on the second attempt. When the client is already connected, the AP may also advise the client to move to another AP with more resources.

Client Balancing Behavior Prior to MR 29.X Firmware

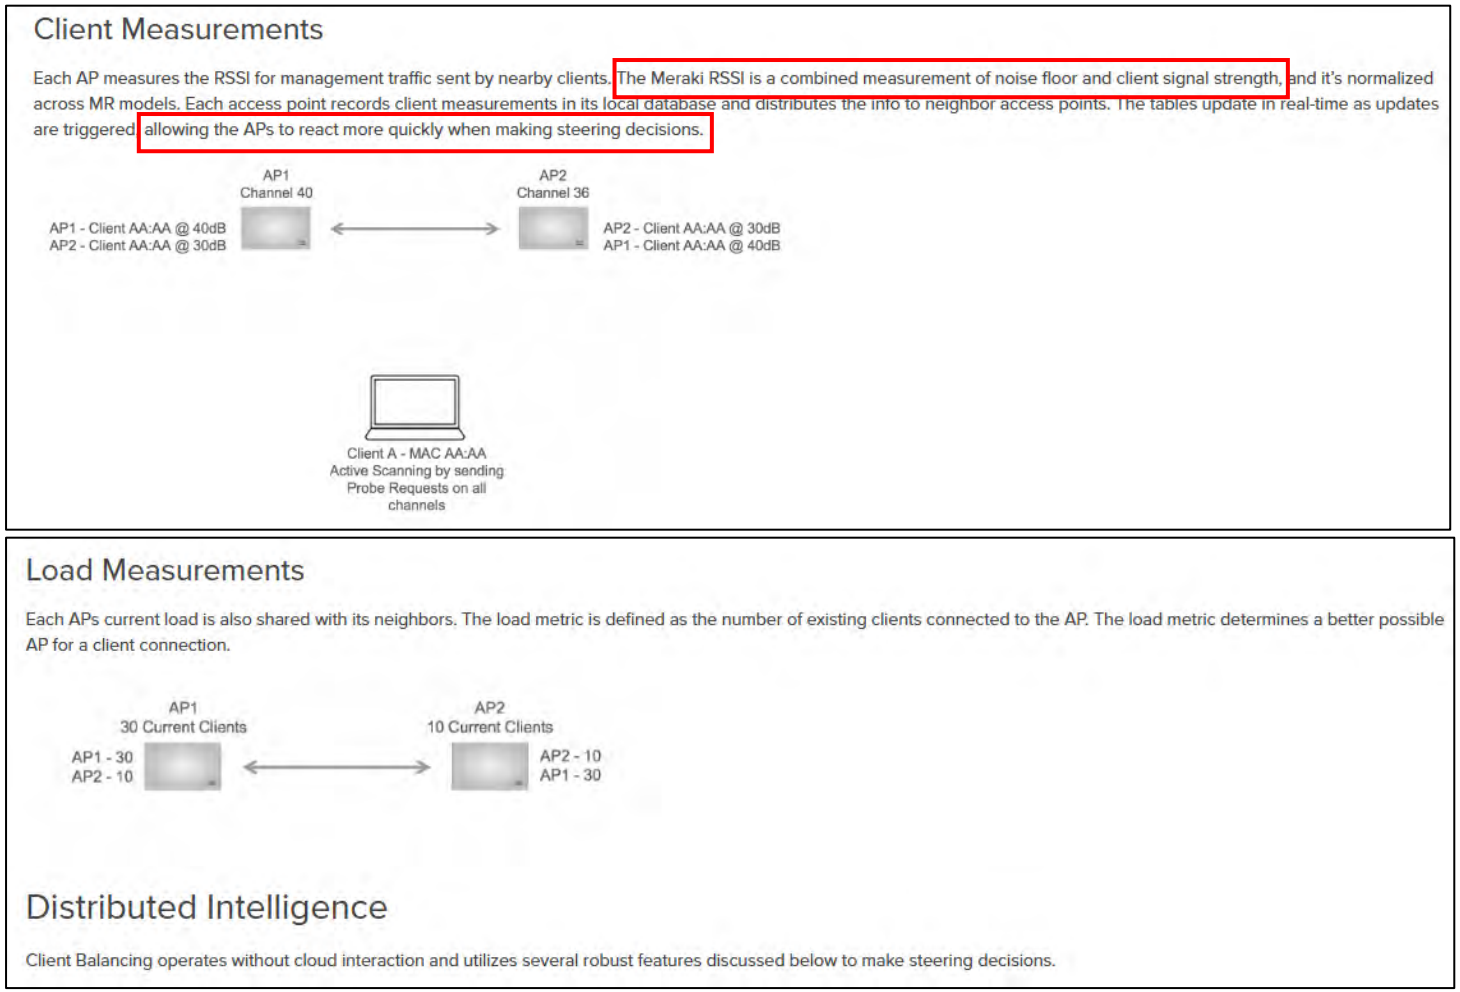
Metrics

Several metrics are gathered by each AP to ensure that the best client and AP match is made. RF environment metrics and the algorithm's wireless clients heard by APs are factors. These measurements are captured in real-time and shared amongst nearby APs in the Meraki network.



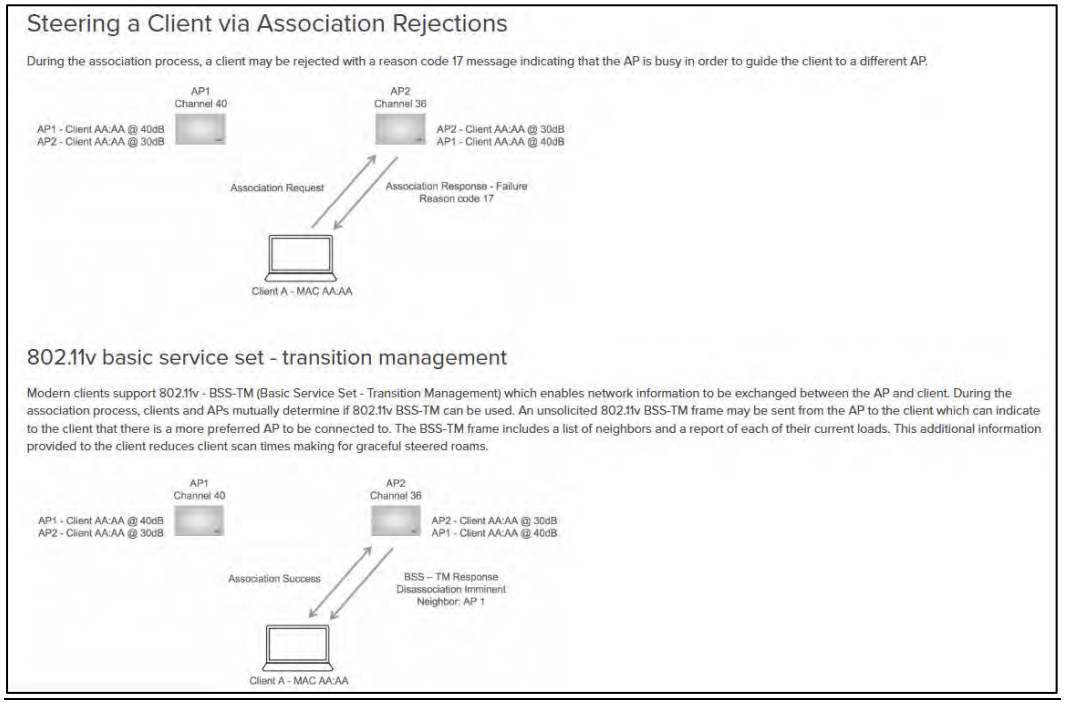
Source:
https://documentation.meraki.com/MR/Operate_and_Maintain/How_Tos/Features_and_Integrations/Client_Balancing

On information and belief, Cisco's Meraki APs will monitor throughput and will steer to another node as part of client balancing, including client connection limits set at the controller:



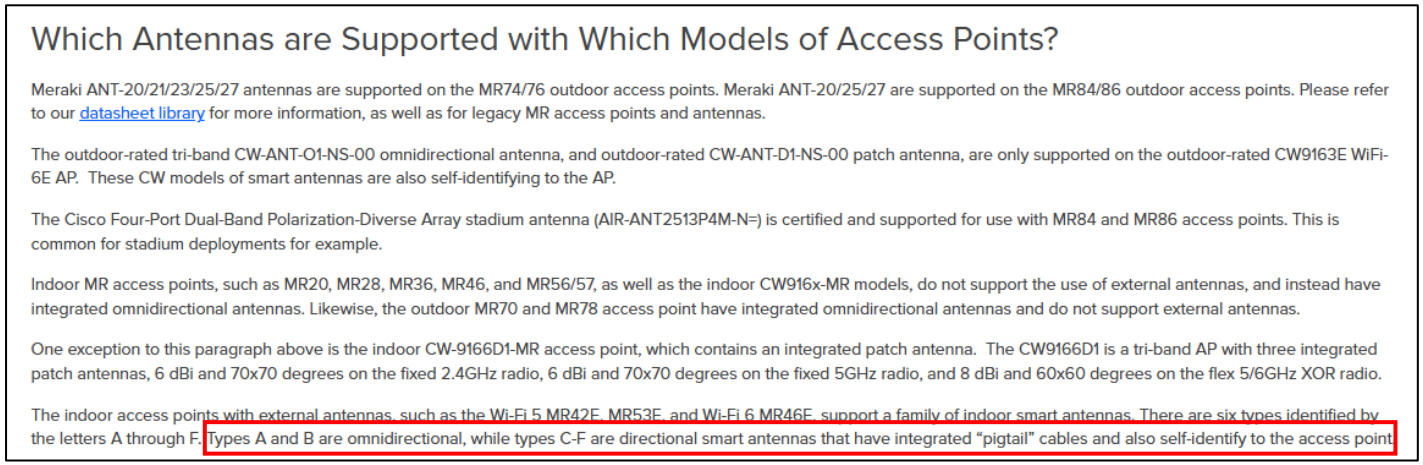
Source:
https://documentation.meraki.com/MR/Operate_and_Maintain/How_Tos/Features_and_Integrations/Client_Balancing

Cisco Meraki access points have a means for adjusting the signal strength of the AP node that decreases the signal strength of the AP node transmitted outward from the antennas towards a client:



Source:
https://documentation.meraki.com/MR/Operate_and_Maintain/How_Tos/Features_and_Integrations/Client_Balancing

On information and belief, Cisco effectively changes the strength of signal to the clients with the available transmit power settings to steer the beam via directional antennas, or by decreasing the transmit power settings strength of the signal to dissuade clients from connecting to the AP node. For example, when traffic from one or more clients is high and the directional antennas are steering the RF beam to those clients, other clients see a reduced signal strength which may cause those clients to seek a different AP node. Additionally, the AP node connected to the client may decrease the output signal strength to encourage clients to seek a different AP node to select.



Which Antennas are Supported with Which Models of Access Points?

Meraki ANT-20/21/23/25/27 antennas are supported on the MR74/76 outdoor access points. Meraki ANT-20/25/27 are supported on the MR84/86 outdoor access points. Please refer to our datasheet library for more information, as well as for legacy MR access points and antennas.

The outdoor-rated tri-band CW-ANT-O1-NS-00 omnidirectional antenna, and outdoor-rated CW-ANT-D1-NS-00 patch antenna, are only supported on the outdoor-rated CW9163E WiFi-6E AP. These CW models of smart antennas are also self-identifying to the AP.

The Cisco Four-Port Dual-Band Polarization-Diverse Array stadium antenna (AIR-ANT2513P4M-N=) is certified and supported for use with MR84 and MR86 access points. This is common for stadium deployments for example.

Indoor MR access points, such as MR20, MR28, MR36, MR46, and MR56/57, as well as the indoor CW916x-MR models, do not support the use of external antennas, and instead have integrated omnidirectional antennas. Likewise, the outdoor MR70 and MR78 access point have integrated omnidirectional antennas and do not support external antennas.

One exception to this paragraph above is the indoor CW-9166D1-MR access point, which contains an integrated patch antenna. The CW9166D1 is a tri-band AP with three integrated patch antennas, 6 dBi and 70x70 degrees on the fixed 2.4GHz radio, 6 dBi and 70x70 degrees on the fixed 5GHz radio, and 8 dBi and 60x60 degrees on the flex 5/6GHz XOR radio.

The indoor access points with external antennas, such as the Wi-Fi 5 MR42E, MR53E, and Wi-Fi 6 MR46E, support a family of indoor smart antennas. There are six types identified by the letters A through F. Types A and B are omnidirectional, while types C-F are directional smart antennas that have integrated "pigtail" cables and also self-identify to the access point.

Source: https://documentation.meraki.com/MR/Operate_and_Maintain/FAQs/Frequently_Asked_Questions_regarding_Cisco_Meraki_Antennas

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting an appropriate AP for a client in the same way by maintaining a list of neighboring AP's and using parameters or criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The Meraki Mesh Algorithm and criteria find the best path back to the gateway, and the repeater learns all available neighbors and determines the route path to the gateway. Client balancing aims to steer devices to associate with the best available APs. Meraki Mesh Algorithm's calculations are continuously updated, allowing network connectivity and paths to change as the traffic patterns on wireless links change. Meraki Mesh Algorithm thus inherently addresses congestion by selecting parents and paths that are less utilized and more reliable. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

To the extent this is considered an element under 35 U.S.C. §112 (¶ 6), in accordance with P.R. 3-1(c), the identity of the structure(s), act(s), or material(s) in the Accused Instrumentality that performs the claimed function include, *e.g.*, the software algorithms in Cisco's proprietary Meraki Mesh Algorithm, and Client balancing selection logic for client selection identified above. MeshDynamics's investigation into Cisco's proprietary algorithms remains ongoing and MeshDynamics will supplement or amend its P.R. 3-1(c) identification following discovery.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 9	
A wireless mesh network comprising:	*See* (1pre) above, the contents of which are incorporated herein by reference.
an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	*See* (1a) above, the contents of which are incorporated herein by reference.
one or more root nodes connected to said access server and an external network;	*See* (1b) above, the contents of which are incorporated herein by reference.
one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria or the associated parent node is a root node	*See* (1c) above, the contents of which are incorporated herein by reference.

within wireless communication range;	
wherein an AP node is in wireless communication with zero or more clients; and	*See* (1d) above, the contents of which are incorporated herein by reference.
wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets	*See* (1e) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node	*See* (1g) above, the contents of which are incorporated herein by reference.
wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available parent node is	*See* (1f) above, the contents of which are incorporated herein by reference.

selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset; and	
wherein one of the datasets contained in an AP node comprises a dataset of child node identifiers wherein the dataset of child node identifiers is a dataset identifying each AP node in wireless communication with the AP.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, comprises wherein one of the datasets contained in an AP node comprises a dataset of child node identifiers wherein the dataset of child node identifiers is a dataset identifying each AP node in wireless communication with the AP. As confirmed in Cisco's documentation, a list of nodes identifying each AP node in RF range, including those in wireless communication, is obtained and maintained on each access point. APs within the Cisco Meraki Wireless Mesh Network operate as either a gateway or repeater and are capable of relaying traffic from a child node associated with the parent node. In Cisco's Meraki Wireless Mesh Network, APs can be in one of two states: gateway or repeater. Repeater APs are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide Both gateways and repeaters can serve clients. Repeaters will automatically choose the gateway to which it has the strongest connection.

Note: Both gateways and repeaters can serve wireless clients. It is possible to have multiple gateways in a mesh network, and repeaters will automatically choose the gateway to which it has the strongest connection.

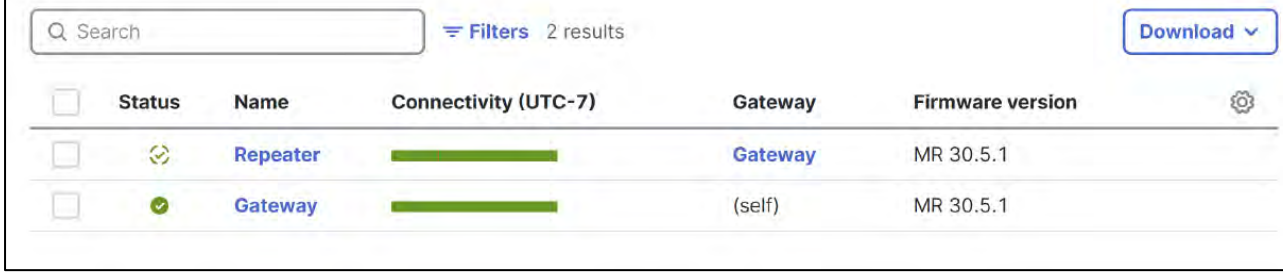
Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

Repeaters will identify the other AP in the network that it is using as its gateway, and its route back to the gateway, as well identify and track neighbors. Repeaters know the distance to its neighbors as well as the link quality. This is shown at the access server using identifier values to identify the repeaters, *e.g.*, "Name."

Identifying a Repeater access point vs Gateway access point

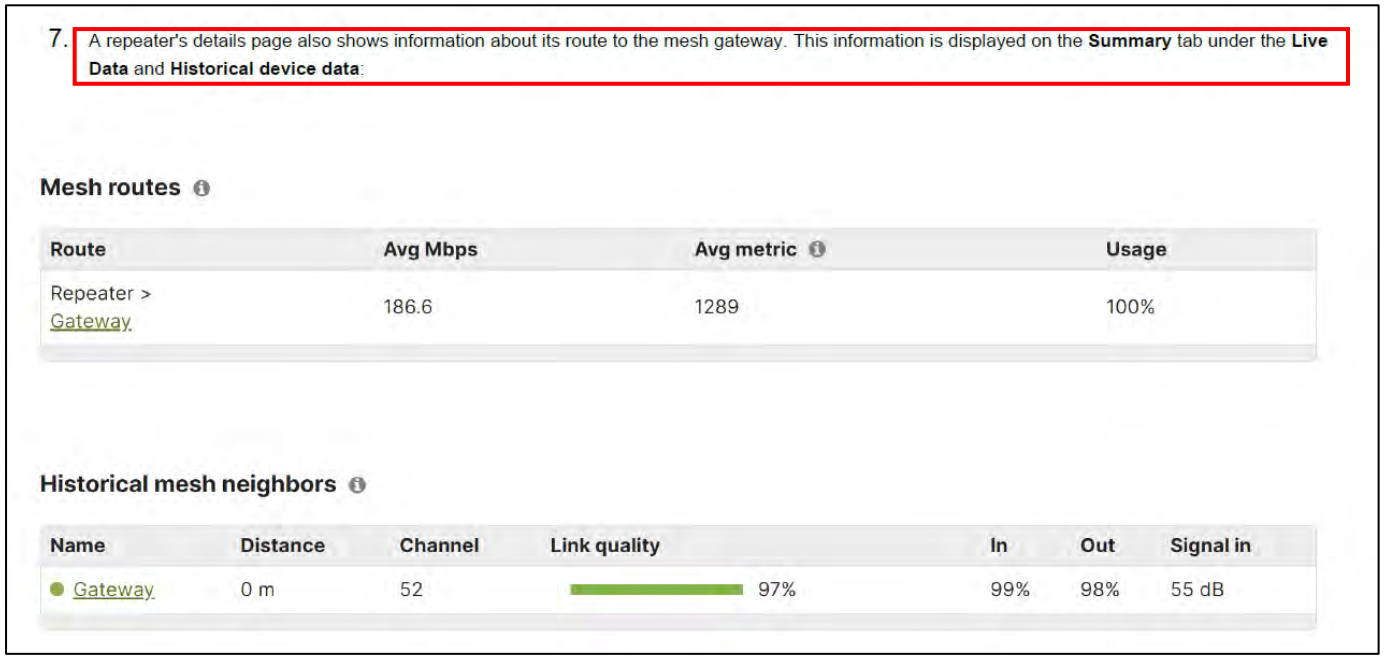
1. Navigate to **Wireless > Monitor > Access points**
2. Click the Gear icon on the right-hand side and make sure the **Gateway** option is selected:

3. A gateway access point will be listed as "(self)" under the **Gateway** column, while a repeater access point will list some other access point in the network (denoting that it is using that access point as its gateway):



Source: Wireless Mesh Networking

7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:



Mesh routes

Route	Avg Mbps	Avg metric	Usage
Repeater > Gateway	186.6	1289	100%

Historical mesh neighbors

Name	Distance	Channel	Link quality	In	Out	Signal in
Gateway	0 m	52	97%	99%	98%	55 dB

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network communicate traffic among themselves and back to the gateway using a proprietary protocol. Each AP is available for connections from repeaters if connected as a gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

The APs scan to automatically detect each other and keep routing tables and list of available neighbors, and have a dataset of "child node identifiers" identifying each AP node in wireless communication with the AP.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Mesh Routes

The Routes Table shows the routes used by different flows over time. As new routes are selected, they are added to the routes table. The overall amount of traffic per-route over the time period selected is shown in the Usage column. The Metric is also displayed in this table, representing a combination of loss and packet delivery times. Avg. Mbps throughput values are also provided for customers to gauge the capacity of that particular mesh route.

Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Each access point in the wireless mesh network maintains a list of nodes within the RF range of each other and thus includes all child nodes in the network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

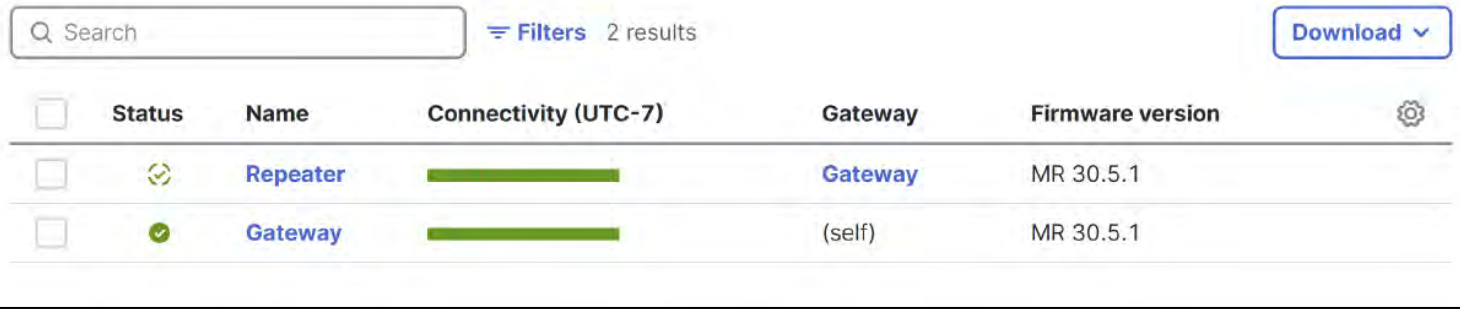
Claim 10	
10. The wireless mesh network of claim 9 wherein the dataset of child node identifiers is accessible to each child node.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 9 as shown above, comprises wherein the dataset of child node identifiers is accessible to each child node. As discussed above, each access point in the wireless mesh network is maintaining a list of nodes within the RF range of each other, and in wireless communication, and thus includes all child nodes in the network. Access points communicate among themselves, and on information and belief, the list of child node identifiers that identifies each AP node in wireless communication with the AP is accessible to each child node. APs within the Cisco Meraki Wireless Mesh Network operate as either a gateway or repeater and are capable of relaying traffic from a child node associated with the parent node. In Cisco's Meraki Wireless Mesh Network, APs can be in one of two states: gateway or repeater. Repeater APs are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide Both gateways and repeaters can serve clients. Repeaters will automatically choose the gateway to which it has the strongest connection. Note: Both gateways and repeaters can serve wireless clients. It is possible to have multiple gateways in a mesh network, and repeaters will automatically choose the gateway to which it has the strongest connection. Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

Repeaters will identify the other AP in the network that it is using as its gateway, and its route back to the gateway, as well identify and track neighbors. Repeaters know the distance to its neighbors as well as the link quality. This is shown at the access server using identifier values to identify the repeaters, *e.g.*, "Name."

Identifying a Repeater access point vs Gateway access point

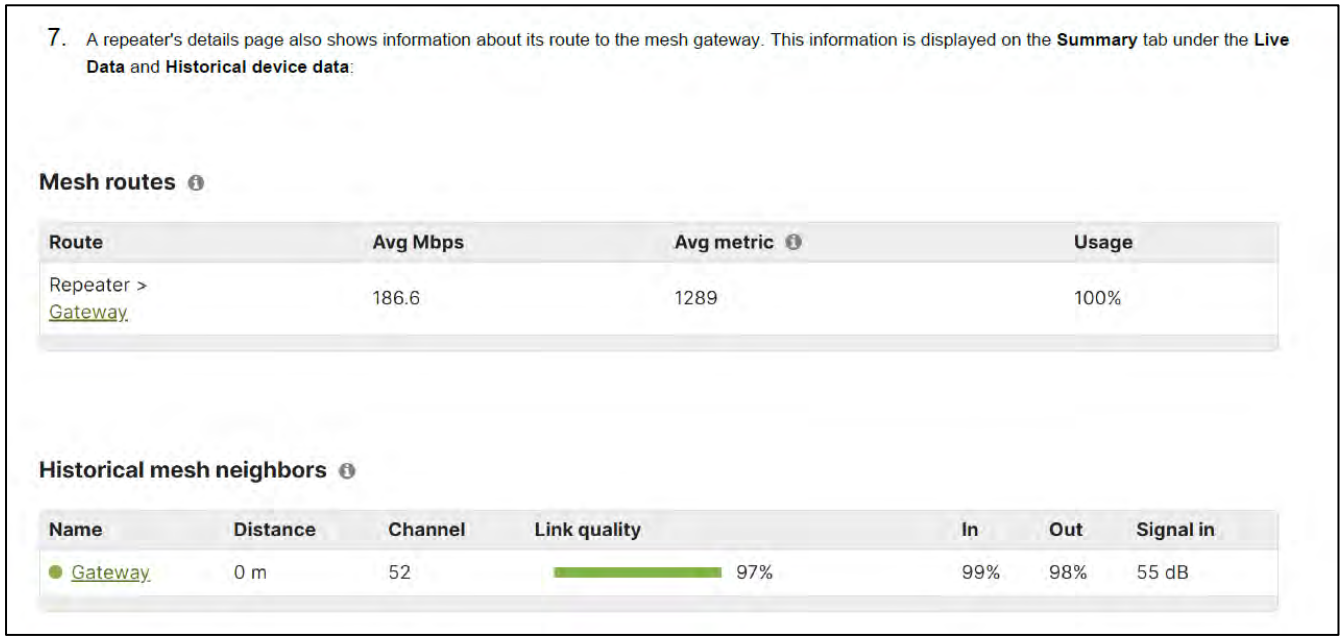
1. Navigate to **Wireless > Monitor > Access points**
2. Click the Gear icon on the right-hand side and make sure the **Gateway** option is selected:

3. A gateway access point will be listed as "(self)" under the **Gateway** column, while a repeater access point will list some other access point in the network (denoting that it is using that access point as its gateway):



Source: Wireless Mesh Networking

7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:



Mesh routes

Route	Avg Mbps	Avg metric	Usage
Repeater > Gateway	186.6	1289	100%

Historical mesh neighbors

Name	Distance	Channel	Link quality	In	Out	Signal in
Gateway	0 m	52	97%	99%	98%	55 dB

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network communicate traffic among themselves and back to the gateway using a proprietary protocol. Each AP is available for connections from repeaters if connected as a gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

The APs scan to automatically detect each other and keep routing tables and list of available neighbors. All information is accessible to each AP as the APs are constantly communicating.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Mesh Routes

The Routes Table shows the routes used by different flows over time. As new routes are selected, they are added to the routes table. The overall amount of traffic per-route over the time period selected is shown in the Usage column. The Metric is also displayed in this table, representing a combination of loss and packet delivery times. Avg. Mbps throughput values are also provided for customers to gauge the capacity of that particular mesh route.

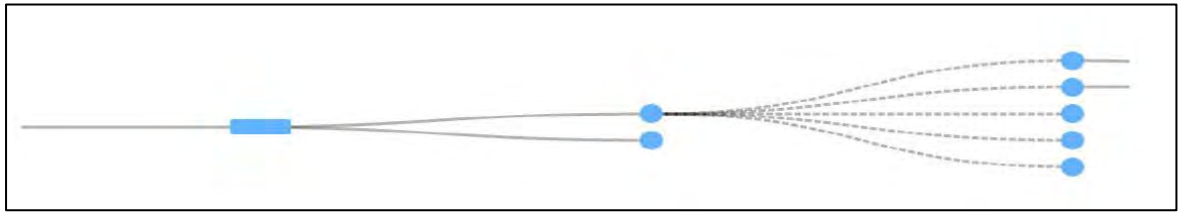
Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Each access point in the wireless mesh network is maintaining a list of nodes within the RF range of each other and thus includes all child nodes in the network. Access points communicate among themselves. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

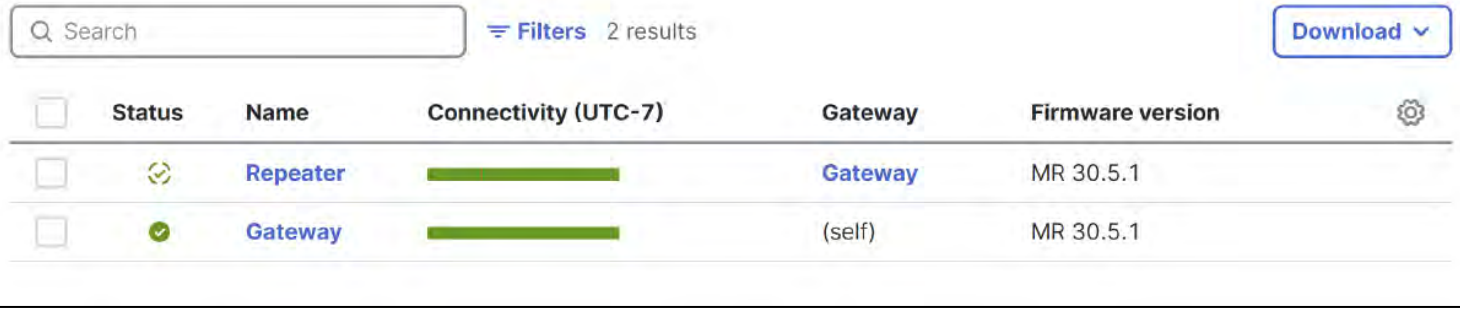
Claim 11	
11. The wireless mesh network of claim 10 wherein the AP node dataset of child nodes contains two or more child nodes wherein the zero or more clients in communication with a first child node sends data wherein a destination of the data is a second child node, the first child node sends the data directly to the second child node.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 10 as shown above, comprises wherein the AP node dataset of child nodes contains two or more child nodes wherein the zero or more clients in communication with a first child node sends data wherein a destination of the data is a second child node, the first child node sends the data directly to the second child node. As discussed above, repeaters are aware of neighbors and thus can communicate directly over wireless. It is possible for a repeater to send data to a destination within a second repeater and that data will go directly from repeater 1 to repeater 2. Cisco's Meraki Wireless Mesh Network comprise a tree-shaped network. Multiple repeaters may connect to a single gateway as shown in the example deployment mode below.  Source: Mesh Deployment Guide Cisco's Meraki Wireless Mesh Network utilizes a mesh algorithm to determine the best route between APs and a gateway, and the network topology takes on a tree-like structure. The gateways form the "root" of the tree, providing the wired connection. Repeaters branch out from the gateways, creating wireless links to extend coverage and will connect to another repeater. The mesh algorithm finds the most efficient path, and the overall structure will resemble a tree. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: Wireless Mesh Networking;

Repeaters will identify the other AP in the network that it is using as its gateway, and its route back to the gateway, as well identify and track neighbors. Repeaters know the distance to their neighbors as well as the link quality. This is shown at the access server using identifier values to identify the repeaters, *e.g.*, "Name."

Identifying a Repeater access point vs Gateway access point

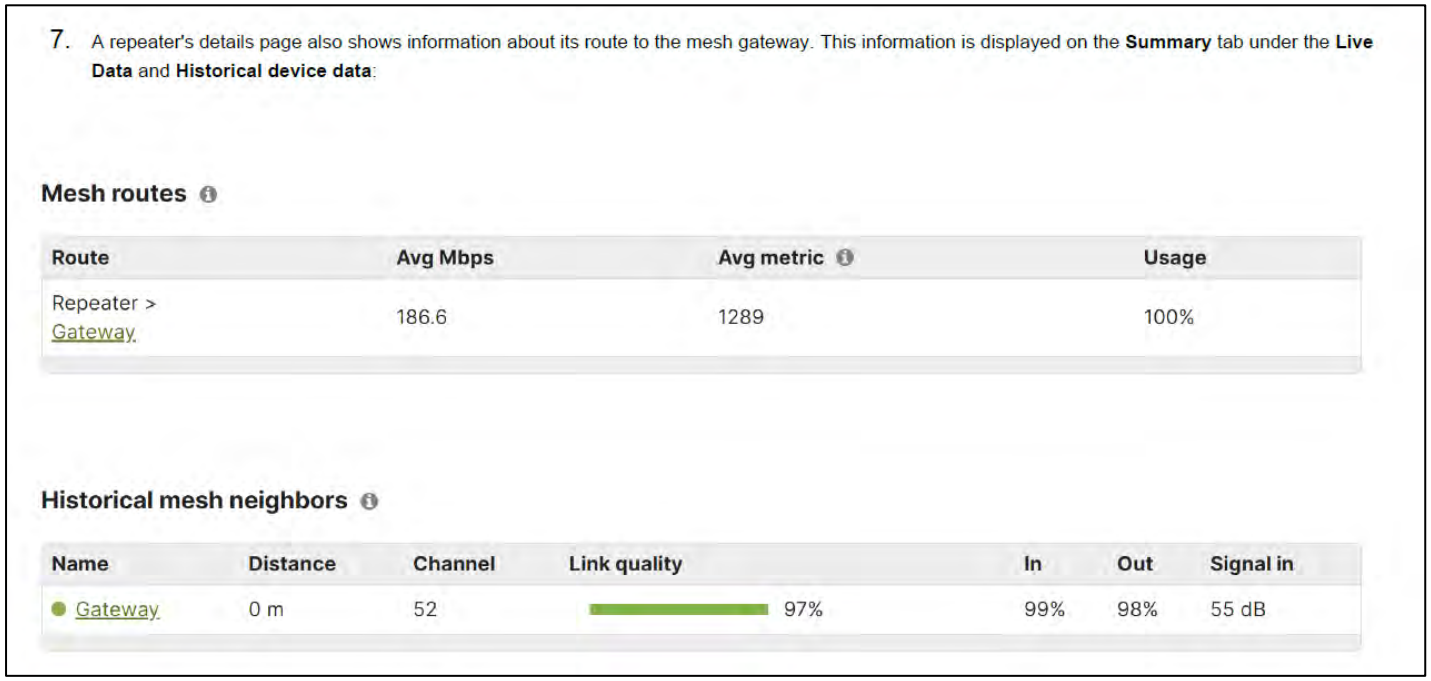
1. Navigate to **Wireless > Monitor > Access points**
2. Click the Gear icon on the right-hand side and make sure the **Gateway** option is selected:

3. A gateway access point will be listed as "(self)" under the **Gateway** column, while a repeater access point will list some other access point in the network (denoting that it is using that access point as its gateway):



Source: Wireless Mesh Networking

7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:



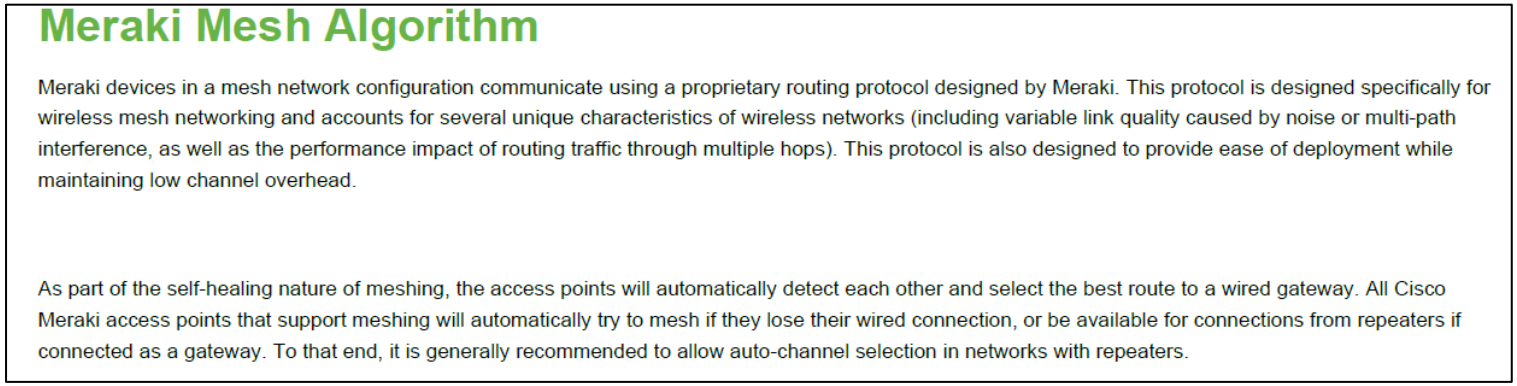
Mesh routes

Route	Avg Mbps	Avg metric	Usage
Repeater > Gateway	186.6	1289	100%

Historical mesh neighbors

Name	Distance	Channel	Link quality	In	Out	Signal in
Gateway	0 m	52	97%	99%	98%	55 dB

Source: Wireless Mesh Networking



Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of sending data traffic through mesh network in same way by

	maintaining a list of neighbors and finding the best path to the data traffic destination and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Rather than forcing data back up the hierarchical chain to a root node, and then back down the chain to a destination child node, the mesh is cooperative and since it is aware of the destination node will take the quickest path there (*i.e.*, directly from repeater 1 to repeater 2). A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 12	
A wireless mesh network comprising:	*See* (1pre) above, the contents of which are incorporated herein by reference.
an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	*See* (1a) above, the contents of which are incorporated herein by reference.
one or more root nodes connected to said access server and an external network	*See* (1b) above, the contents of which are incorporated herein by reference.
one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication	*See* (1c) above, the contents of which are incorporated herein by reference.

range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria or the associated parent node is a root node within wireless communication range;	
wherein an AP node is in wireless communication with zero or more clients; and	*See* (1d) above, the contents of which are incorporated herein by reference.
wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets	*See* (1e) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node	*See* (1g) above, the contents of which are incorporated herein by reference.

wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset; and	*See* (1f) above, the contents of which are incorporated herein by reference.
wherein the communication criteria further comprises instructions for the AP node to associate with a single suitable parent node wherein the route path dataset of the parent node is the shortest route path dataset of all available parent nodes.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprises wherein the communication criteria further comprises instructions for the AP node to associate with a single suitable parent node wherein the route path dataset of the parent node (*e.g.*, a dataset comprising "the best path back to a gateway") is the shortest route path dataset (*e.g.*, the route path dataset with the lowest number of hops back to the gateway) of all available parent nodes. The Meraki Mesh Algorithm criteria to choose the best parent includes instructions for the repeater or child node to associate with a gateway or parent, which on information and belief, as shown in Cisco's proprietary code, is the shortest route or path, *e.g.*, the best route with the least number of hops back to the gateway. Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead. APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric (which includes the least number of hops or shortest route) equals the better route.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

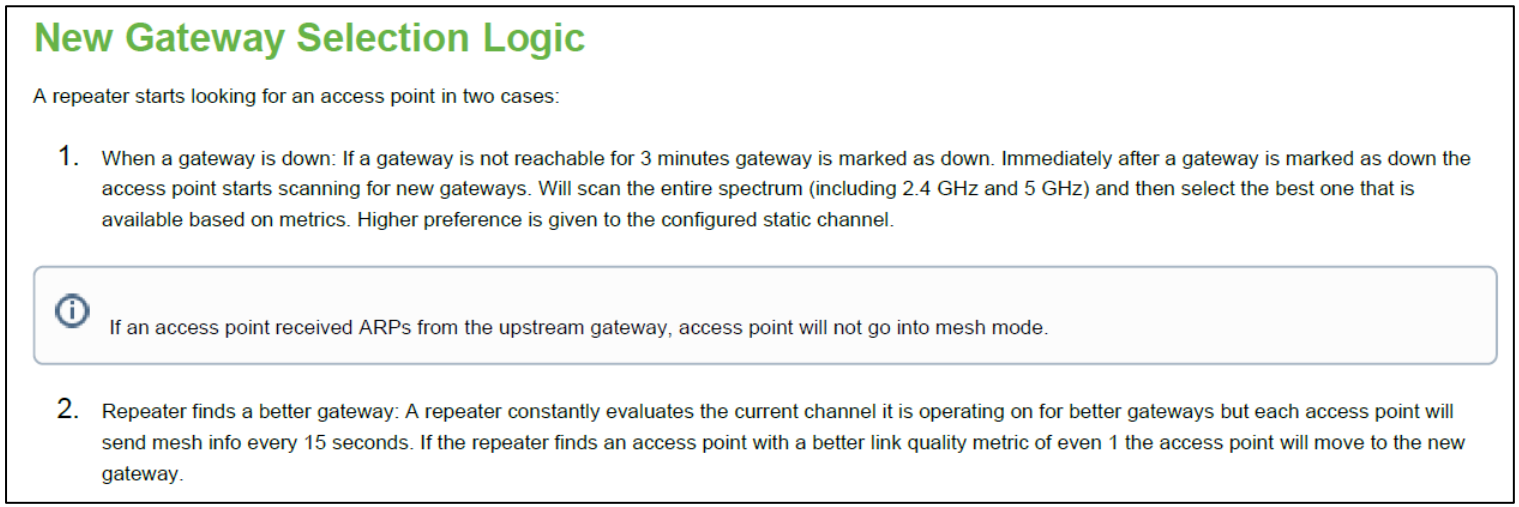
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AP node in a Cisco Meraki Wireless Mesh Network associates with a single suitable parent node when the route path of the parent node is the best path of all available parent nodes. It is the best path and shortest when it is the least number of hops, for example. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 13	
A wireless mesh network comprising:	*See* (1pre) above, the contents of which are incorporated herein by reference.
an access server wherein the access server sets one or more functioning parameters of the wireless mesh network;	*See* (1a) above, the contents of which are incorporated herein by reference.
one or more root nodes connected to said access server and an external network;	*See* (1b) above, the contents of which are incorporated herein by reference.
one or more AP nodes wherein each AP node is in wireless two-way data communication with an associated parent node wherein said associated parent node is selected from all available parent nodes wherein an available parent node is another AP node within wireless communication range of the AP node and the associated parent node is an available parent node meeting one or more communication criteria or the associated	*See* (1c) above, the contents of which are incorporated herein by reference.

parent node is a root node within wireless communication range;	
wherein an AP node is in wireless communication with zero or more clients; and	*See* (1d) above, the contents of which are incorporated herein by reference.
wherein an AP node includes a means for switching two-way data communication from a first associated parent node to a second associated parent node based on the functioning parameters of the wireless mesh network and wherein an AP node contains one or more datasets	*See* (1e) above, the contents of which are incorporated herein by reference.
wherein one of the datasets contained in an AP node comprises a route path dataset comprising an identifier for the associated parent node appended to the route path dataset for the associated parent node	*See* (1g) above, the contents of which are incorporated herein by reference.
wherein the communication criteria comprises instructions for the AP node to select the associated parent node wherein an available	*See* (1f) above, the contents of which are incorporated herein by reference.

parent node is selected to become the associated parent node if the available parent node is in wireless communication with a root node or if a root node is contained in the available parent node's route path dataset; and	
wherein the access server functioning parameters includes a latency modifier wherein the AP node means for switching from the first associated parent node to a second associated parent node result in selection of the second associated parent wherein the route path of the second associated parent node is shorter than the first associated route path by a value related to the latency modifier.	Cisco Meraki and Catalyst CW Access Points in Cisco's Meraki Wireless Mesh Network comprises wherein the access server functioning parameters includes a latency modifier wherein the AP node means for switching from the first associated parent node to a second associated parent node result in selection of the second associated parent wherein the route path of the second associated parent node is shorter than the first associated route path by a value related to the latency modifier. As discussed above in claim 1(a) above, functioning parameters set at the access server includes the signal strength and number of hops ("latency modifier"). Meraki Mesh Algorithm results in the selection of a parent with a best or shorter route path shorter by signal strength and number of hops. Meraki Mesh Algorithm determines the best route to a wired gateway on the unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference possible using channel selection algorithms. Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in. Source: Manually Changing Channels A repeater can also determine (*e.g.*, for parent optimization and upon the loss of a parent connection) that it should change channels and connect with a better potential parent across any channel to maintain its uplink with the best parent (*i.e.*, the best path back to the gateway based on signal strength and number of hops). The channel selection will choose the best parent.

Repeaters in a Cisco Wireless Mesh Network look for an AP in two cases: when the gateway is down or it finds a better gateway. AP will scan the entire spectrum in the first case. In the second, a repeater is constantly evaluating the current channel for better gateways. If a repeater finds an AP with a better link quality metric, it will move to the new gateway.



New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

ⓘ If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network scan all channels to collect information from its neighbors. If a valid neighbor is found, the AP switches to the channel.

Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. The number of hops is considered as well, with preference given to gateways with the lower number of hops. An AP finalizes its route based on link quality and number of hops.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

Devices in a Cisco Meraki Wireless Mesh Network communicate using a proprietary routing protocol. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics, including the performance impact of routing traffic through multiple hops. The protocol is also designed to provide ease of deployment while maintaining low channel overhead.

APs in Cisco's Meraki Wireless Mesh Network automatically detect each other and select the best route to the gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: Wireless Mesh Networking

APs in Cisco's Meraki Wireless Mesh Network constantly update its routing tables with the optimal path to the network gateway. Traffic will flow via the best-known path. The APs use a routing metric to determine the best path, where the lower metric (including preference for lower hops or shorter path set at the cloud) equals the better route.

	Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Cisco Meraki Wireless Mesh Network performs the same function of selecting a parent for a mesh access point in the same way by maintaining a list of neighbors and using the number of hops or signal strength as criteria to find the best path back to the external network and achieves the same result by creating a wireless mesh network with optimal performance. Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. The AP node in a Cisco Meraki Wireless Mesh Network associates with a single suitable parent node when the route path of the parent node is the best path of all available parent nodes. It is the best path and shortest when it is the least number of hops such as when a gateway can connect directly to a repeater, for example. This path would be the shortest number of hops and the least latency. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Exhibit 7 - U.S Patent No. 11,368,537 (Cisco Wireless Mesh Networks with Catalyst and Aironet Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Wireless Mesh Network including but not limited to Cisco's Catalyst and Aironet Access Points, with related controllers and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1-19 (the "Asserted Claims") of U.S. Patent No. 11,368,537 (the "'537 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '537 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringed by testing, setting up, and using Cisco Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Wireless Mesh Network in the United States.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the '537 Patent and how to make and use the infringing Cisco Wireless Mesh Network. Specifically, Cisco's Wireless Mesh Access Points, Design and Deployment Guides, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco's website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-19 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Wireless Mesh Network, and the infringing configuration and operation of the controller, access points, and software that comprise the infringing Cisco Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Wireless Mesh Network (and

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g.*, https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

directly infringe literally and/or under the doctrine of equivalents). MeshDynamics put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’537 Patent, or in the alternative was willfully blind to the ’537 Patent and the infringement, following the MeshDynamics letter in 2022.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’537 Patent. Cisco Accused Products come enabled as a mesh access point and automatically connect to the controller to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particular suited for Cisco’s Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’537 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the ’537 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends, in the alternative, that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics’s preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts

presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco discovery, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco's Wireless Mesh Network with Catalyst and Aironet Access Points
1.a. A wireless mesh network comprising: a plurality of Wi-Fi nodes in the wireless mesh network organized in a tree shape;	Cisco's Wireless Mesh Network (a component of the Cisco Unified Wireless Network solution) comprises a plurality of Wi-Fi nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Access Points) organized in a tree shape. Cisco Catalyst 9100 and Aironet Access Points support Wi-Fi capabilities. The Catalyst series is a follow-up from the Cisco Aironet series access points, and the Aironet and Catalyst Access Points contain the same mesh networking capabilities, as shown below in Cisco's documentation and guides that apply to both product lines: Q. Which Cisco access points support Wi-Fi 6? A. Cisco Catalyst 9100 enterprise-class access points support Wi-Fi 6 capabilities, extending the power of intent-based networking with hardware and software innovations. The Cisco Catalyst 9100 APs offer resiliency and security while enabling intelligence at the network edge. Q. Which Cisco access points support Wi-Fi 6E? A. Currently, the Cisco Catalyst 9136, 9166, 9166D1, 9164, 9162, and outdoor 9163E Series Access Points and the Cisco Meraki MR57 are the only APs from Cisco that support Wi-Fi 6E. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html; *see also* https://www.cisco.com/site/us/en/products/networking/wireless/access-points/aironet-1000-series/index.html

Q. What is the portfolio transition from the Cisco Aironet Series to the Cisco Catalyst 9100 platform?

A. The Cisco Catalyst 9100 Access Points consist of the Cisco Catalyst 9105, 9115, 9120, 9130, 9136, 9162, 9164, 9166 and 9166D1, and outdoor 9163E Series. These access points are the follow-ups to the Cisco Aironet 1815, 1850, 2800, and 3800 Series Access Points, respectively. The Cisco Catalyst 9100 Access Points come equipped with Wi-Fi 6/6E capabilities, a better industrial design, and improved RF performance, and deliver reliability, security, and intelligence at scale.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

Cisco Aironet 1560 Series Outdoor Access Points

Product Type	Outdoor and Industrial Wireless
Status	End of Sale EOL Details
Series Release Date	05-JUL-2016
End-of-Sale Date	31-JAN-2023
End-of-Support Date	31-JAN-2028

This product is supported by Cisco, but is no longer being sold. Ready for an upgrade? The Cisco Catalyst 9124AX Series Access Points offer greater performance and functionality.



Source: https://www.cisco.com/c/en/us/support/wireless/aironet-1560-series/series.html

In addition to the Wave 1 and Wave 2 access points, mesh deployments are supported on the Cisco Catalyst 9100 Access Points, including, *e.g.*, Catalyst 9124 Access Points, Catalyst 9130 Access Points and "AX" access points:

Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points?

A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

Families of mesh APs

The two families of mesh APs available are the Cisco Wave 1 APs and the Cisco Wave 2 APs + the AX APs. These two families of mesh APs have different capabilities and features, and they are not fully compatible with each other. The Cisco Wave 1 APs do not support the same feature set as the Cisco Wave 2 + AX APs.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg/mesh-access-points.html

From this release, mesh support is included in the Cisco Catalyst 9130AX Series Access Points. All traditional capabilities of mesh are included in the Cisco Catalyst 9130AX Series APs operating in Cisco IOS XE Dublin 17.12.1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html

Cisco Catalyst 9100 Access Points (like the Cisco Aironet series access points) can be deployed in a Cisco Wireless Mesh Network with a controller and software, such as the Catalyst 9800 Controller and Cisco DNA and Prime Infrastructure:

12. Mesh

This section provides design and deployment guidelines for the deployment of secure enterprise, campus, and metropolitan Wi-Fi networks within Cisco wireless mesh networking, a component of the Cisco Unified Wireless Network solution.

For more detailed information about Cisco wireless mesh networking, including configuration and deployment, refer to the Mesh Deployment Guide for Cisco Catalyst 9800 Series Wireless Controllers: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf.

Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network.

End-to-end security within the mesh network is supported by employing AES encryption between the wireless mesh APs and WPA2/WPA3 clients. This section also outlines RF components to consider when designing an outdoor network.

The features described in this section are for the following outdoor AP products:

- Cisco Catalyst 9124AX Series outdoor mesh access points
- Cisco Aironet 1560 (1562) Series outdoor mesh access points
- Cisco Aironet 1540 (1542) Series outdoor mesh access points
- Cisco Aironet 1570 (1572) Series outdoor mesh access points

Note: 1570 Series outdoor APs are not supported in Cisco DNA Center

- Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series.
- Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series.
- Mesh features in Cisco Catalyst 9800 Series wireless controllers.
- Mesh features in Cisco Prime Infrastructure and Cisco DNA Center.

Mesh network components

The Cisco wireless mesh network has four core components:

- Cisco Catalyst 9800 Series and Cisco IOS XE.
- Cisco Aironet and Catalyst access points.
- Cisco Prime Infrastructure and Cisco DNA Center.
- Mesh software architecture.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco Wireless Mesh Networks organize the wireless access points (such as the Cisco Catalyst and Aironet Access Points) in a tree shape topology as shown in red below:

This figure shows the relationship between RAPs and MAPs in a mesh network.

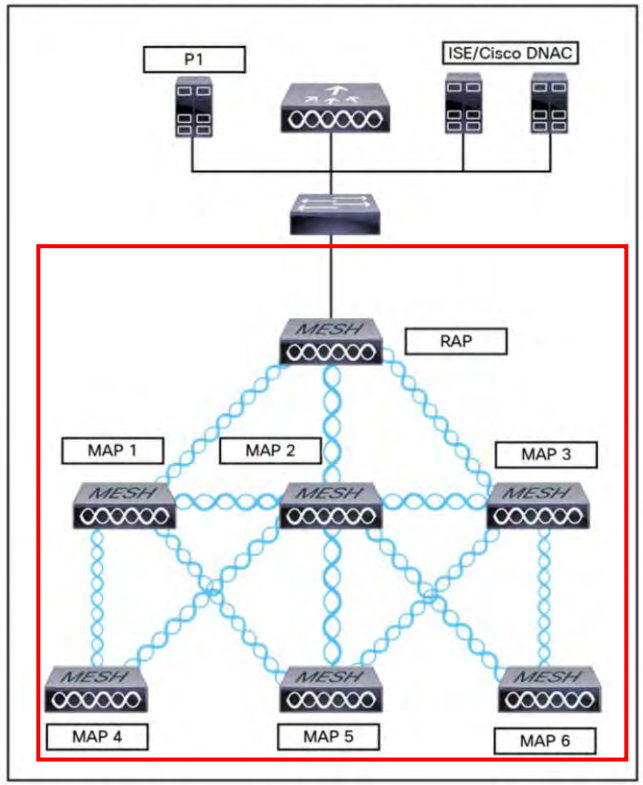
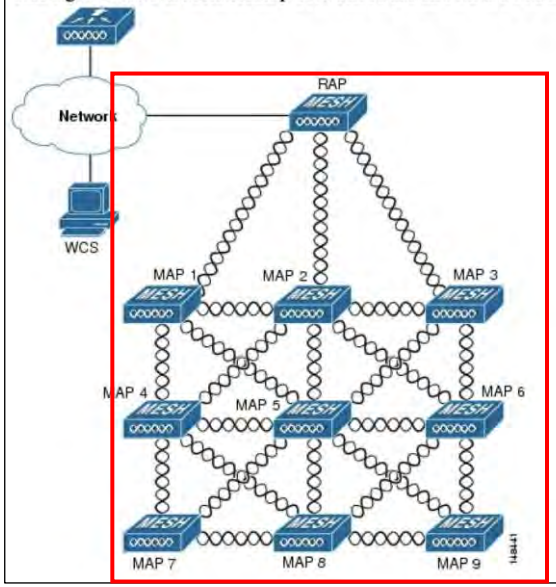


Figure 1: Simple Mesh Network Hierarchy

This figure shows the relationship between RAPs and MAPs in a mesh network.



Caution We recommend that you do not deploy RAPs and MAPs belonging to different domains on the same mesh tree.

Note The overall throughput of backhaul radio decreases by half for each hop of a mesh tree. When the Ethernet-bridged clients are used in MAPs and heavy traffic is passed, it may result in a high throughput consumption, which may cause the downlink MAPs to disassociate from the network due to throughput starvation.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (Wi-Fi communication networking) in the same way (by deploying a plurality of Wi-Fi nodes organized in a "mesh tree") and achieves the same result (a wireless mesh network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses a plurality of Wi-Fi nodes to create a mesh network organized as a "mesh tree." A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.
1.b. an access server in data communication with the nodes,	Cisco's Wireless Mesh Network includes an access server, *e.g.*, Catalyst 9800 series wireless controllers or other Cisco Wireless LAN Controllers (WLC), in data communication with the Catalyst 9100 and Aironet Access Points. The Cisco Wireless Mesh Network includes a Catalyst 9800 Series wireless controller, which manages the connection of the mesh access points/nodes, *e.g.*, Catalyst 9100 and Aironet Access Points, to the network: **12. Mesh** This section provides design and deployment guidelines for the deployment of secure enterprise, campus, and metropolitan Wi-Fi networks within Cisco wireless mesh networking, a component of the Cisco Unified Wireless Network solution. For more detailed information about Cisco wireless mesh networking, including configuration and deployment, refer to the Mesh Deployment Guide for Cisco Catalyst 9800 Series Wireless Controllers: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf. Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network. End-to-end security within the mesh network is supported by employing AES encryption between the wireless mesh APs and WPA2/WPA3 clients. This section also outlines RF components to consider when designing an outdoor network. The features described in this section are for the following outdoor AP products: • Cisco Catalyst 9124AX Series outdoor mesh access points • Cisco Aironet 1560 (1562) Series outdoor mesh access points • Cisco Aironet 1540 (1542) Series outdoor mesh access points • Cisco Aironet 1570 (1572) Series outdoor mesh access points Note: 1570 Series outdoor APs are not supported in Cisco DNA Center • Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series. • Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series. • Mesh features in Cisco Catalyst 9800 Series wireless controllers. • Mesh features in Cisco Prime Infrastructure and Cisco DNA Center. **Mesh network components** The Cisco wireless mesh network has four core components: • Cisco Catalyst 9800 Series and Cisco IOS XE. • Cisco Aironet and Catalyst access points. • Cisco Prime Infrastructure and Cisco DNA Center. • Mesh software architecture.

Cisco wireless controllers

The wireless mesh solution with Cisco IOS XE 17.9.1 is supported on Catalyst 9800-CL virtual and hardware appliances such the 9800-80, 9800-40, and 9800-L.

Wireless mesh network

In a Cisco wireless outdoor mesh network, multiple mesh APs make up a network that provides a secure, scalable outdoor WLAN.

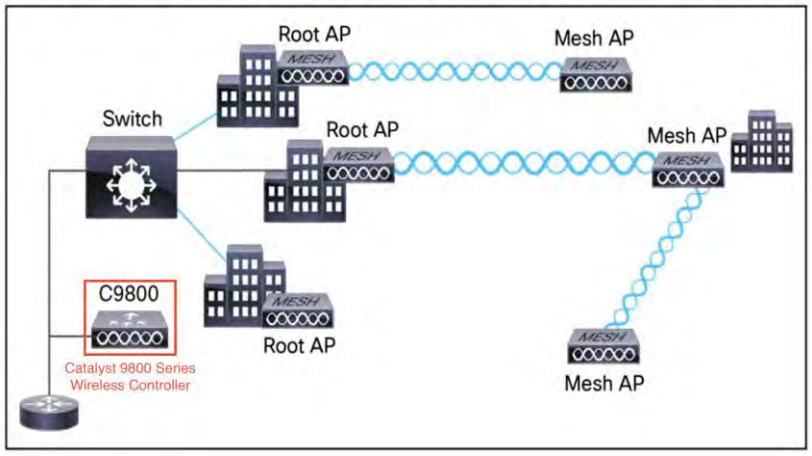


Figure 75.
Wireless Mesh Topology

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The Catalyst 9800 Series wireless controller uses the Control and Provisioning of Wireless Access Points protocol (CAPWAP) to manage access points, *e.g.*, Catalyst 9100 and Aironet Access Points, in the network:

CAPWAP

The Internet Engineering Task Force (IETF) Control and Provisioning of Wireless Access Points (CAPWAP) protocol is the underlying protocol used in the Cisco WLAN architecture. CAPWAP provides the configuration and management of APs and WLANs in addition to encapsulation and forwarding of WLAN client traffic between an AP and a WLAN Controller (WLC).

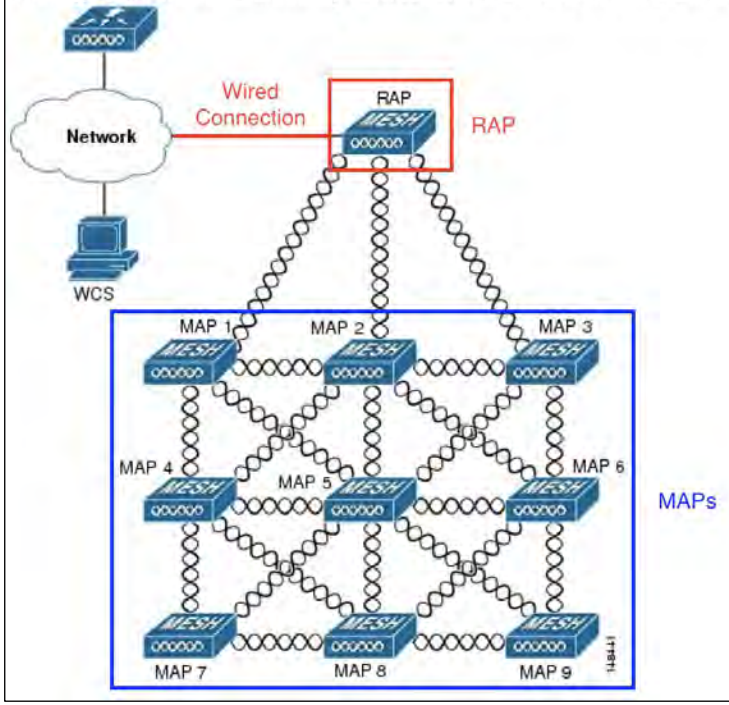
Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Control and Provisioning of Wireless Access Points

Control and provisioning of wireless access points (CAPWAP) is the provisioning and control protocol used by the controller to manage access points (mesh and nonmesh) in the network.

	Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf The Catalyst 9800 Series wireless controller is in data communication with the access points/nodes: While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. A mesh access point establishes an AWPP link with a parent mesh AP that is already connected to the controller before starting CAPWAP discovery. **Network access** Wireless mesh networks can simultaneously carry two different traffic types: • WLAN client traffic • MAP Ethernet port traffic WLAN client traffic terminates on the controller, and the Ethernet traffic terminates on the Ethernet ports of the mesh APs. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (control of access points/node) in the same way (by deploying a controller in data communication with the access points/node) and achieves the same result (control of access points/node). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses a controller in data communication with the mesh network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.c. wherein the nodes are comprised of: one or more root access point (RAP) nodes having a wired	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more root access point (RAP) nodes having a wired connection to an external network and one or more mesh access point (MAP) nodes:

connection to an external network; one or more mesh access point (MAP) nodes;	**Mesh access point roles** Mesh networking employs Cisco Aironet outdoor mesh APs and indoor mesh APs along with the Cisco WLC and Cisco Prime Infrastructure and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network. For Mesh mode, an AP should be configured with Bridge mode. Access points within a mesh network operate in one of the following two ways: • Root AP (RAP) • Mesh AP (MAP) Note: All APs are configured and shipped as mesh APs. To use an APs as a RAP, you must reconfigure the mesh AP as a RAP. In all mesh networks, ensure that there is at least one RAP. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh The RAPs have a wired connection to the controller and external network, and the MAPs have wireless connections to the controller via the RAP or another MAP: While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

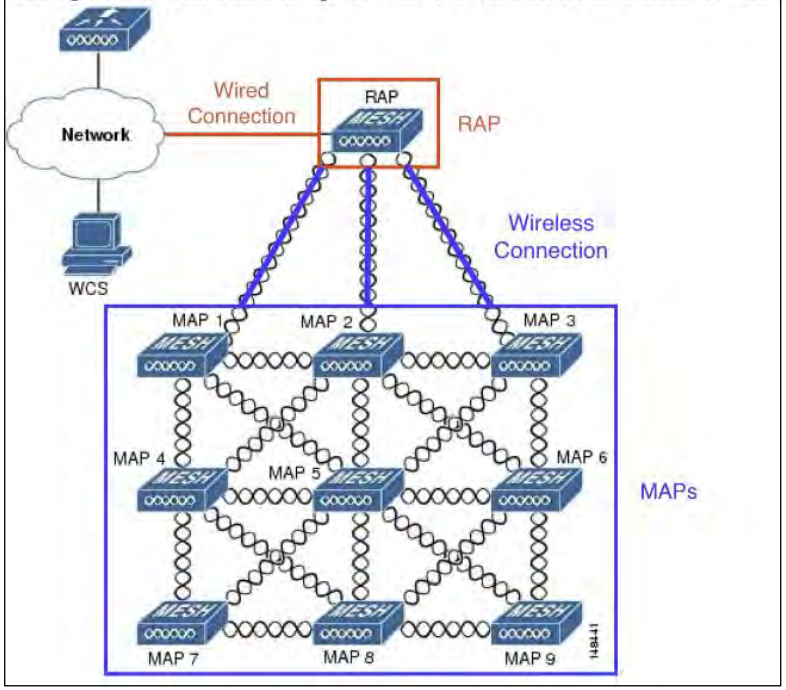
	Figure 1: Simple Mesh Network Hierarchy This figure shows the relationship between RAPs and MAPs in a mesh network.  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (wireless mesh network topology) in the same way (by deploying RAPs having a wired connection to an external network and MAPs) and achieves the same result (wireless mesh network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses a controller in data communication with the RAP and a wired connection to the external network, with one or more MAPs. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.
1.d. wherein a MAP node is in a data communication with the external network through an associated RAP node;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein the MAP nodes are in data communication with the external network through an associated RAP node. The MAPs have wireless connections to the controller and external network via an associated RAP:

While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Figure 1: Simple Mesh Network Hierarchy

This figure shows the relationship between RAPs and MAPs in a mesh network.



s

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (data communication between MAPs and the external network) in the same way (through wireless connection with an associated RAP) and achieves the same result (data communication between MAPs and the external network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses a RAP to connect MAPs to

	the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.e. wherein the MAP node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node;	Cisco's Catalyst 9100 and Aironet Series Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein the MAP node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node. The MAPs have wireless connections to the controller and external network via an associated RAP or a wireless connection to another MAP: While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. A mesh access point establishes an AWPP link with a parent mesh AP that is already connected to the controller before starting CAPWAP discovery. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh The relationships among mesh access points in the Cisco Wireless Mesh Network are a parent, child, or neighbor. A parent access point offers the best route back to the RAP based on ease values, and a parent can be either the RAP or another MAP:

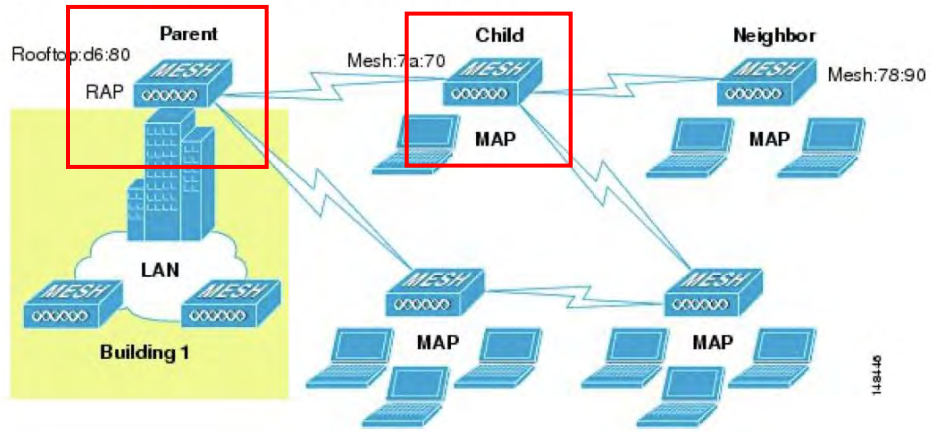
Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
  - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.
- Mesh networks are half duplex, meaning after the first hop (RAP to MAP) overall throughput is decreased by 50% for each additional hop (MAP to MAP). Where Ethernet bridging clients are used in MAPs and heavy traffic is passed, high throughput consumption nay result, which may cause the downlink MAPs to disassociate from the network due to throughput starvation.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Figure 6: Parent, Child, and Neighbor Access Points



Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (data communication between MAPs and the external network) in the same way (through data communication with an associated parent node that is either the associated RAP or another MAP in data communication with the external network through its associated RAP node) and achieves the same result (data communication between MAPs and the external network). Any differences between the accused Cisco Wireless

	Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses a RAP to connect MAPs to the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.f. wherein the MAP node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node, process for the automatic connection comprising:	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein the MAP node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node. Catalyst 9100 and Aironet Access Points in Cisco Wireless Mesh Networks use Adaptive Wireless Path Protocol (AWPP) to determine the best path through the mesh access points (MAPs) to the controller. This includes selection of parents and establishing a parent-child relationship. The MAPs establish a wireless connection with a parent node and have a wireless connection to the controller (via the RAP or another MAP). While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. A mesh access point establishes an AWPP link with a parent mesh AP that is already connected to the controller before starting CAPWAP discovery. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh AWPP connects a MAP to an associated parent through a process for automatic connection designed specifically for wireless mesh networking and to provide ease of deployment, fast convergence and minimal resource consumption. AWPP enables remote access points to dynamically find the best path back to a RAP.

Adaptive Wireless Path Protocol (AWPP)

AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Using the AWPP, the MAPs scan to determine the potential parent nodes that are within a direct wireless communication range:

Criteria for choosing the best parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

1. A list of channels with a neighbor is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
2. The channels with a neighbor are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor.
3. The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state.
4. Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and it is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

	Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (establishing a parent-child relationship between an MAP and an associated parent node) in the same way (by selecting a parent node based on certain criteria) and achieves the same result (selecting a parent node that offers the best route to the controller and external network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses certain criteria to select a parent that is the best route back to the RAP and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.g. selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node; and	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein automatic connection to the associated parent node comprises selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria (*e.g.*, an ease calculation criteria) defined by the access server, wherein the parent node so selected is the associated parent node. The MAP selects an associated parent node from one or more potential patent nodes using the AWPP process shown in Cisco's documentation below, using "ease to determine the best path": **Criteria for choosing the best parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: 1. A list of channels with a neighbor is generated by passive scanning in the scan state, which is a subset of all backhaul channels. 2. The channels with a neighbor are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor. 3. The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state. 4. Parent maintenance and optimization occurs in the maintain state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and it is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The ease selection criteria is calculated using SNR and hop value as explained and illustrated below:

Ease calculation

Ease is calculated using the SNR and hop value of each neighbor and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

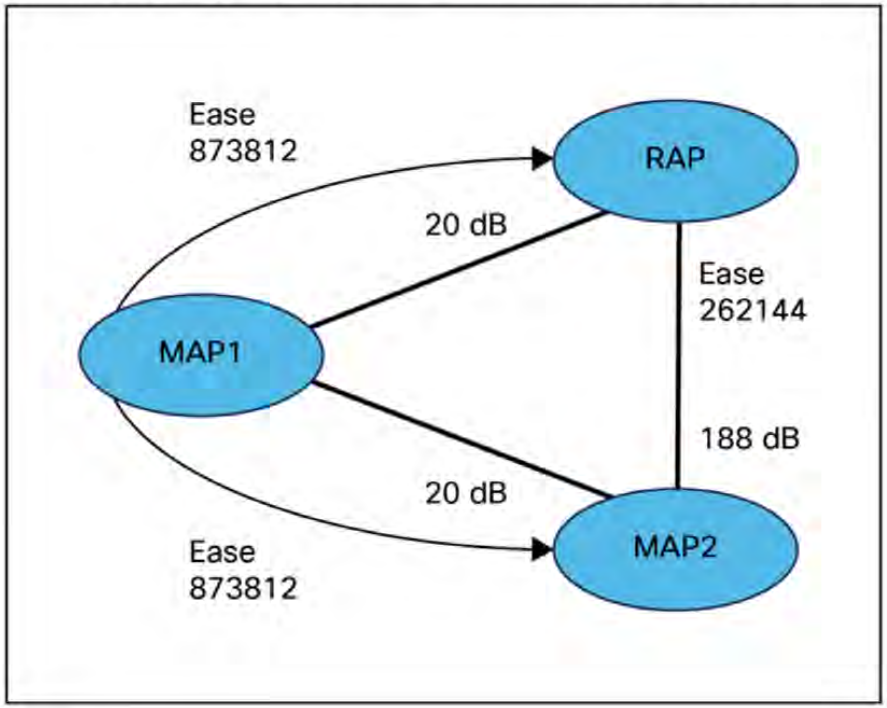


Figure 74.

Parent path selection

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Mesh AP roaming

Mesh APs can roam from one parent mesh AP to a new parent mesh AP. A parent mesh AP and the Catalyst 9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

The Catalyst 9800 Series will support MAP roaming between parent mesh APs within the same controller and parent mesh APs across different controllers. MAP roaming across parent mesh APs connected to Aire-OS and Catalyst 9800 Series controllers in the same mobility group will be supported.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The access server provides the framework for APs to communicate and choose optimal paths, including defining criteria for selecting a parent. For example, Cisco sets or recommends the number of hops or number of MAPs per RAP.

Note With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

- Number of backhaul hops is limited to eight but we recommend three to four hops.

  The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.

- Number of MAPs per RAP.

  There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

To avoid changing route paths based on moment-to-moment fluctuations in a RF link, a parent is given 20% of bonus ease on top of the parent's calculated ease, such that a potential parent must be significantly better than the current parent for a child to make a switch.

SNR smoothing

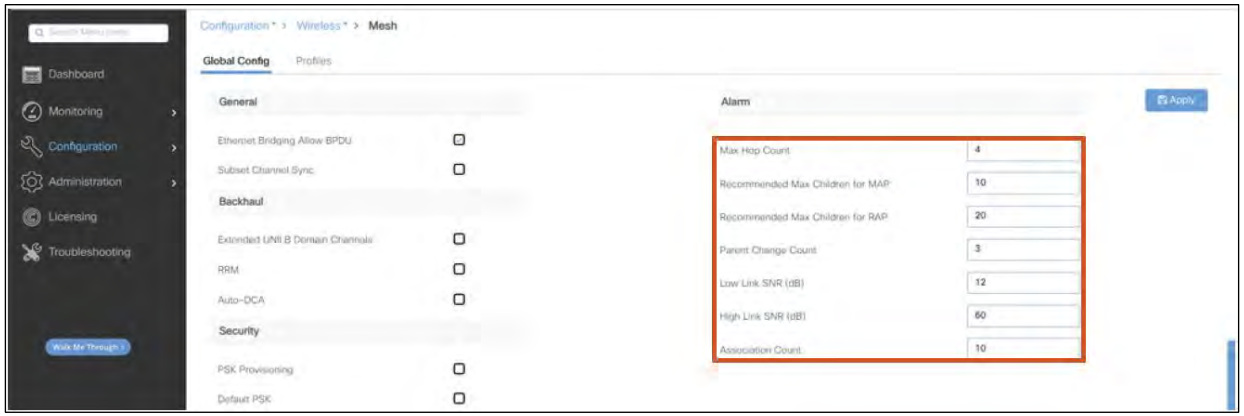
One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch.

Parent switching is transparent to CAPWAP and other higher-layer functions.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Access points ship as mesh access points, and can be configured by the controller to be a RAP. The controller can also be used to configure various mesh parameters, including "Max Hop Count," "Low Link SNR," and "High Link SNR."



Source: https://www.cisco.com/c/en/us/support/docs/wireless/embedded-wireless-controller-on-catalyst-9124ax-access-points/221557-configure-point-to-point-mesh-link-with.html; *see also* https://youtu.be/ESNFzlGTueU?si=I6wUvN7ECTIlKBhx&t=422

The controller can further be used to configure additional mesh parameters, including setting the data rate and convergence time for parent selection at the controller.

Defining Mesh Access Point Roles

By default, the AP1500s are shipped with a radio role set to MAP. Therefore, you must change the radio role on a mesh access point for it to function as RAP.

You can change this configuration on the mesh access point by statically setting them as rooftop access points or mesh access points with the **config ap role** {*rootAP* | *mesh AP* | *default*} command:

To change the radio role can also be changed using the GUI, follow these steps:

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-3/config-guide/b_cg83/troubleshooting_mesh_access_points.pdf

Configuring Global Mesh Parameters

This section provides instructions to configure the mesh access point to establish a connection with the controller including:

- Setting the maximum range between RAP and MAP (not applicable to indoor MAPs).
- Enabling a backhaul to carry client traffic.
- Defining if VLAN tags are forwarded or not.
- Defining the authentication mode (EAP or PSK) and method (local or external) for mesh access points including security settings (local and external authentication).

You can configure the necessary mesh parameters using either the GUI or the CLI. All parameters are applied globally.

Source: https//www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/connecting-mesh-aps.pdf

Configuring Wireless Backhaul Data Rate

Backhaul is used to create only the wireless connection between the access points. The backhaul interface vary between 802.11a/n/ac/ax rates depending upon the access point. The rate selection is important for effective use of the available RF spectrum. The rate can also affect the throughput of client devices, and throughput is an important metric used by industry publications to evaluate vendor devices.

Using the controller, you can configure mesh convergence methods per mesh AP (MAP) or for all mesh APs. This enables you to choose the convergence methods based on deployment without affecting the existing convergence mechanism. The default setting is the existing convergence mechanism.

Mesh Convergence	Parent Loss Detection / Keep Alive Timers	Channel Scan / Seek	DHCP / CAPWAP Information
Standard	21 / 3 seconds	Scan/Seek all 5-GHz channels	Renew/Restart CAPWAP
Fast	7 / 3 seconds	Scan/Seek only preset channels	Maintain DHCP and CAPWAP
Very Fast	4 / 1.5 seconds	Scan/Seek only preset channels	Maintain DHCP and CAPWAP

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-10/config-guide/b_cg810.pdf

The controller also enables MAPs to be assigned to "Bridge Group Names (BGNs)." When "Strict Match BGN" is enabled on a MAP, it will preferentially select parents within the same BGN.

Bridge Group Names

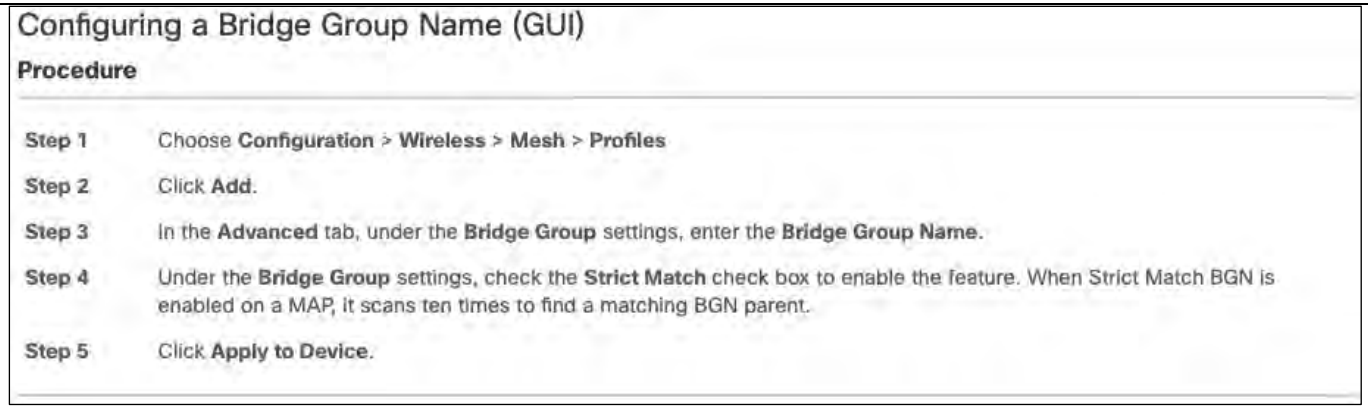
Bridge group names (BGNs) control the association of MAPs to the parent mesh AP. BGNs can logically group radios to avoid two networks on the same channel from communicating with each other. The setting is also useful if you have more than one RAP in your network in the same sector (area). BGN is a string comprising a maximum of 10 characters.

A BGN of *NULL VALUE* is assigned by default during manufacturing. Although not visible to you, it allows a MAP to join the network prior to your assignment of your network-specific BGN.

If you have two RAPs in your network in the same sector (for more capacity), we recommend that you configure the two RAPs with the same BGN, but on different channels.

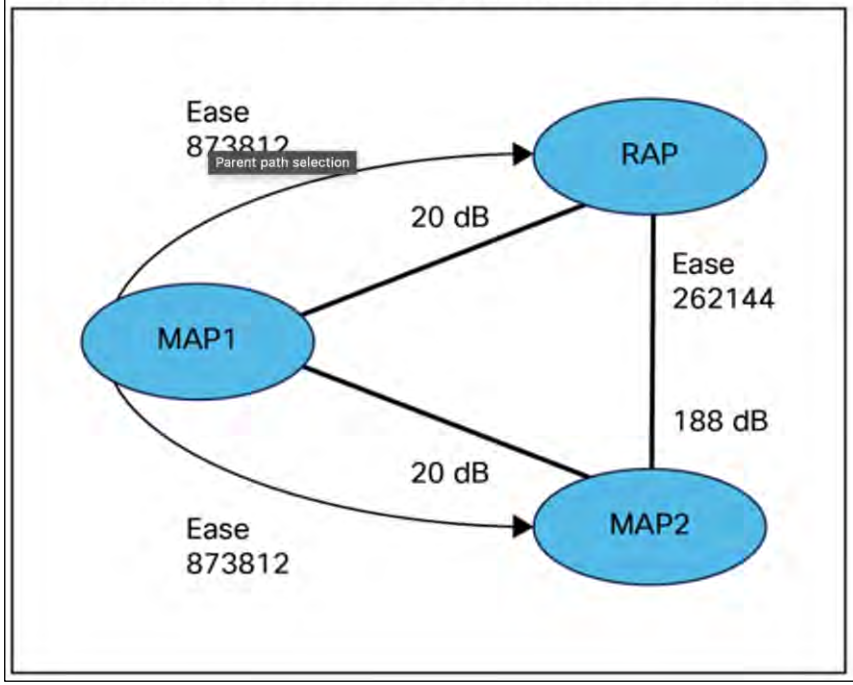
When Strict Match BGN is enabled on a MAP, it will scan ten times to find a matching BGN parent. After ten scans, if the AP does not find the parent with matching BGN, it will connect to the nonmatched BGN and maintain the connection for 15 minutes. After 15 minutes, the AP will again scan ten times, and this cycle continues. The default BGN functionalities remain the same when Strict Match BGN is enabled.

In Cisco Catalyst 9800 Series Wireless Controller, the BGN is configured on the mesh profile. Whenever a MAP joins the controller, the controller pushes the BGN that is configured on the mesh profile to the AP.

	 Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-6/config-guide/ewc_cg_17_6/mesh_access_points.html If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (establishing a parent-child relationship between an MAP and an associated parent node) in the same way (by selecting a parent node based on certain criteria defined by the controller) and achieves the same result (selecting a parent node that offers the best route to the controller and external network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses certain criteria to select a parent that is the best route back to the RAP and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.h. establishing a direct wireless data connection between the MAP node and the associated parent node;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein automatic connection to the associated parent node comprises establishing a direct wireless data connection between the MAP node and the associated parent node. Following AWPP, the MAP establishes a direct wireless data connection with the associated parent node. While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

	Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
1.i. wherein the automatic connection process establishes the MAP node's routing path to the external network;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein automatic connection process to the associated parent node establishes the MAP node's routing path to the external network. The selection of a parent access point (either the RAP itself or another MAP) and the automatic connection process establishes the best route back to the RAP and the external network. While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. **Adaptive Wireless Path Protocol (AWPP)** AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability. A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (establishing the MAP node's routing path to the external network) in the same way (by selecting a parent node based on certain criteria) and achieves the same result (selecting a parent node and associated routing path that offers the best route to the controller and external network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses certain criteria to select a parent that is the best route back to the RAP and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1.j. the one or more parent selection criteria comprising: a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein the one or more parent selection criteria comprises a throughput value representing signal strength between two nodes and a latency value representing number of hops in a routing path Cisco Wireless Mesh Networks use AWPP to provide ease of deployment, fast convergence and minimal resource consumption. The AWPP enables a MAP to dynamically find the best path back to the RAP and external network. It does this by calculating the cost of each path in terms of signal strength and number of hops. **Adaptive Wireless Path Protocol (AWPP)** AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Ease calculation

Ease is calculated using the SNR and hop value of each neighbor and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

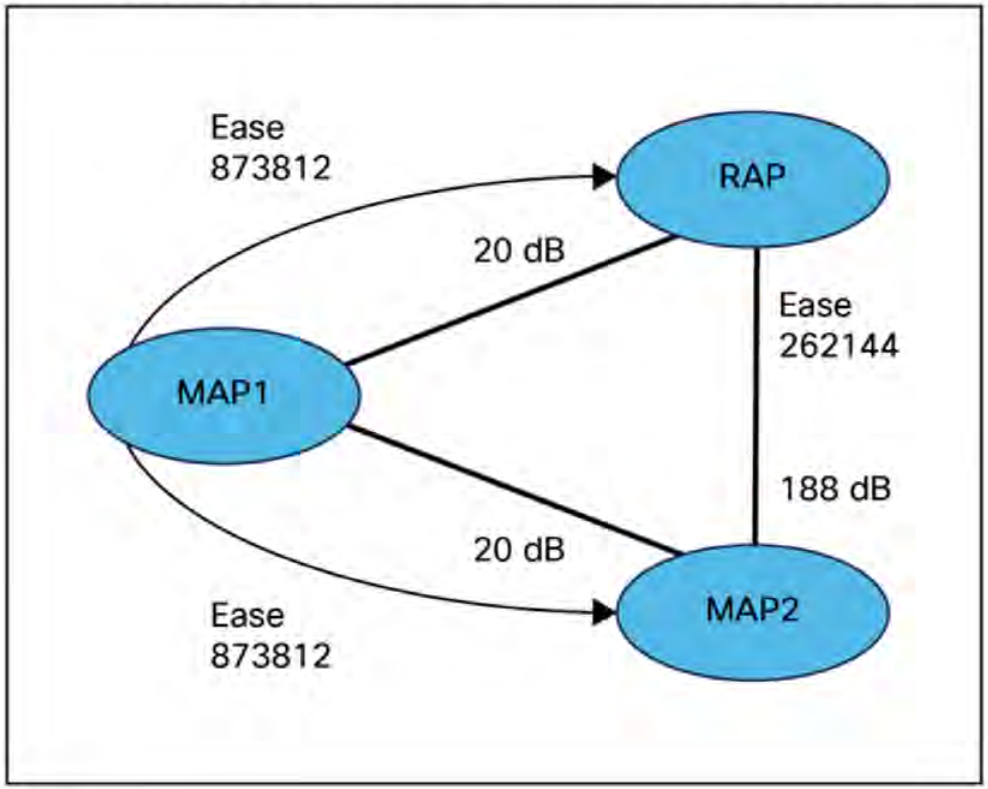


Figure 74.

Parent path selection

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Mesh AP roaming

Mesh APs can roam from one parent mesh AP to a new parent mesh AP. A parent mesh AP and the Catalyst 9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

The Catalyst 9800 Series will support MAP roaming between parent mesh APs within the same controller and parent mesh APs across different controllers. MAP roaming across parent mesh APs connected to Aire-OS and Catalyst 9800 Series controllers in the same mobility group will be supported.

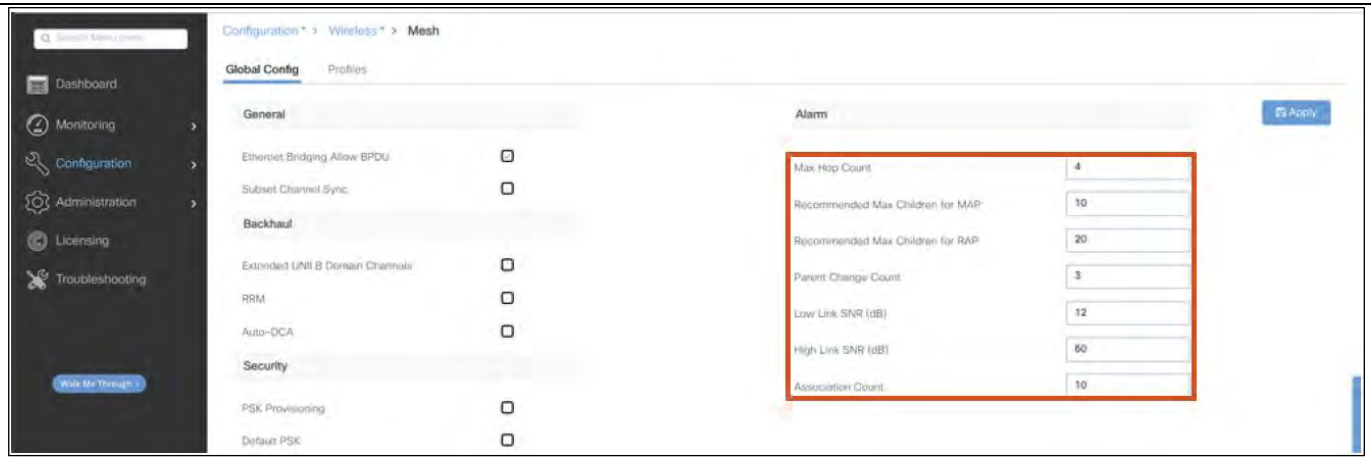
Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
- Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The access server provides the framework for APs to communicate and defines criteria for selecting a parent. For example, the access server can be used to configure max hop counts and SNR parameters.



Source: https://www.cisco.com/c/en/us/support/docs/wireless/embedded-wireless-controller-on-catalyst-9124ax-access-points/221557-configure-point-to-point-mesh-link-with.html; *see also* https://youtu.be/ESNFzlGTueU?si=I6wUvN7ECTIlKBhx&t=422

Note With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

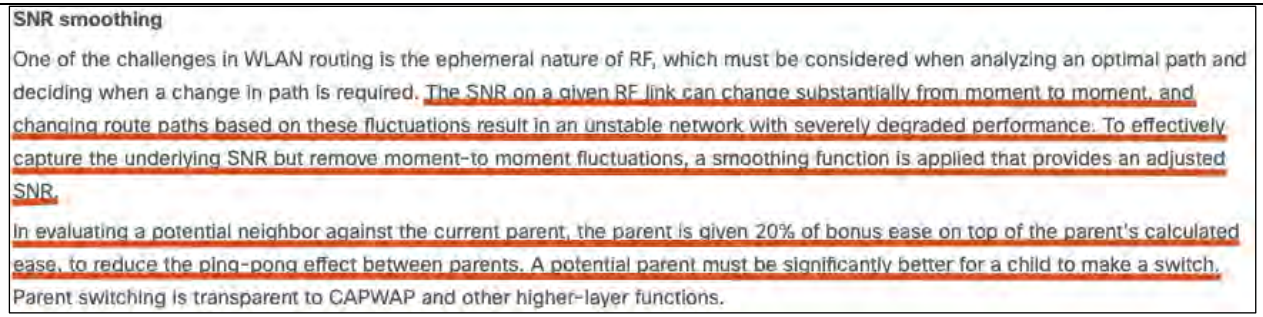
- Number of backhaul hops is limited to eight but we recommend three to four hops.

  The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.

- Number of MAPs per RAP.

  There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

	 **SNR smoothing** One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR. In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent node for a MAP) in the same way (based on parent selection criteria, *e.g.*, signal strength and number of hops) and achieves the same result (selecting a parent node and associated routing path that offers the best route to the controller and external network). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses certain criteria including signal strength and number of hops in a routing path to select a parent that is the best route back to the RAP and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.k. the MAP node comprising: one or more identifier values representing a unique identity of the MAP node in the wireless mesh network;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) comprising one or more identifier values representing a unique identity of the MAP node in the wireless mesh network. Each MAP in the Cisco Wireless Mesh Network comprises a unique Media Access Control (MAC) address. A MAP's MAC address is used to add the MAP to the Cisco Wireless Mesh Network:

Adding Mesh Access Points to the Mesh Network

This section assumes that the controller is already active in the network and is operating in Layer 3 mode.

Note Controller ports that the mesh access points connect to should be untagged.

Before adding a mesh access point to a network, do the following:

Step 1 Add the MAC address of the mesh access point to the controller's MAC filter. See the Adding MAC Addresses of Mesh Access Points to MAC Filter section.

Step 2 Define the role (RAP or MAP) for the mesh access point. See the Defining Mesh Access Point Role section.

Step 3 Verify that Layer 3 is configured on the controller. See the Verifying Layer 3 Configuration section.

Adding the MAC Address of the Mesh Access Point to the Controller Filter List (GUI)

To add a MAC filter entry for the mesh access point on the controller using the controller GUI, follow these steps:

Step 1 Choose **Security > AAA > MAC Filtering**. The MAC Filtering page appears.

Figure 29: MAC Filtering Page



Step 2 Click **New**. The MAC Filters > New page appears.

Step 3 Enter the radio MAC address of the mesh access point.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Step 3. Provision APs to the site

Once the wireless network and RF characteristics are set up, APs can be added to the local or remote site, either by using static AP MAC address assignment or by assigning already joined APs to a specific location.

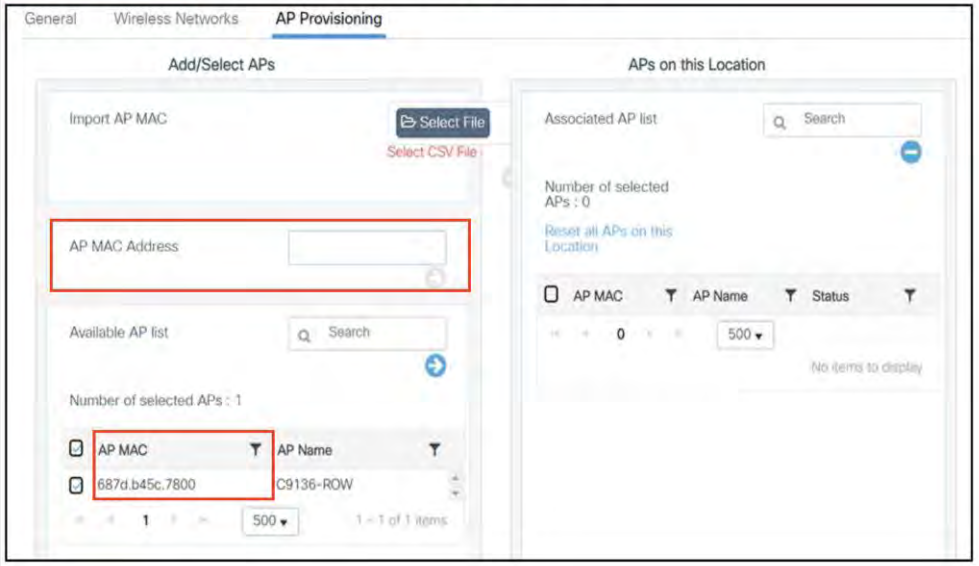


Figure 96.

Basic setup wizard – AP provisioning

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The Cisco Wireless Mesh Network uses a MAP's MAC address to, *inter alia*, manage traffic flow and prevent routing loops:

Traffic flow

The traffic flow within the wireless mesh can be divided into three components:

- Overlay CAPWAP traffic that flows within a standard CAPWAP access point deployment; that is, CAPWAP traffic between the CAPWAP AP and the CAPWAP controller.
- Wireless mesh data frame flow.
- AWPP exchanges.

As the CAPWAP model is well known and AWPP is a proprietary protocol, only the wireless mesh data flow is described. The key to the wireless mesh data flow is the address fields of the 802.11 frames being sent between mesh APs.

An 802.11 data frame can use up to four address fields: receiver, transmitter, destination, and source. The standard frame from a WLAN client to an AP uses only three of these address fields because the transmitter address and the source address are the same. However, in a WLAN bridging network, all four address fields are used because the source of the frame might not be the transmitter of the frame, because the frame might have been generated by a device behind the transmitter.

The figure below shows an example of this type of framing. The source address of the frame is MAP:03:70, the destination address of this frame is the controller (the mesh network is operating in Layer 2 mode), the transmitter address is MAP:D5:60, and the receiver address is RAP:03:40.

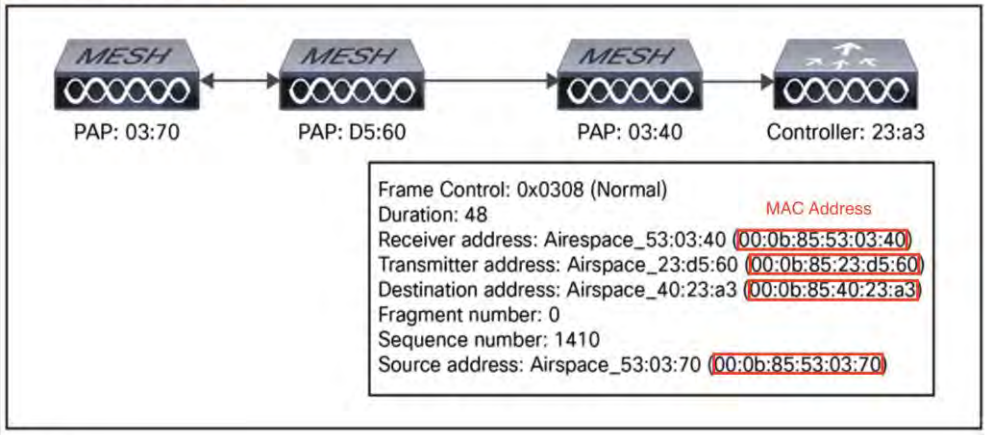


Figure 71.
Wireless mesh frame

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (identifying a MAP) in the same way (by assigning a unique identifier, *e.g.*, MAC and/or IP address to the MAP) and achieves the same result (data flow amongst a plurality of MAPs). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network uses MAC addresses to identify and add access points to the mesh network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1.l. two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes comprising two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network. Catalyst 9100 and Aironet Access Points have two or more radios for wirelessly communicating in a Cisco Wireless Mesh Network with another MAP or RAP, or a client device such as a laptop, including, *e.g.*, the Aironet 4800, Catalyst 9120 series, and Catalyst 9130 series access points. Source: https://www.cisco.com/site/us/en/products/networking/wireless/access-points/catalyst-9100-series/index.html (overview of the Catalyst 9100 Access Points stating two or more radio). **Table 3. AP models and types of hardware managed by FRA** | AP model | FRA radios | Functions | |---|---|---| | Cisco Aironet 2800 Series Access Points | 2.4/5 XOR | 2.4-GHz and 5-GHz or dual 5-GHz operations | | Cisco Aironet 3800 Series Access Points | 2.4/5 XOR | 2.4-GHz and 5-GHz or dual 5-GHz operations | | Cisco Aironet 4800 Series Access Points | 2.4/5 XOR | 2.4-GHz and 5-GHz or dual 5-GHz operations | | Cisco Catalyst 9120 Series Access Points | 2.4/5 XOR | 2.4-GHz and 5-GHz or dual 5-GHz operations | | Cisco Catalyst 9130 Series Access Points | 5-GHz Tri-Radio | 2.4-GHz 4x4 and single 5-GHz 8x8, or 2.4-GHz 4x4 and dual 5-GHz 4x4 | | Cisco Catalyst Wireless 9166 Access Points | 5/6-GHz XOR | 2.4-GHz 4x4 and dual 5-GHz 4x4, or 5-GHz 4x4 and 6-GHz 4x4 |

Mesh network segmentation

Membership in the WLAN mesh network for mesh APs is controlled by the Bridge Group Names (BGNs). Mesh APs can be placed in similar bridge groups to manage membership or provide network segmentation.

Enterprise 11n/ac mesh is added to the Catalyst 9800 controller features to work with the 802.11n/ac APs. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but add higher backhaul and client access speeds. The 802.11ac indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP, and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4-GHz radios in Release 17.1 can be used for local (client) access as well as backhaul. Enterprise 11ac mesh supports point-to-point, point-to-multipoint, and mesh types of architectures.

In the current architecture, while an autonomous AP functions as a WGB, only one radio interface is used for controller connectivity, one Ethernet interface for wired client connectivity, and the other radio interface for wireless client connectivity. dot11radio 1 (5 GHz) can be used to connect to a controller (using the mesh infrastructure) and Ethernet interface for wired clients. dot11radio 0 (2.4 GHz) can be used for wireless client connectivity. Depending on the requirement, dot11radio 1 or dot11radio 0 can be used for client association or controller connectivity.

With two radios, one radio can be used for client access and the other radio can be used for accessing the APs. Having two independent radios performing two independent functions provides you better control and lowers the latency. Also, wireless clients on the second radio for the WGB do not get disassociated by the WGB when an uplink is lost or in a roaming scenario. One radio must be configured as a RAP (radio role), and the second radio must be configured as a WGB (radio role).

Mesh serial backhaul

Serial backhaul enables us to use different channels on different backhauls, which helps in reducing interference and maximizing throughput. The intention is that in units with third and fourth radios, two of the radios are reserved for backhaul.

By keeping radios on exclusive channels or frequency bands, we avoid the need to use the same shared wireless medium between the northbound and southbound traffic in a mesh tree-based network.

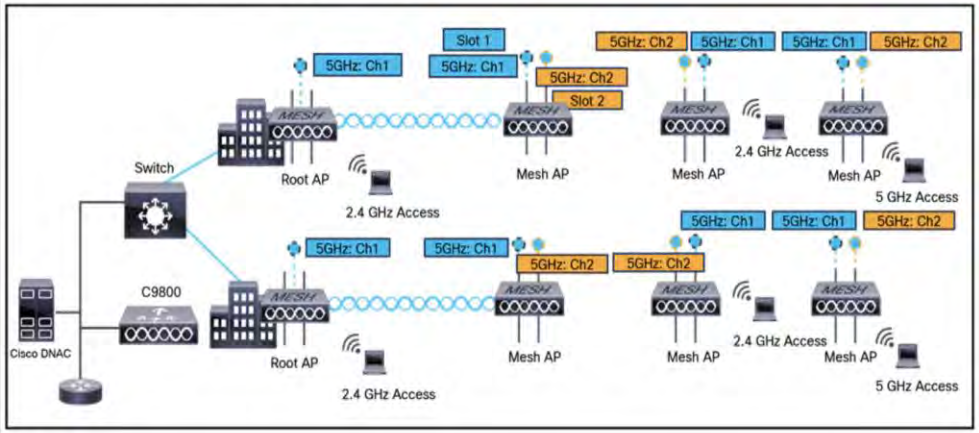


Figure 79.

Mesh serial backhaul design

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network performs the same function (wireless communication) in the same way (with two or more radios) and achieves the same result (wireless communication in two or more frequency bands and/or with two or more AP/client devices). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's access points use two or more radios to communicate in the mesh network and to communicate with client devices. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 2	
2.a. The wireless mesh network of claim 1 further comprising: one or more network health monitoring criteria defined by the access server, wherein the one or more network health monitoring criteria comprises one or more of: whether or not the associated parent node is offline or can be found in the wireless mesh network; whether or not data can be routed to the external network through the MAP node's routing path; whether or not the one or more parent selection criteria is satisfied; and	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein one or more network health monitoring criteria is defined by the access server, *e.g.,* the Catalyst 9800 Wireless Controller, and comprises one or more of: whether or not the associated parent node is offline or can be found in the wireless mesh network; whether or not data can be routed to the external network through the MAP node's routing path; whether or not the one or more parent selection criteria is satisfied. As discussed for claim 1 above, Cisco Wireless Mesh Networks use AWPP to provide ease of deployment, fast convergence and minimal resource consumption. The AWPP enables a MAP to dynamically find the best path back to the RAP and external network by calculating the cost of each path in terms of signal strength and number of hops. The access server provides the framework for APs to communicate and choose optimal paths, including defining criteria for selecting a parent and other criteria. *See supra* discussion for claim elements 1.e-1.i. The AWPP further comprises processes for parent maintenance and optimization. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel:

Criteria for choosing the best parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

1. A list of channels with a neighbor is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
2. The channels with a neighbor are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor.
3. The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state.
4. Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and it is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The AWPP enables parent loss detection. If there are consecutive losses of a parent node's NEIGHBOR_RESPONSE message to a child node's NEIGHBOR_REQUEST, the parent is declared to be lost, offline, or otherwise cannot be found in the network. In case of parent loss, and the routing path involving the lost parent can no longer be used to route data to the external network, the child MAP will try roaming to the next best potential neighbor or try to find a new parent:

Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

Moreover, when a MAP loses the backhaul connection from its current parent, AWPP's mesh convergence feature enables the MAP to reestablish the connection to the Catalyst 9800 controller:

Mesh convergence

Mesh convergence is an important feature in reestablishing the connection to the Catalyst 9800 whenever a MAP loses the backhaul connection from its current parent, which could be a RAP or a successive level MAP.

Standard convergence

In this method, to detect a parent loss, MAP takes 21 seconds and for seek it takes 3 seconds per channel. The parent/neighbor keep-alive time will be 3 seconds.

Fast convergence

In this method, parent loss detection is reduced to 7 seconds and for seek 2 seconds per channel. By using the subset of channels, MAP will scan only in a subset of channels, which reduces the overall seek time. The parent/neighbor keep-alive time will be 3 seconds.

Very fast convergence

In this method, parent loss detection is reduced to 4 seconds and for seek 2 seconds per channel only on the subset of channels. The parent/neighbor keep-alive time will be 1.5 seconds.

Table 17. Mesh convergence time

Convergence method	Parent loss adjTimerMN-seconds	Seek per channel adjTimer1-seconds	Parent, neighbor keep alive adjTimerMP-seconds
Standard	21	3	3
Fast	7	2	3
Very Fast	4	2	1.5

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

The AWPP also includes SNR smoothing, wherein the current parent is given 20 percent of bonus-ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. Thus, when performing parent optimization, a potential parent must be significantly better than the current parent (based on ease) for a child node to make a switch:

SNR smoothing

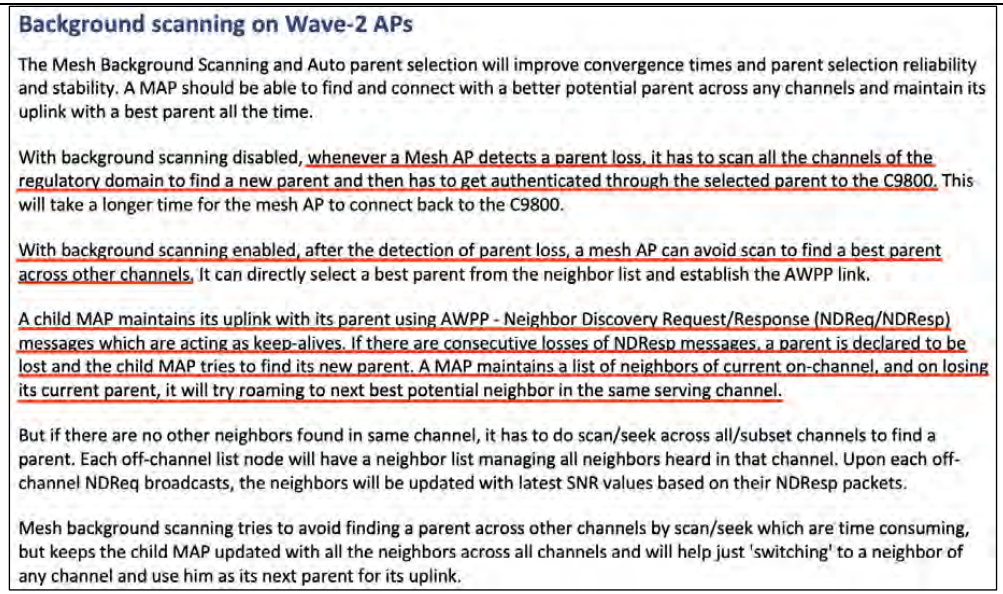
One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (monitor network health, *e.g.*, connection between parent and child node) in the same way (by monitoring whether a parent is offline) and achieves the same result (roaming to a different routing path and/or finding a new parent). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network monitors the connection between the parent and child and will select a new parent as needed. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.b. testing whether the MAP node's routing path fails to satisfy the one or more network health monitoring criteria;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein the one or more MAP node tests whether its routing path fails to satisfy the one or more network health monitoring criteria. The AWPP enables MAPs to test whether its associated parent node is offline or can be found in the wireless mesh network. If parent loss is detected, its routing path involving the lost parent cannot be used to route data to the external network: **Criteria for choosing the best parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: 1. A list of channels with a neighbor is generated by passive scanning in the scan state, which is a subset of all backhaul channels. 2. The channels with a neighbor are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor. 3. The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state. 4. Parent maintenance and optimization occurs in the maintain state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and it is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh



Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

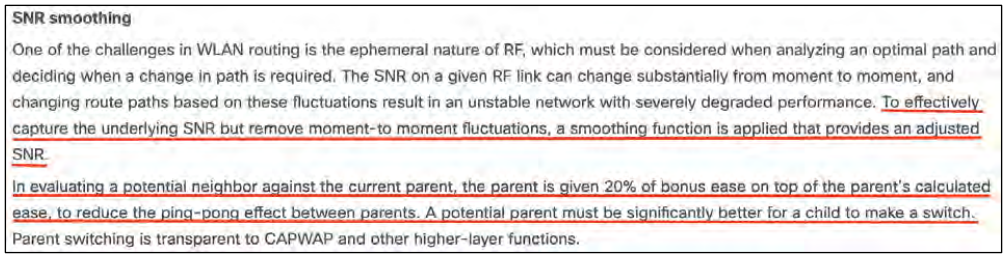
A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

The AWPP also enables MAPs to test whether its associated parent node is significantly worse than a potential parent based on SNR smoothing. If the ease of a potential parent is higher than the bonus-adjusted ease of the current parent, AWPP enables MAPs to switch from the current parent to the potential parent:



SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (test the network health of a MAP node's routing path) in the same way (by monitoring whether a parent is lost, offline, cannot be found, or has significantly worse ease than another potential

	parent) and achieves the same result (roaming to a different routing path and/or finding a new parent). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will monitor and test the MAPs routing path back to the RAP. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.c. when it is determined that the MAP node's routing path does not meet one or more of the one or more network health monitoring criteria, establishing a new routing path from the MAP node to the external network by connecting the MAP node to a new associated parent node through the automatic connection process.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein when it is determined that the MAP node's routing path does not meet one or more of the one or more network health monitoring criteria, establishing a new routing path from the MAP node to the external network by connecting the MAP node to a new associated parent node through the automatic connection process. When a MAP detects a parent loss, and the routing path involving the lost parent can no longer be used to route data to the external network, the MAP will establish a new routing path by roaming to the next best potential neighbor or finding a new parent: **Criteria for choosing the best parent** AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul: 1. A list of channels with a neighbor is generated by passive scanning in the scan state, which is a subset of all backhaul channels. 2. The channels with a neighbor are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor. 3. The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state. 4. Parent maintenance and optimization occurs in the maintain state. This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and it is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information. Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE. Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel. A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

Moreover, when a MAP loses the backhaul connection from its current parent, AWPP's mesh convergence feature enables the MAP to reestablish the connection to the Catalyst 9800 controller:

Mesh convergence

Mesh convergence is an important feature in reestablishing the connection to the Catalyst 9800 whenever a MAP loses the backhaul connection from its current parent, which could be a RAP or a successive level MAP.

Standard convergence

In this method, to detect a parent loss, MAP takes 21 seconds and for seek it takes 3 seconds per channel. The parent/neighbor keep-alive time will be 3 seconds.

Fast convergence

In this method, parent loss detection is reduced to 7 seconds and for seek 2 seconds per channel. By using the subset of channels, MAP will scan only in a subset of channels, which reduces the overall seek time. The parent/neighbor keep-alive time will be 3 seconds.

Very fast convergence

In this method, parent loss detection is reduced to 4 seconds and for seek 2 seconds per channel only on the subset of channels. The parent/neighbor keep-alive time will be 1.5 seconds.

Table 17. Mesh convergence time

Convergence method	Parent loss adjTimerMN-seconds	Seek per channel adjTimer1-seconds	Parent, neighbor keep alive adjTimerMP-seconds
Standard	21	3	3
Fast	7	2	3
Very Fast	4	2	1.5

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

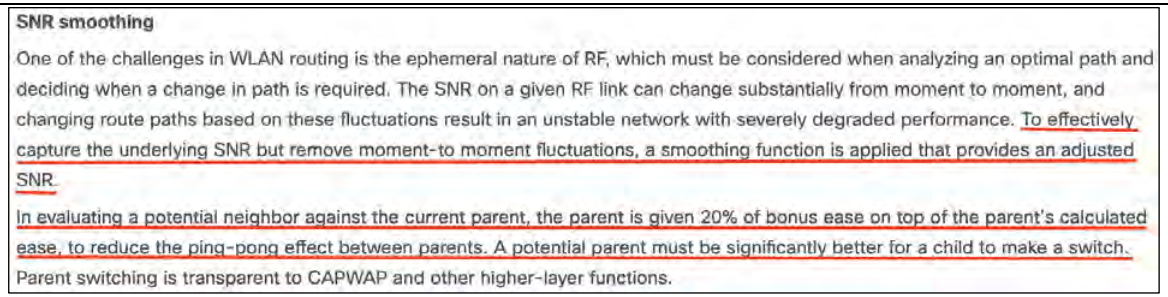
Parent loss Detection

Parent loss detection is based on failure to get response for AWPP neighbor request, which evaluates current parent (AE_UPDATED flag) every 21 seconds in the standard method. Each neighbor is sent a unicast request every 3 seconds along with a request to parent. Failure to get response from parent, initiates a roam if there is any best neighbor available on same channel (OR) a full scan for a new parent. Considering a parent loss on 1st second of an update cycle, it takes another 20 seconds to detect parent loss to initiate roam / full channel scan in standard convergence method. For Fast convergence method, timer adjTimerMN was reduced to 7 seconds from 21 seconds. If there is a failure in getting response for unicast AWPP request, after a retry, change of parent logic will be triggered.

For Very Fast convergence method, timer adjTimerMN was reduced to 4 seconds. Also, Parent, neighbor keep alive timer (adjTimerMP) was reduced to 1.5 seconds to keep updated with parent and adjacent neighbors in same channel. This results in increase of number of request messages (1 request per 3 sec to 1 request per 1.5 sec) for parent and 5 best neighbors.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf

Furthermore, when the bonus-adjusted ease of a MAP's current parent is lower than the ease of another potential parent, AWPP enables MAPs to switch from the current parent to the potential parent:

	 **SNR smoothing** One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR. In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (establishing a new routing path from the MAP node to the external network upon detection of parent loss) in the same way (by connecting the MAP node to a new associated parent node) and achieves the same result (data communication between the MAP node and the external network via the new routing path). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will monitor and test the MAPs routing path back to the RAP and automatically connect the MAP to the best path. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 3	
3. The wireless mesh network of claim 2 wherein the access server further comprises causing an update to the wireless mesh network, the update comprising one or more of: a change to the one or more parent selection criteria,	Cisco's Catalyst 9800 Series Controller in the Cisco Wireless Mesh Network, in addition to the elements of claim 2 as shown above, comprises causing an update to the wireless mesh network, the update comprising one or more of: a change to the one or more parent selection criteria, a change to the one or more network health monitoring criteria, and an upgrade of software or firmware residing on at least some of the nodes. The Catalyst 9800 Series Controller supports upgrade of software or firmware residing on the APs, including rollouts of critical AP bug fixes using an AP Service Pack (APSP):

a change to the one or more network health monitoring criteria, and an upgrade of software or firmware residing on at least some of the nodes.	**Rolling AP update infrastructure** The Cisco Catalyst 9800 Series supports rollouts of critical AP bug fixes using an AP Service Pack (APSP). When APs need to be upgraded to the new image, the 9800 Series supports doing so in a staggered fashion, such that an appropriate number of APs are always up and running in the network and providing RF coverage to clients. This is referred to as a rolling AP upgrade. The AP service pack, which is for AP-specific fixes, will be independent of the SMU timeline and will be available on non-MD releases as well after Release 16.10. Three main highlights to this feature are: • Rolling updates are supported natively on the wireless controller using the UI or CLI. • Rolling updates support automatic candidate selection using the RRM- based AP neighbor information. The device auto-selects the candidate APs to be upgraded in each iteration based on a chosen percentage per iteration (5%, 15%, or 25%, with the default being 15%) and RRM AP neighbor information. • Clients from candidate APs are actively steered away using 802.11v packets with the "dissociation imminent" field set to help ensure seamless network connectivity as APs are being upgraded. If clients do not honor this setting, they will be de-authenticated before AP reload. **Per-site AP service pack rollout** At the time of AP service pack activation, the user selects the sites where the AP service pack should be rolled out. All APs on these sites will be updated with the designated service pack, including any new APs that join the site after the filter is applied. This allows the user to control the propagation of a service pack in the network. It should be noted that this enhancement allows for activating service packs on sites incrementally but requires that all sites be brought to the same service pack level before a new service pack can be rolled out to a subset of sites. **Per-AP-model service pack rollout** An AP service pack can also be built with a subset of AP images. These enable a pre-download only to the affected AP models. Similarly, these service packs are activated only on the AP models affected, in conjunction with any site-based filters, as mentioned earlier. Again, it should be noted that if, for example, three model images were included in an AP service pack, all future AP service packs in that release for any of these three AP images will contain all three of them. This helps subsequent service packs to supersede older ones. These two capabilities work in conjunction with each other, meaning that you can select specific sites in a campus and then apply the fix to specific AP models within those sites, as designated by the service pack. This enables controlled propagation of the fix with minimal or no service disruption because the fix is predownloaded and rolled out only to affected AP models. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh Moreover, AWPP's mesh convergence feature allows changing between 3 convergence methods—standard convergence, fast convergence, and very fast convergence:

Mesh convergence

Mesh convergence is an important feature in reestablishing the connection to the Catalyst 9800 whenever a MAP loses the backhaul connection from its current parent, which could be a RAP or a successive level MAP.

Standard convergence

In this method, to detect a parent loss, MAP takes 21 seconds and for seek it takes 3 seconds per channel. The parent/neighbor keep-alive time will be 3 seconds.

Fast convergence

In this method, parent loss detection is reduced to 7 seconds and for seek 2 seconds per channel. By using the subset of channels MAP will scan only in a subset of channels, which reduces the overall seek time. The parent/neighbor keep-alive time will be 3 seconds.

Very fast convergence

In this method, parent loss detection is reduced to 4 seconds and for seek 2 seconds per channel only on the subset of channels. The parent/neighbor keep-alive time will be 1.5 seconds.

Table 17. Mesh convergence time

Convergence method	Parent loss adjTimerMN-seconds	Seek per channel adjTimer1-seconds	Parent, neighbor keep alive adjTimerMP-seconds
Standard	21	3	3
Fast	7	2	3
Very Fast	4	2	1.5

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Parent loss Detection

Parent loss detection is based on failure to get response for AWPP neighbor request, which evaluates current parent (AE_UPDATED flag) every 21 seconds in the standard method. Each neighbor is sent a unicast request every 3 seconds along with a request to parent. Failure to get response from parent, initiates a roam if there is any best neighbor available on same channel (OR) a full scan for a new parent. Considering a parent loss on 1st second of an update cycle, it takes another 20 seconds to detect parent loss to initiate roam / full channel scan in standard convergence method. For Fast convergence method, timer adjTimerMN was reduced to 7 seconds from 21 seconds. If there is a failure in getting response for unicast AWPP request, after a retry, change of parent logic will be triggered.

For Very Fast convergence method, timer adjTimerMN was reduced to 4 seconds. Also, Parent, neighbor keep alive timer (adjTimerMP) was reduced to 1.5 seconds to keep updated with parent and adjacent neighbors in same channel. This results in increase of number of request messages (1 request per 3 sec to 1 request per 1.5 sec) for parent and 5 best neighbors.

	Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Catalyst 9800 Series Controller performs the same function (updating the wireless mesh network) in the same way (by rolling on the update using the Catalyst 9800 Series Controller) and achieves the same result (wireless mesh network with updated settings, software, or firmware). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will provide updates, such as images, to the access points. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 4	
4. The wireless mesh network of claim 2 wherein the at least one of the one or more network health monitoring criteria overlap with at least one of the one or more parent selection criteria.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 2 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the at least one of the one or more network health monitoring criteria overlap with at least one of the one or more parent selection criteria. As discussed for claim 2 above, a MAP tests whether a potential parent's ease is higher than the current parent's bonus-adjusted ease to determine whether to switch from the current parent to the potential parent. *See supra* discussion for claim 2. The current parent's bonus-adjusted ease and the potential parent's ease overlaps with the ease calculation discussed for claim 1. *See supra* discussion for claim 1. **SNR smoothing** One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR. In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

	Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh (annotation added). MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 5	
5. The wireless mesh network of claim 1 wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their latency values.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their latency values. As discussed for claim 1 above, Cisco Wireless Mesh Networks use AWPP to provide ease of deployment, fast convergence and minimal resource consumption. The AWPP enables a MAP to dynamically find the best path back to the RAP and external network. It does this by comparing the ease of each path (based on signal strength and number of hops), where higher number of hops (*i.e.*, latency value) reduces the ease of a path. **Adaptive Wireless Path Protocol (AWPP)** AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Ease calculation

Ease is calculated using the SNR and hop value of each neighbor and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

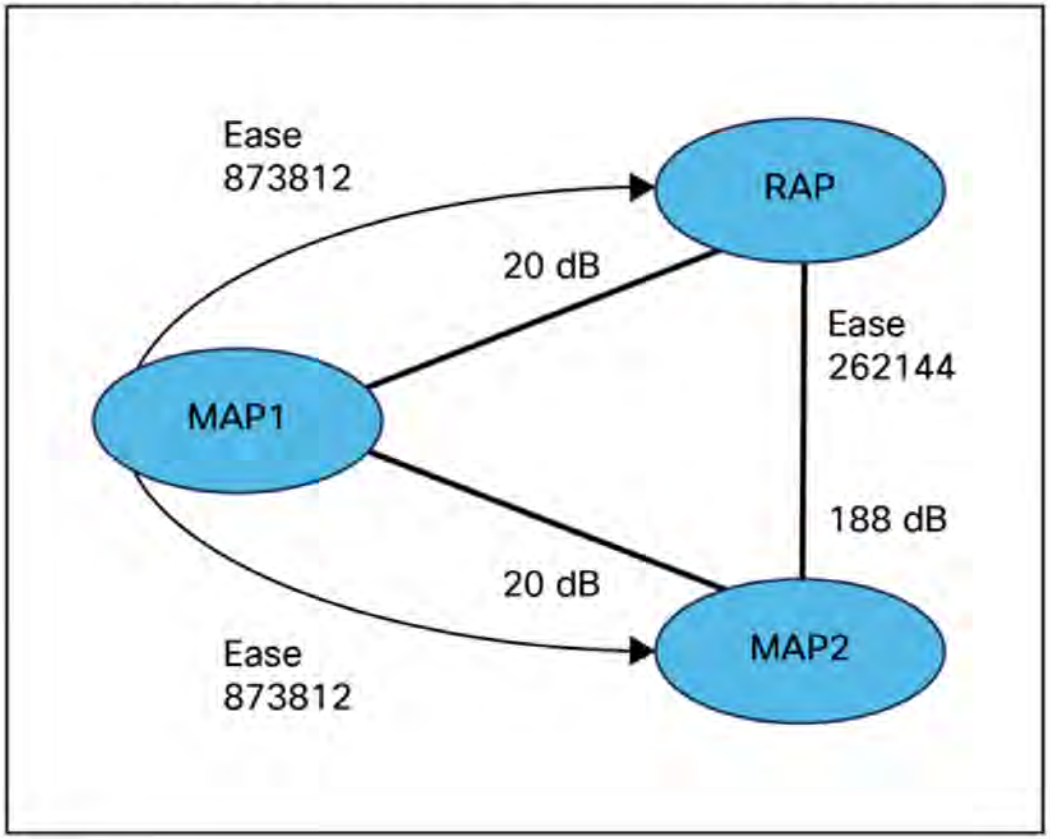


Figure 74.

Parent path selection

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Mesh AP roaming

Mesh APs can roam from one parent mesh AP to a new parent mesh AP. A parent mesh AP and the Catalyst 9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

The Catalyst 9800 Series will support MAP roaming between parent mesh APs within the same controller and parent mesh APs across different controllers. MAP roaming across parent mesh APs connected to Aire-OS and Catalyst 9800 Series controllers in the same mobility group will be supported.

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
- Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent) in the same way (using criteria such as latency or delay) and achieves the same result (best path back to the RAP). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will select the parent node based at least in part on a comparison of their latency values. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 6	
6. The wireless mesh network of claim 5 wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a lower latency value than the another potential parent node.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 5 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a lower latency value than the another potential parent node. As discussed for claims 1 and 5 above, Cisco Wireless Mesh Networks use AWPP to provide ease of deployment, fast convergence and minimal resource consumption. The AWPP enables a MAP to dynamically find the best path back to the RAP and external network. It does this by comparing the ease of each path (based on signal strength and number of hops), where higher number of hops (*i.e.*, latency value) reduces the ease of a path. Thus, an MAP would tend to select a parent node that provides a routing path with less number of hops (*i.e.*, lower latency value): **Adaptive Wireless Path Protocol (AWPP)** AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Ease calculation

Ease is calculated using the SNR and hop value of each neighbor and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

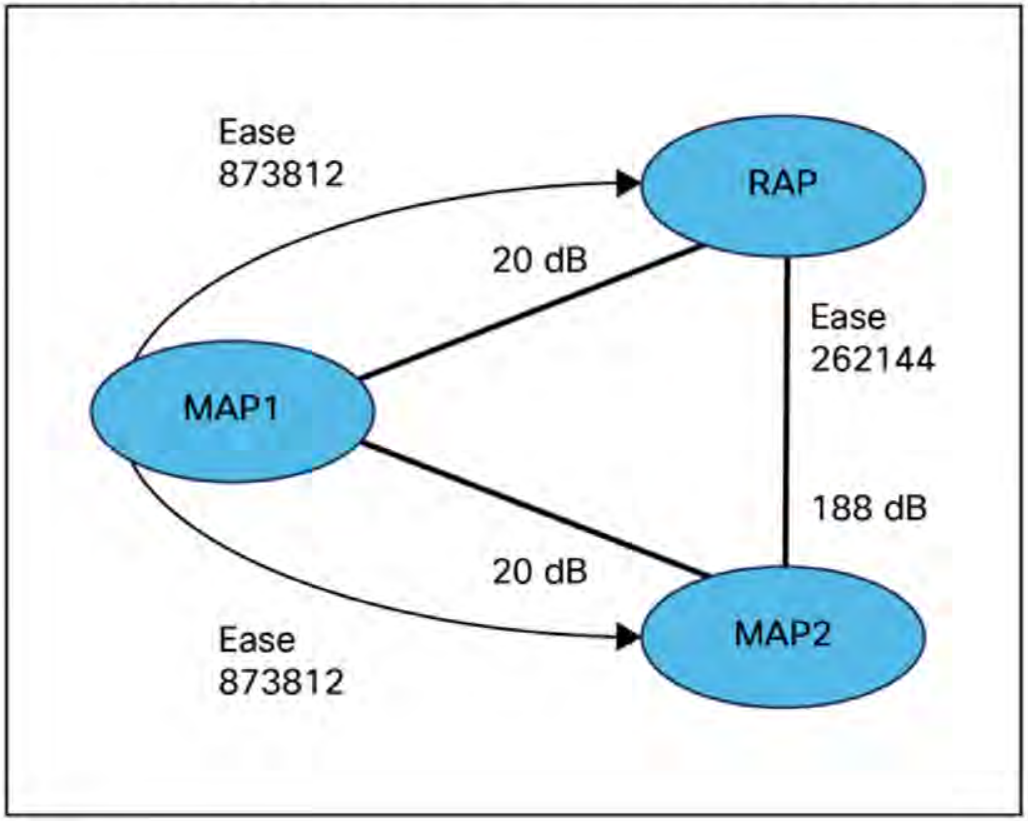


Figure 74.

Parent path selection

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Mesh AP roaming

Mesh APs can roam from one parent mesh AP to a new parent mesh AP. A parent mesh AP and the Catalyst 9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

The Catalyst 9800 Series will support MAP roaming between parent mesh APs within the same controller and parent mesh APs across different controllers. MAP roaming across parent mesh APs connected to Aire-OS and Catalyst 9800 Series controllers in the same mobility group will be supported.

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
- Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent) in the same way (using criteria such as latency or signal strength or number of hops) and achieves the same result (best path back to the RAP). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will select the parent node based at least in part on a comparison of their latency values. A person

	skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 7	
7. The wireless mesh network of claim 1 wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their throughput values.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their throughput values. As discussed for claim 1 above, Cisco Wireless Mesh Networks use AWPP to provide ease of deployment, fast convergence and minimal resource consumption. The AWPP enables a MAP to dynamically find the best path back to the RAP and external network. It does this by comparing the ease of each path (based on SNR and number of hops), where higher SNR (*i.e.*, throughput value) increases the ease of a path. **Adaptive Wireless Path Protocol (AWPP)** AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Ease calculation

Ease is calculated using the SNR and hop value of each neighbor and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

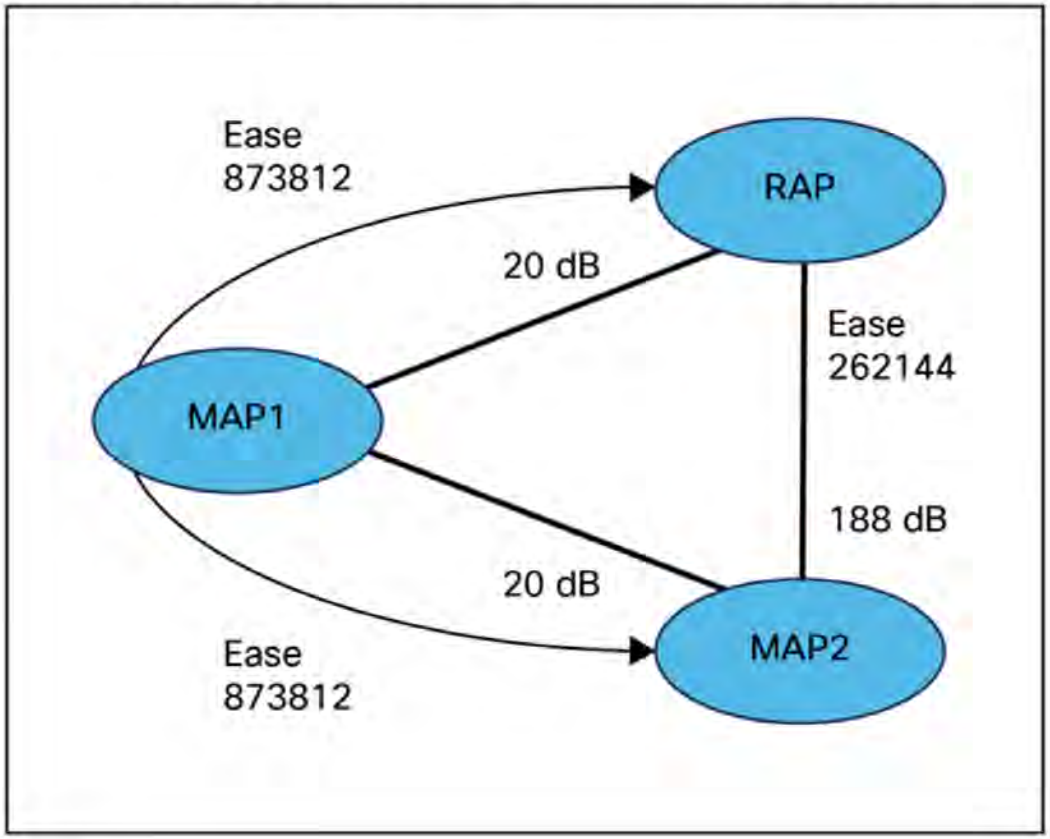


Figure 74.

Parent path selection

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Mesh AP roaming

Mesh APs can roam from one parent mesh AP to a new parent mesh AP. A parent mesh AP and the Catalyst 9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

The Catalyst 9800 Series will support MAP roaming between parent mesh APs within the same controller and parent mesh APs across different controllers. MAP roaming across parent mesh APs connected to Aire-OS and Catalyst 9800 Series controllers in the same mobility group will be supported.

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
- Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent) in the same way (using criteria such as throughput or signal strength or number of hops) and achieves the same result (best path back to the RAP). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will select the parent node based at least in part on a comparison of their throughput values. A person

	skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 8	
8. The wireless mesh network of claim 7 wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a higher throughput value than the another potential parent node.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 7 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a higher throughput value than the another potential parent node. As discussed for claims 1 and 7 above, Cisco Wireless Mesh Networks use AWPP to provide ease of deployment, fast convergence and minimal resource consumption. The AWPP enables a MAP to dynamically find the best path back to the RAP and external network. It does this by comparing the ease of each path (based on SNR and number of hops), where higher SNR (*i.e.*, throughput value) increases the ease of a path. Thus, an MAP would tend to select a parent node that has a higher SNR (*i.e.*, higher throughput value): **Adaptive Wireless Path Protocol (AWPP)** AWPP is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption. AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP. AWPP enables a remote AP to dynamically find the best path back to a RAP for each MAP that is part of the RAP BGN. Unlike traditional routing protocols, AWPP takes RF details into account. To optimize the route, a MAP actively solicits neighbor MAPs. During the solicitation, the MAP learns all the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops. AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Ease calculation

Ease is calculated using the SNR and hop value of each neighbor and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows a parent path selection process. In the figure, MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater than the ease value (262144) of the direct path from MAP2 to RAP.

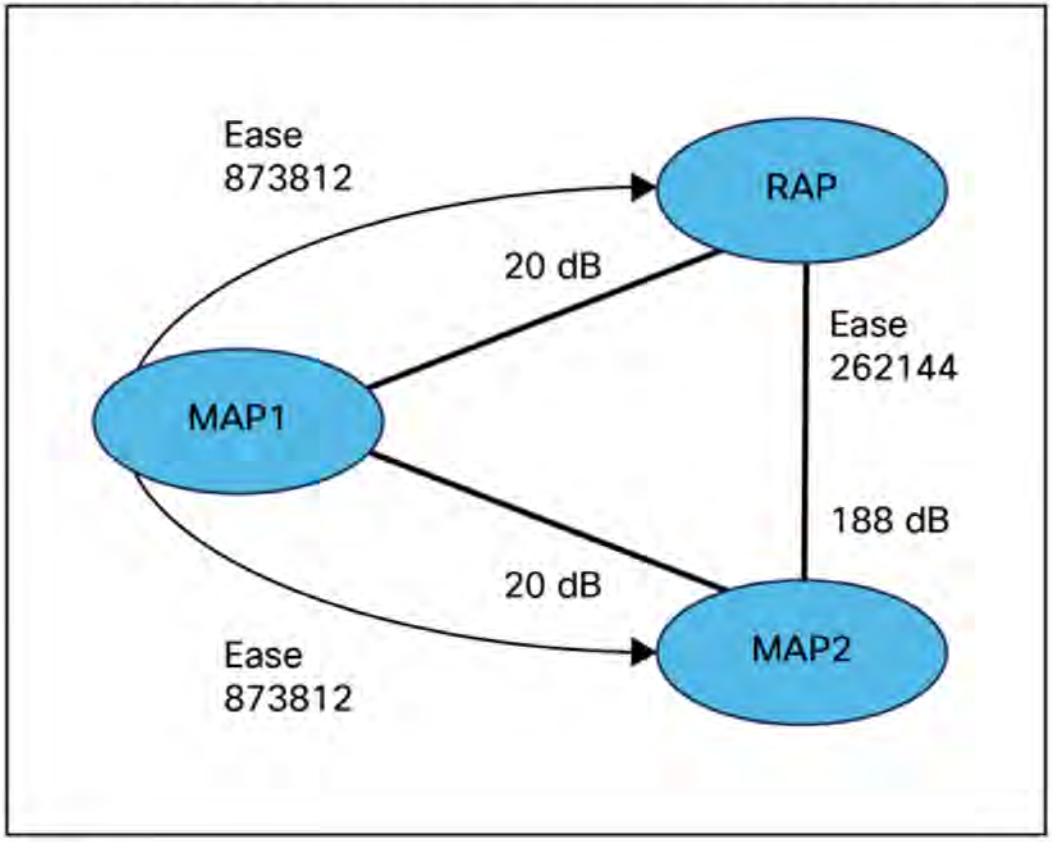


Figure 74.

Parent path selection

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

SNR smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Mesh AP roaming

Mesh APs can roam from one parent mesh AP to a new parent mesh AP. A parent mesh AP and the Catalyst 9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

The Catalyst 9800 Series will support MAP roaming between parent mesh APs within the same controller and parent mesh APs across different controllers. MAP roaming across parent mesh APs connected to Aire-OS and Catalyst 9800 Series controllers in the same mobility group will be supported.

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
- Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent) in the same way (using criteria such as throughput or signal strength or number of hops) and achieves the same result (best path back to the RAP). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless

	Mesh Network will select the parent node based at least in part on a comparison of their throughput values. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 9	
9. The wireless mesh network of claim 1 wherein the selection of the parent node further comprises not selecting a potential parent node if the potential parent node's wireless signal strength with the MAP node does not meet a threshold throughput value defined by the access server.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises not selecting a potential parent node if the potential parent node's wireless signal strength with the MAP node does not meet a threshold throughput value defined by the access server. The AWPP includes SNR smoothing, wherein the current parent is given a pre-defined 20 percent of bonus-ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. Thus, when performing parent optimization, a potential parent will not be selected if its ease does not exceed the current parent's bonus-adjusted-ease (*i.e.*, threshold throughput value): **SNR smoothing** One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations result in an unstable network with severely degraded performance. To effectively capture the underlying SNR but remove moment-to moment fluctuations, a smoothing function is applied that provides an adjusted SNR. In evaluating a potential neighbor against the current parent, the parent is given 20% of bonus ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent) in the same way (using criteria such as throughput or signal strength) and achieves the same result (best path back to the RAP). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network will select the parent node based at least in part on its throughput value and whether it exceeds the 20% bonus ease. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 10	
10. The wireless mesh network of claim 1 wherein the MAP node further comprises a data store having data representing the MAP node's routing path to the external network.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) node which comprises a data store having data representing the MAP node's routing path to the external network. Cisco's Catalyst 9100 and Aironet Access Points comprise system memory (*i.e.*, data store). Source: https://www.cisco.com/site/us/en/products/networking/wireless/access-points/catalyst-9100-series/index.html; *see also* https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-4800-access-point/nb-09-air-4800-acces-ds-cte.html The Catalyst 9100 and Aironet Access Points maintains a list of neighbors of current on-channel, which enables roaming to the next best potential neighbor in the event that the current parent is lost. A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh The Catalyst 9100 and Aironet Access Points also maintain routing information which contains the MAC address of each hop to the RAP. **Loop prevention** To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Wireless Mesh Network performs the same function (selecting a parent) in the same way (maintaining a list of neighbors and routing information) and achieves the same result (best path back to the RAP). Any differences between the accused Cisco Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Wireless Mesh Network, using AWPP, maintains routing information and a list of neighbors to use to determine the best potential neighbor or parent back to the RAP. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 11	
11. The wireless mesh network of claim 1 wherein the MAP node further comprises: an ethernet port through which a direct wired connection can be established with another device.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) node which comprises an ethernet port through which a direct wired connection can be established with another device: Source: https://www.cisco.com/site/us/en/products/networking/wireless/access-points/catalyst-9100-series/index.html (product overview showing Ethernet connection). **Network access** Wireless mesh networks can simultaneously carry two different traffic types: • WLAN client traffic • MAP Ethernet port traffic WLAN client traffic terminates on the controller, and the Ethernet traffic terminates on the Ethernet ports of the mesh APs. Access to the WLAN mesh for mesh APs is managed by the following authentication methods: • MAC authentication: Mesh APs are added to a database that can be referenced to ensure that they are provided access to a given controller and mesh network. • External RADIUS authentication: Mesh APs can be externally authorized using a RADIUS server such as Cisco ISE that supports the client authentication type of EAP-FAST with certificates and WPA2/PSK on the Catalyst 9800 Series. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Claim 12	
12. The wireless mesh network of claim 1 wherein the one or more identifier values of the MAP comprises a MAC address or an IP address.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) wherein the one or more identifier values of the MAP comprises a MAC address or an IP address. As discussed and shown above for claim element 1.j, each MAP in the Cisco Wireless Mesh Network comprises a unique Media Access Control (MAC) address and an IP address. *See supra* discussion for claim element 1.j.
Claim 13	
13. The wireless mesh network of claim 1 wherein a first radio of the MAP is adapted for communicating with the client device and wherein a second radio is adapted for communicating with another node.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes comprising two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network, wherein a first radio of the MAP is adapted for communicating with the client device and wherein a second radio is adapted for communicating with another node. As discussed for claim element 1.k above, Catalyst 9100 and Aironet Access Points have two or more radios for wirelessly communicating in a Cisco Wireless Mesh Network with another MAP or RAP, or a client device such as a laptop. *See supra* discussion for claim element 1.k. The first radio of the MAP is adapted for communicating with the client device and the second radio is adapted for communicating with another node: **Mesh network segmentation** Membership in the WLAN mesh network for mesh APs is controlled by the Bridge Group Names (BGNs). Mesh APs can be placed in similar bridge groups to manage membership or provide network segmentation. Enterprise 11n/ac mesh is added to the Catalyst 9800 controller features to work with the 802.11n/ac APs. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but add higher backhaul and client access speeds. The 802.11ac indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP, and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4-GHz radios in Release 17.1 can be used for local (client) access as well as backhaul. Enterprise 11ac mesh supports point-to-point, point-to-multipoint, and mesh types of architectures.

	In the current architecture, while an autonomous AP functions as a WGB, only one radio interface is used for controller connectivity, one Ethernet interface for wired client connectivity, and the other radio interface for wireless client connectivity. dot11radio 1 (5 GHz) can be used to connect to a controller (using the mesh infrastructure) and Ethernet interface for wired clients. dot11radio 0 (2.4 GHz) can be used for wireless client connectivity. Depending on the requirement, dot11radio 1 or dot11radio 0 can be used for client association or controller connectivity. With two radios, one radio can be used for client access and the other radio can be used for accessing the APs. Having two independent radios performing two independent functions provides you better control and lowers the latency. Also, wireless clients on the second radio for the WGB do not get disassociated by the WGB when an uplink is lost or in a roaming scenario. One radio must be configured as a RAP (radio role), and the second radio must be configured as a WGB (radio role). Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 14	
14. The wireless mesh network of claim 13, wherein the second radio is configurable to be used as a dual-purpose radio for communicating with both the client and another node.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 13 as shown above, are comprised of one or more mesh access point (MAP) nodes comprising two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network, wherein a first radio of the MAP is adapted for communicating with the client device and wherein a second radio is configurable to be used as a dual-purpose radio for communicating with both the client and another node. As discussed for claim elements 1.k and 13 above, Catalyst 9100 and Aironet Access Points comprise a first radio adapted for communicating with the client device and a second radio adapted for communicating with another node. *See supra* discussion for claims 1.k and 13. The second radio is further configurable to be used as a dual-purpose radio for communicating with both the client and another node:

Mesh network segmentation

Membership in the WLAN mesh network for mesh APs is controlled by the Bridge Group Names (BGNs). Mesh APs can be placed in similar bridge groups to manage membership or provide network segmentation.

Enterprise 11n/ac mesh is added to the Catalyst 9800 controller features to work with the 802.11n/ac APs. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but add higher backhaul and client access speeds. The 802.11ac indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP, and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4-GHz radios in Release 17.1 can be used for local (client) access as well as backhaul. Enterprise 11ac mesh supports point-to-point, point-to-multipoint, and mesh types of architectures.

In the current architecture, while an autonomous AP functions as a WGB, only one radio interface is used for controller connectivity, one Ethernet interface for wired client connectivity, and the other radio interface for wireless client connectivity. dot11radio 1 (5 GHz) can be used to connect to a controller (using the mesh infrastructure) and Ethernet interface for wired clients, dot11radio 0 (2.4 GHz) can be used for wireless client connectivity. Depending on the requirement, dot11radio 1 or dot11radio 0 can be used for client association or controller connectivity.

With two radios, one radio can be used for client access and the other radio can be used for accessing the APs. Having two independent radios performing two independent functions provides you better control and lowers the latency. Also, wireless clients on the second radio for the WGB do not get disassociated by the WGB when an uplink is lost or in a roaming scenario. One radio must be configured as a RAP (radio role), and the second radio must be configured as a WGB (radio role).

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

Cisco's Catalyst 9100 and Aironet Access Points support multiple radios, including one or more radios for client access and one or more radios for wireless backhaul and flexible radio assignment. For instance, Cisco Catalyst 9100 and Aironet Access Points can operate in tri-radio mode (known as split 5 GHz mode) or XOR mode (the ability to operate a single radio interface in many modes), which allows one 5 GHz radio to be split into two separate 5 GHz radios. The access point carries out the radio reassignment, and this enables the access point to utilize three radios simultaneously: one 2.4 GHz radio and two 5 GHz radios. The radios can serve different user groups on different channels.

Table 3. AP models and types of hardware managed by FRA

AP model	FRA radios	Functions
Cisco Aironet 2800 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Aironet 3800 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Aironet 4800 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Catalyst 9120 Series Access Points	2.4/5 XOR	2.4-GHz and 5-GHz or dual 5-GHz operations
Cisco Catalyst 9130 Series Access Points	5-GHz Tri-Radio	2.4-GHz 4x4 and single 5-GHz 8x8, or 2.4-GHz 4x4 and dual 5-GHz 4x4
Cisco Catalyst Wireless 9166 Access Points	5/6-GHz XOR	2.4-GHz 4x4 and dual 5-GHz 4x4, or 5-GHz 4x4 and 6-GHz 4x4

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html

For example, Cisco Catalyst 9130AX Wi-Fi 6 access point is tri-radio (2.4 GHz + 5 GHz with dual-5GHz using Flexible Radio Assignment) that supports mesh. One radio can be used for client access, the others used flexibly for backhaul (uplink and downlink) at different channels and client access.

Flexible Radio Assignment (FRA) with tri-radio mode	FRA allows the access points to intelligently determine the operating mode of serving radios based on the RF environment and traffic demands. The access points can operate in the following modes: • Dual-radio mode: One 8x8 5 GHz and one 4x4 2.4 GHz. One radio will serve clients on the 5-GHz band, while the other serves clients on the 2.4-GHz band. • Tri-radio mode: * Dual 4x4 5 GHz and one 4x4 2.4 GHz. With two 4x4 5-GHz and one 4x4 2.4-GHz radios (tri-radio) inside the access point, client device capacity can be increased on demand. The access point's default mode is dual radio with 8x8 5 GHz and 4x4 2.4 GHz. It has the ability to split the 8x8 radio into two separate 4x4 5-GHz radios through software, thereby enabling the benefits of FRA while allowing the 2.4-GHz radio to remain active.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat-9130-ser-ap-ds-cte-en.html; see also https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-4800-access-point/nb-09-air-4800-acces-ds-cte.html (Aironet 4800 with four radios); https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat9124-ser-ap-ds-cte-en.html (Catalyst 9124AX with similar radios); https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html (mesh support included in 9130AX access points)

Cisco Access Points with Tri-Radio

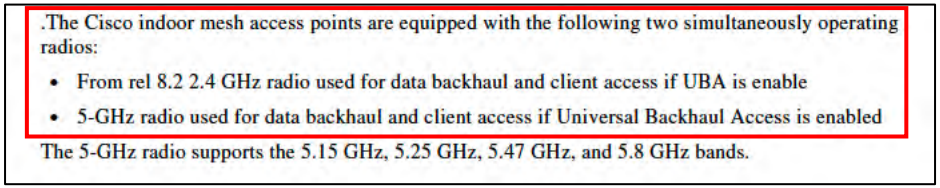
This topic describes the Tri-Radio feature for Cisco Access Points (APs).

Access Points with three radios are designed for high density environments. The APs by default run one dedicated 2.4-GHz 4x4 mode radio and one 5-GHz 8x8 mode radio. In the default mode, the radios are managed by the Flexible Radio Assignment (FRA), and the Dual Radio Mode is in the disabled state indicating that the radios have either been assigned as client serving 8x8 radio or have not yet been evaluated by FRA.

When you enable the dual radio mode setting, the 8x8 radio is split to two independent 5-GHz 4x4 radios. In this mode, slot 1 and slot 2 are active independent 4x4 radio interfaces. They can serve different user groups with different assigned channels.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-2/config-guide/b_wl_17_2_cg/cisco_access_points_with_tri_radio.pdf

In addition, while 5 GHz is often used for the primary backhaul (communication between access points), both the 2.4 GHz and 5GHz radios are configurable to be used for backhaul and client access (*e.g.*, "dual purpose radio"), for example, when Universal Backhaul Access (UBA) is enabled.



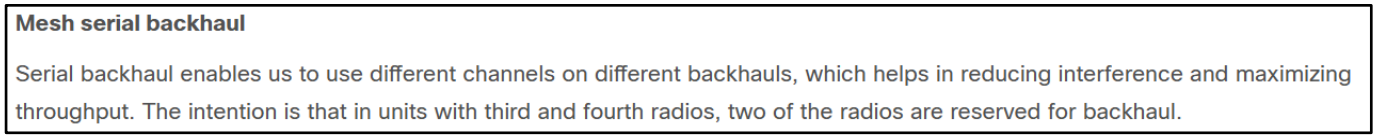
.The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios:

- From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable
- 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled

The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands.

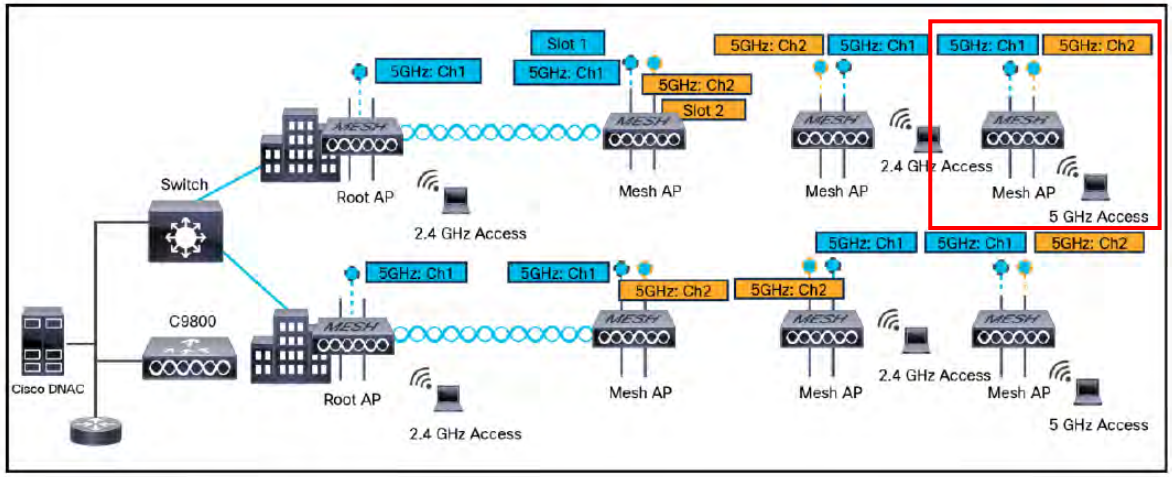
Source: https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-6/b_mesh_86/m_mesh-network-components.html; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/technotes/8-5/cisco-enterprise-mobility-design-guide-8-5.pdf

With Mesh Serial Backhaul, two of the radios are reserved for backhaul. The Mesh Serial Backhaul feature in a mesh access point (MAP) allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. Basic client access functionality is offered on the 2.4-GHz radio and the 5-GHz radio, which are not used in serial backhaul. Universal access is made available on the downlink radio (*e.g.*, "dual purpose radio").



Mesh serial backhaul

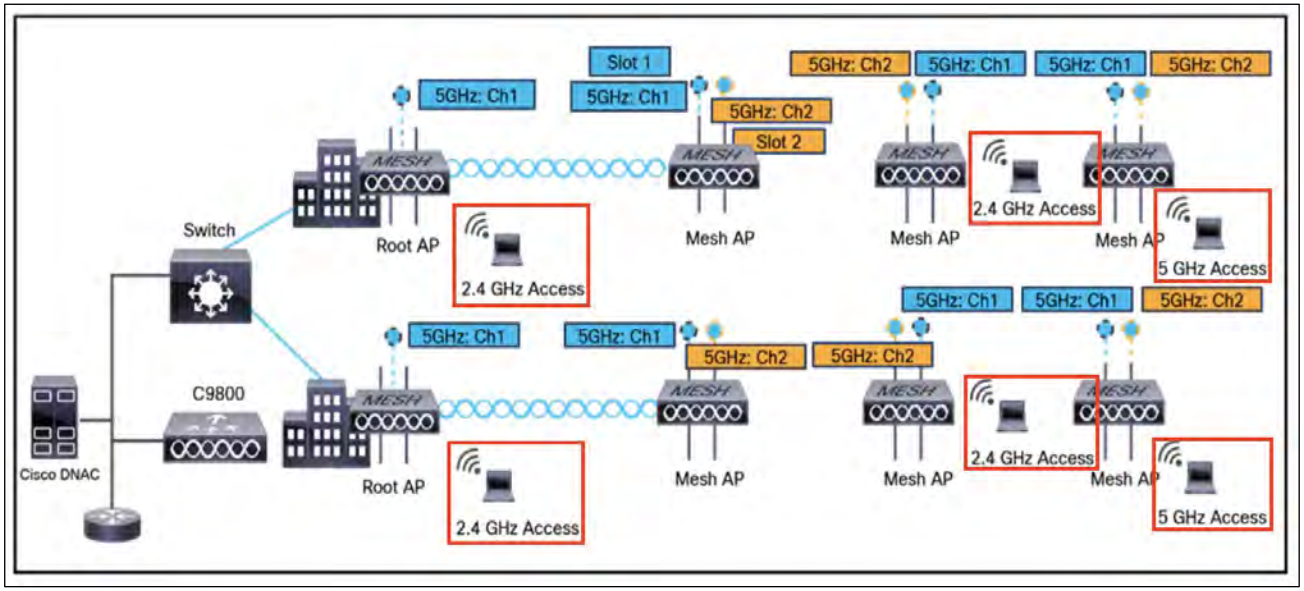
Serial backhaul enables us to use different channels on different backhauls, which helps in reducing interference and maximizing throughput. The intention is that in units with third and fourth radios, two of the radios are reserved for backhaul.



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-7/config-guide/b_wl_17_7_cg/m_mesh_ewlc.html#mesh-backhaul-24ghz-5ghz

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 15	
15. The wireless mesh network of claim 1 further comprising a client device wirelessly connected to the MAP node, wherein the client device exchanges data with the external network through the MAP's routing path to the external network.	The Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises a client device wirelessly connected to the MAP node, wherein the client device exchanges data with the external network through the MAP's routing path to the external network, as shown in Cisco's documentation below: **CAPWAP** The Internet Engineering Task Force (IETF) Control and Provisioning of Wireless Access Points (CAPWAP) protocol is the underlying protocol used in the Cisco WLAN architecture. CAPWAP provides the configuration and management of APs and WLANs in addition to encapsulation and forwarding of WLAN client traffic between an AP and a WLAN Controller (WLC). CAPWAP brings additional security with Datagram Transport Layer Security (DTLS) for both management and client traffic. CAPWAP uses the User Datagram Protocol (UDP) and can operate over either IPv4 or IPv6. Table 1 lists the protocol and port implemented for each CAPWAP version.

	CAPWAP data encryption is optional and is enabled per AP. Data encryption requires a DTLS license to be installed on the WLC prior to being enabled on an AP. When enabled, all WLAN client traffic is encrypted at the AP before being forwarded to the WLC and vice versa. DTLS data encryption is automatically enabled for teleworker APs but is disabled by default for all other APs. Most APs are deployed in a secure network where data encryption is not necessary. In contrast, traffic exchanged between a teleworker AP and WLC is forwarded over an unsecured public network, where data encryption is important. **Network access** Wireless mesh networks can simultaneously carry two different traffic types: • WLAN client traffic • MAP Ethernet port traffic WLAN client traffic terminates on the controller, and the Ethernet traffic terminates on the Ethernet ports of the mesh APs.  Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh
Claim 16	
16.a. The wireless mesh network of claim 1 further comprising one or more data	The Cisco Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises one or more data types being routed through the routing path, wherein the one or more data types include at least two of: voice data, video data, and other data type that is not the voice or video data, as shown in Cisco's documentation below:

types being routed through the routing path, wherein the one or more data types include at least two of: voice data, video data, and other data type that is not the voice or video data;	**Different deployment types for WLAN coverage** The amount of WLAN coverage you plan for in the design of your wireless network depends largely on the usage and density of clients you require. With limited exceptions, all designs should be deployed to minimize retransmission and data rate shifting while supporting good client roaming and throughput. Wireless networks can be deployed for data-only, voice, video, and location-aware services or, more frequently these days, a combination of all of these. The difference between these application types is minimal today, with the requirements of each largely describing good, solid, capacity-based coverage. Location-aware services add some AP placement criteria for good location triangulation and guidelines on hyperlocation technologies. Real-time multimedia (voice and video) applications have different latency requirements for two-way live implementations. But by and large, all describe a minimum coverage level needed to make the application viable for the number of users you expect in any given area. For most campuses and enterprise installations, coverage and capacity are the primary concerns, and these are easily achievable. High-density client implementations or high-interference locations such as shopping malls or apartment buildings may require additional equipment such as external antennas to properly implement the network to scale. The sections that follow provide in-depth information on application-specific guidelines, recommendations, and configurations.
16.b. wherein routing a priority data type with a higher priority than another type of data, wherein the priority data type may be one or both of: the voice data or the video data.	The Cisco Wireless Mesh Network comprises one or more data types being routed through the routing path, wherein routing a priority data type with a higher priority than another type of data, wherein the priority data type may be one or both of: the voice data or the video data. The Cisco Wireless Mesh Network supports Quality of Service model which provides priority processing of data based on data type, including voice and video data: **DSCP trust** “DSCP trust” is the QoS model supported by the Catalyst 9800 Series. This means that all QoS processing (queuing and policies) applied to the wireless traffic within the AP and WLC is based on the client DSCP value and not the 802.11 User Priority (UP). For example, for a centralized switching SSID in the downstream direction (wired to wireless traffic), the AP takes the DSCP value from the received CAPWAP header and uses it for internal QoS processing and mapping (received DSCP > UP > access category). The DSCP value is mapped to the UP value in the frame to the wireless client using the data in the table below according to RFC 8325. It is a modular approach, as the name Modular QoS CLI implies. First, classify the traffic to consider. Classification can be done by matching on a user-defined condition (specified with an access list or a DSCP value) or by matching on an application using the Network-Based Application Recognition (NBAR) engine. In the example, the user first defines an ACL to match the CAPWAP control traffic (port UDP = 5246 for CAPWAP between AP and WLC, and UDP = 1666 for mobility group tunnels). The next step is to define a class map to match based on the ACL and DSCP = EF for voice traffic. Then the actual policy is defined: assign priority queue 1 for traffic identified as CAPWAP control; for voice traffic, use priority queue 2. The last step is to apply the policy to the target; in this case, it is the uplink port in the output or egress direction.

Table 10. DSCP to UP map

IETF DiffServ service class	DSCP	802.11 user priority	801.11 access category
Network control	CS6, (CS7)	0	AC_BE
IP telephony	EF	6	AC_VO
VOICE-ADMIT	44	6	AC_VO
Signaling	CS5	5	AC_VI
Multimedia conferencing	AF4x	4	AV_VI
Real-time interactive	CS4	5	AC_VI
Multimedia streaming	AF3x	4	AC_VI
Broadcast video	CS3	4	AC_VI
Low-latency data (transactional)	AF2x	3	AC_BE
Operations, Administration, and Maintenance (OAM)	CS2	0	AC_BE
High-throughput data (bulk data)	AF1x	2	AC_BK
Low-priority data (scavenger)	CS1	1	AC_BK
Remaining	Remaining	0	AC_BE

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#6QualityofServiceQoS

	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 17	
17.a. A wireless mesh network comprising: a plurality of Wi-Fi nodes in the wireless mesh network organized in a tree shape;	*See supra* discussion for claim element 1.a., the contents of which are incorporated herein by reference.
17.b. an access server in data communication with the nodes,	*See supra* discussion for claim element 1.b., the contents of which are incorporated herein by reference.
17.c. wherein the nodes are comprised of: one or more root access point (RAP) nodes having a wired connection to an external network; one or more mesh access point (MAP) nodes;	*See supra* discussion for claim element 1.c., the contents of which are incorporated herein by reference.
17.d. wherein a MAP node is in a data communication with the external network through an associated RAP node;	*See supra* discussion for claim element 1.d., the contents of which are incorporated herein by reference.
17.e. wherein the MAP node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data	*See supra* discussion for claim element 1.e., the contents of which are incorporated herein by reference.

communication with the external network through its associated RAP node;	
17.f. wherein the MAP node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node, process for the automatic connection comprising:	*See supra* discussion for claim element 1.f., the contents of which are incorporated herein by reference.
17.g. selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node; and	*See supra* discussion for claim element 1.g., the contents of which are incorporated herein by reference.
17.h. establishing a direct wireless data connection between the MAP node and the associated parent node;	*See supra* discussion for claim element 1.h., the contents of which are incorporated herein by reference.
17.i. wherein the automatic connection process establishes the MAP node's routing path to the external network;	*See supra* discussion for claim element 1.i., the contents of which are incorporated herein by reference.

17.j. the one or more parent selection criteria comprising: a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path;	*See supra* discussion for claim element 1.j., the contents of which are incorporated herein by reference.
17.k. the MAP node comprising: one or more identifier values representing a unique identity of the MAP node in the wireless mesh network;	*See supra* discussion for claim element 1.k., the contents of which are incorporated herein by reference.
17.l. two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network;	*See supra* discussion for claim element 1.l., the contents of which are incorporated herein by reference.
17.m. one or more network health monitoring criteria defined by the access server, wherein the one or more network health monitoring criteria comprises one or more of: whether or not the associated parent node is offline or can be found in the wireless mesh network;	*See supra* discussion for claim element 2.a., the contents of which are incorporated herein by reference.

whether or not data can be routed to the external network through the MAP node's routing path; whether or not the one or more parent selection criteria is satisfied; and whether or not the MAP node's routing path is more congested than a threshold load value representing data congestion level of a routing path;	
17.n. testing whether the MAP node's routing path fails to satisfy the one or more network health monitoring criteria;	*See supra* discussion for claim element 2.b., the contents of which are incorporated herein by reference.
17.o. when it is determined that the MAP node's routing path does not meet one or more of the one or more network health monitoring criteria, establishing a new routing path from the MAP node to the external network by connecting the MAP node to a new associated parent node through the automatic connection process.	*See supra* discussion for claim element 2.c., the contents of which are incorporated herein by reference.
Claim 18	
18.a. A wireless mesh network comprising:	*See supra* discussion for claim element 1.a., the contents of which are incorporated herein by reference.

a plurality of Wi-Fi nodes in the wireless mesh network organized in a tree shape;	
18.b. an access server in data communication with the nodes,	*See supra* discussion for claim element 1.b., the contents of which are incorporated herein by reference.
18.c. wherein the nodes are comprised of: one or more root access point (RAP) nodes having a wired connection to an external network; one or more mesh access point (MAP) nodes;	*See supra* discussion for claim element 1.c., the contents of which are incorporated herein by reference.
18.d. wherein a MAP node is in a data communication with the external network through an associated RAP node;	*See supra* discussion for claim element 1.d., the contents of which are incorporated herein by reference.
18.e. wherein the MAP node is in a direct wireless data communication with an a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node;	*See supra* discussion for claim element 1.e., the contents of which are incorporated herein by reference.
18.f. wherein the MAP node automatically connects to the associated parent node by	*See supra* discussion for claim element 1.f., the contents of which are incorporated herein by reference.

selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node, process for the automatic connection comprising:	
18.g. selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node; and	*See supra* discussion for claim element 1.g., the contents of which are incorporated herein by reference.
18.h. establishing a direct wireless data connection between the MAP node and the associated parent node;	*See supra* discussion for claim element 1.h., the contents of which are incorporated herein by reference.
18.i. wherein the automatic connection process establishes the MAP node's routing path to the external network;	*See supra* discussion for claim element 1.i., the contents of which are incorporated herein by reference.
18.j. the one or more parent selection criteria comprising: a loop prevention criterion which verifies that a potential parent node's routing path to the external	As discussed for claim element 1.j, the one or more parent selection criteria comprising: a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path. *See supra* discussion for claim element 1.i. and 1.j. The one or more parent selection criteria further comprises a loop prevention criterion which verifies that a potential parent node's routing path to the external network does not contain the MAP node:

network does not contain the MAP node;

a throughput value representing signal strength between two nodes; and

a latency value representing number of hops in a routing path;

Traffic flow

The traffic flow within the wireless mesh can be divided into three components:

- Overlay CAPWAP traffic that flows within a standard CAPWAP access point deployment; that is, CAPWAP traffic between the CAPWAP AP and the CAPWAP controller.
- Wireless mesh data frame flow.
- AWPP exchanges.

As the CAPWAP model is well known and AWPP is a proprietary protocol, only the wireless mesh data flow is described. The key to the wireless mesh data flow is the address fields of the 802.11 frames being sent between mesh APs.

An 802.11 data frame can use up to four address fields: receiver, transmitter, destination, and source. The standard frame from a WLAN client to an AP uses only three of these address fields because the transmitter address and the source address are the same. However, in a WLAN bridging network, all four address fields are used because the source of the frame might not be the transmitter of the frame, because the frame might have been generated by a device behind the transmitter.

The figure below shows an example of this type of framing. The source address of the frame is MAP:03:70, the destination address of this frame is the controller (the mesh network is operating in Layer 2 mode), the transmitter address is MAP:D5:60, and the receiver address is RAP:03:40.

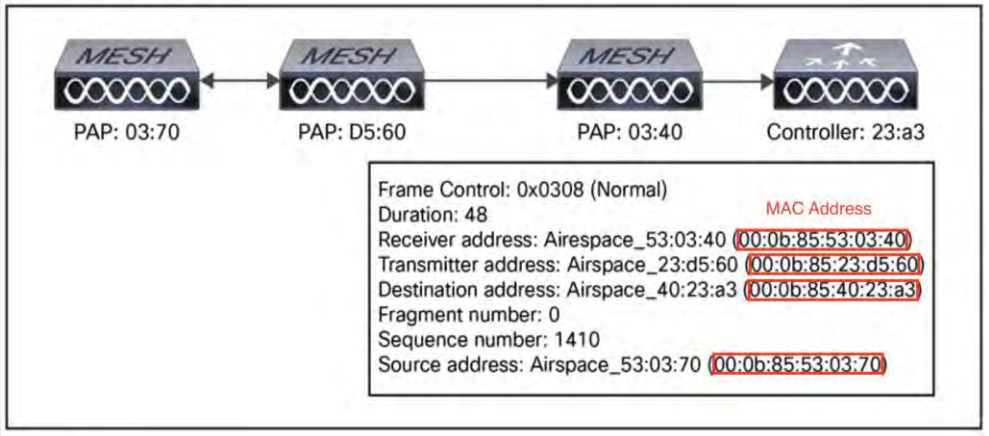


Figure 71.

Wireless mesh frame

Loop prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
18.k. the MAP node comprising: one or more identifier values representing a unique identity of the MAP node in the wireless mesh network;	*See supra* discussion for claim element 1.k., the contents of which are incorporated herein by reference.
18.l. two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network.	*See supra* discussion for claim element 1.l., the contents of which are incorporated herein by reference.
Claim 19	
19. The wireless mesh network of claim 18 wherein the selection of the parent node further comprises not selecting a potential parent node if the MAP node's potential routing path to the external network through the potential parent node exceeds a threshold latency value.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network, in addition to the elements of claim 18 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises not selecting a potential parent node if the MAP node's potential routing path to the external network through the potential parent node exceeds a threshold latency value. The number of hops (*i.e.*, a threshold latency value) is set for the Cisco Wireless Mesh Network, such that a parent will not be selected if the potential routing path back to the external network exceeds a certain number of hops. *See also* claim element 1.g.

Note With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

- Number of backhaul hops is limited to eight but we recommend three to four hops.

  The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.

- Number of MAPs per RAP.

  There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Exhibit 8 - U.S Patent No. 11,368,537 (Cisco Wireless Mesh Networks with Meraki and Catalyst CW Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Meraki Wireless Mesh Network including but not limited to Cisco's Meraki and Catalyst CW Access Points, with related cloud architecture and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1-19 (the "Asserted Claims") of U.S. Patent No. 11,368,537 (the "'537 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '537 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringed by testing, setting up, and using Cisco Meraki Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Meraki Wireless Mesh Network in the United States. As indicated in the Meraki Mesh Networking Guide, each Meraki access point has unique mesh characteristics independent of general technology specifications listed on the datasheet. Cisco recommends a Meraki Systems Engineer as a resource to ensure that an access point's capabilities match the desired design targets.[3] The Meraki Trial Program allows customers to run a real world proof of concept to test the design, which is arranged by a Cisco Meraki account representative. Meraki is also cloud-based architecture, and the infringing mesh networking methods and systems were controlled and put into service by Cisco.[4] Cisco knew how its customers were configuring the specific access points sold, as customers must register to the Meraki cloud and use the invoice or serial number of each purchased access point.[5]

---

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g*., https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

[3] *See, e.g*., https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

[4] *See Centillion Data Systems, LLC v. Qwest Communications International, Inc.*, 631 F.3d 1279 (Fed. Cir. 2011); *NTP, Inc. v. Research in Motion, Ltd*., 418 F. 3d 1282, 1317-1318 (Fed. Cir. 2005); *see also SiRF Tech. Inc. v. ITC*, No. 2009-1262 (Fed. Cir. Apr. 12, 2010) (direct infringement where the claims do not specifically require action by its customers or end users).

[5] *See, e.g*., https://documentation.meraki.com/MR/MR_Installation_Guides/CW9163E-MR_Installation_Guide.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the ’537 Patent and how to make and use the infringing Cisco Meraki Wireless Mesh Network. Specifically, Cisco’s Meraki Mesh Deployment Guides, Meraki Mesh Networking Guide, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco’s website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Meraki Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1-19 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Meraki Wireless Mesh Network, and the infringing configuration and operation of the cloud, access points, and software that comprise the infringing Cisco Meraki Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Meraki Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). MeshDynamics put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Meraki Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’537 Patent, or in the alternative was willfully blind to the ’537 Patent and the infringement, following the MeshDynamics letter in 2022.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Meraki Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’537 Patent. Cisco Accused Products come enabled as a mesh access point and automatically connect to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particular suited for Cisco’s Meraki Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Meraki Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’537 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the ’537 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends, in the alternative, that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference

between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics's preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco discovery, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco's Meraki Wireless Mesh Network with Meraki and Catalyst CW Access Points
1.a. A wireless mesh network comprising: a plurality of Wi-Fi nodes in the wireless mesh network organized in a tree shape;	Cisco's Meraki Wireless Mesh Network comprises a plurality of Wi-Fi nodes (*e.g.*, Cisco Meraki and Catalyst CW Access Points) organized in a tree shape. The Meraki cloud-managed wireless series is a line of enterprise-grade WLAN access points ("APs") that leverage the Cisco Meraki cloud-managed architecture. Both Cisco Meraki and Catalyst CW Access Points support Wi-Fi capabilities and can be paired with the Meraki cloud platform. Meraki AP models start with MR, *e.g.* MR33, MR36, MR46E, and the Catalyst APs with CW, *e.g.* CW9162, CW9164, CW9166, are dual-mode APs which can be used with the Catalyst 9800 controller or the Meraki cloud architecture. Source: Meraki Family Datasheet; Meraki Wi-Fi Models

Cisco Meraki Cloud Management

Pairing the Cisco Catalyst 9166 Series Access Points with the Meraki cloud platform gives organizations a unified IT experience for network monitoring and management. The Meraki dashboard provides an intuitive and interactive web interface connecting your network to the industry's leading cloud IT platform.

Through the dashboard, Meraki provides sophisticated and scalable tools to automate network optimization, deploy policy and segmentation configurations across thousands of sites and devices, and manage a full-stack network from SD-WAN to Access to IoT technologies. The platform supports over 3.5 million active networks around the world.

Working together, the Catalyst 9166 Series and Cisco Meraki offer such features as:

- Cisco Spaces
- Cisco Identity Services Engine
- Meraki Health intelligent optimization and assurance
- Meraki Vision, smart cameras, and sensors for network closet monitoring

Source: CW9166 Datasheet; CW9166

The Meraki and Catalyst CW Access Points contain the same mesh networking capabilities, as shown Cisco's documentation and the Meraki Family Datasheet, and the Meraki Wireless Mesh Networking and Mesh Deployment Guide apply to both product lines for setting up a Cisco Meraki Wireless Mesh Network. Multiple access points communicate over wireless interfaces to form a single network. Meraki's mesh networking functionality is automatic, self-healing and available on all Meraki (including Catalyst CW) Access Points.

cisco Meraki

Wireless Mesh Networking

In a wireless mesh deployment, multiple Access points (with or without Ethernet connections) communicate over wireless interfaces to form a single network. This wireless communication between Access points is called Mesh Networking. Meraki's mesh networking functionality is automatic, self-healing and available on all access points.

A detailed article on designing wireless Mesh networks can be found on Mesh Deployment Guide.

Source: Wireless Mesh Networking; *see also* Meraki Best Practices - Wireless Mesh Networking

Mesh functionality is enabled by default in every Meraki AP. All Meraki APs support wireless meshing without any additional configuration or licensing.

Access Point Capabilities

All Meraki access points support wireless meshing without any additional configuration or licensing. The most successful mesh deployments undergo both pre-deployment design and post-design testing phases.

Source: Mesh Deployment Guide

Class-leading enterprise features

Meraki MR wireless access points come equipped with industry-leading features that make them ideal for demanding enterprise deployments:

- Self-configuring plug-and-play deployment
- 802.11be MU-MIMO with up to sixteen spatial streams built for voice and video
- Integrated Bluetooth® IoT radio
- IoT ready (ESL integration)
- Self-healing, zero-configuration mesh

Source: Meraki Family Datasheet

Mesh networking, included in every Meraki AP, extends coverage to hard-to-wire areas and creates a self-healing network that is resilient to cable and switch failures, continuing to operate despite failures or configuration changes in the rest of the network without the need for manual configuration or optimization.

Source: Meraki Family Datasheet

Cisco's Meraki Wireless Mesh Network organizes the wireless access points (such as the MR and CW series Access Points) in a tree shape topology, as shown below.

Typical Deployment Models

Point-to-Point

In the point-to-point mesh scenario, one gateway access point provides access to one repeater access point. By default, the repeater access point will broadcast the same SSIDs as the gateway access point. Access point tags can be used to disable specific SSIDs on the repeater access point; Using Tags to Broadcast SSIDs from Specific APs has more information about this. The repeater access point can extend an Ethernet segment from the gateway access point at the same time as providing wireless access to clients. Further discussion on Ethernet bridging can be found in Ethernet Bridging below.

Applications for which this deployment scenario is appropriate include:

- Bridge link between a main office and a remote temporary office
- Extend the reach of infrastructure-based wireless network at a campus to less critical, hard-to-reach areas like sports fields or parking lots
- Extend a secure network from an industrial core to a non-mission-critical temporary roadhouse or satellite building



Point to Multi-Point / Multi-Point to Multi-Point

In the point to multi-point scenario, one or more gateways serve multiple repeaters. All gateways and repeaters may broadcast the same SSIDs and serve wireless clients.

Applications for which this deployment scenario is appropriate include:

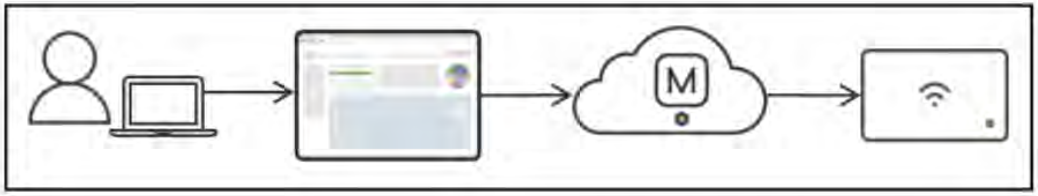
- Provide backhaul for multiple temporary buildings
- Reach remote wireless access point clusters
- Cover low bandwidth subscribers over a large area

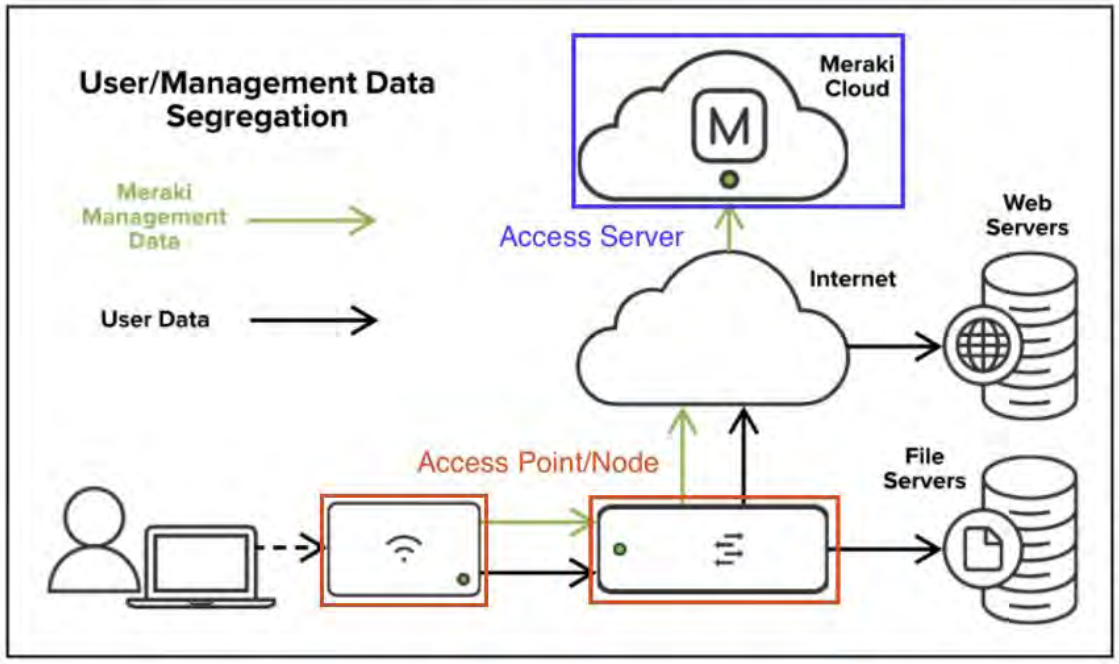


ⓘ Any gateway may be used by a mesh repeater in the same Meraki dashboard network. A separate network may be used if more control is required.

Source: Mesh Deployment Guide

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (Wi-Fi communication networking) in the same way (by deploying a plurality of Wi-Fi nodes organized in a "mesh tree") and achieves the same result (a wireless mesh network). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses a plurality of Wi-Fi nodes to create a

	mesh network organized as a "mesh tree." A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.
1.b. an access server in data communication with the nodes,	Cisco's Meraki Wireless Mesh Network includes an access server, *e.g.*, the Meraki cloud server, in data communication with the Cisco Meraki and Catalyst CW Access Points. The Cisco Meraki Wireless Mesh Network includes a Meraki cloud server, which is in data communication with the access points/nodes, *e.g.*, the MR and CW series Access Points, and is used to manage, configure, and monitor the access points/nodes: The Meraki cloud solution is a centralized management service that allows users to manage all of their Meraki network devices via a single, simple and secure platform.  Users are able to deploy, monitor, and configure their Meraki devices via the Meraki dashboard web interface or via APIs. Once a user makes a configuration change, the change request is sent to the Meraki cloud and is then pushed to the relevant device(s). **Management data:** The data (configuration, statistics, monitoring, etc.) that flows from Meraki devices (wireless access points, switches, security appliances) to the Meraki cloud over a secure internet connection. Management Data The Meraki cloud gathers and stores certain types of "management" data to enable its solutions. All forms of data are encrypted in transit to and from Meraki servers. There are four major types of data stored in the Meraki cloud: **User records:** Includes account email and company name or other optional information such as user name and address. **Configuration data:** Includes network settings and configurations made by customers in the Meraki dashboard. **Analytics data:** Includes client, traffic, and location analytics data, providing visualizations and network insights into traffic patterns across customer sites. **Customer-uploaded assets:** Includes custom floor plans and splash logos.

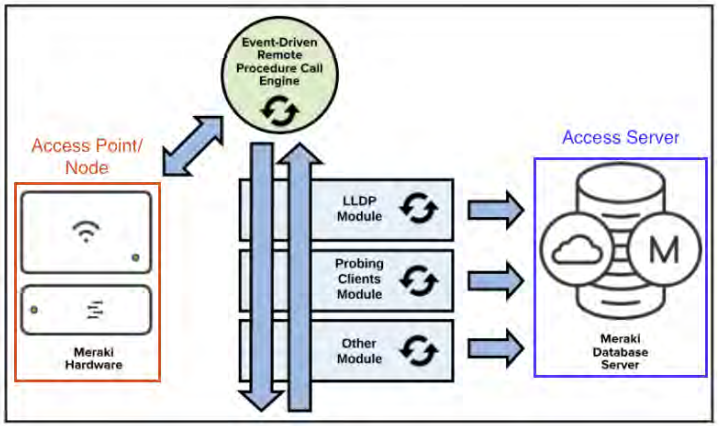


Meraki Cloud Architecture

The Meraki cloud is the backbone of the Meraki management solution. This "cloud" is a collection of highly reliable multi-tenant servers strategically distributed around the world at Meraki data centers and select public cloud service providers. For the remainder of this document we will refer to both the Meraki data centers and the public cloud service providers' locations we use as "data centers". The servers at these data centers are powerful hosting computers comprised of many separate user accounts. They are called multi-tenant servers because the accounts share (equal) computing resources on their host (the server). However, even though these accounts share resources, Meraki ensures that customer information is kept secure by restricting organization access based on account authentication, as well as hashing authentication information such as user passwords or API keys.

Meraki Device-to-Cloud Communications

Meraki uses an event-driven remote procedure call (RPC) engine for Meraki devices to communicate to the dashboard and for Meraki servers to send and receive data. Meraki hardware devices act as the server/receiver as the Meraki cloud initiates calls to the devices for data collection and configuration deployment. The cloud infrastructure is the initiator, so configurations can be executed in the cloud before the devices are actually online or even physically deployed.



Communication Process

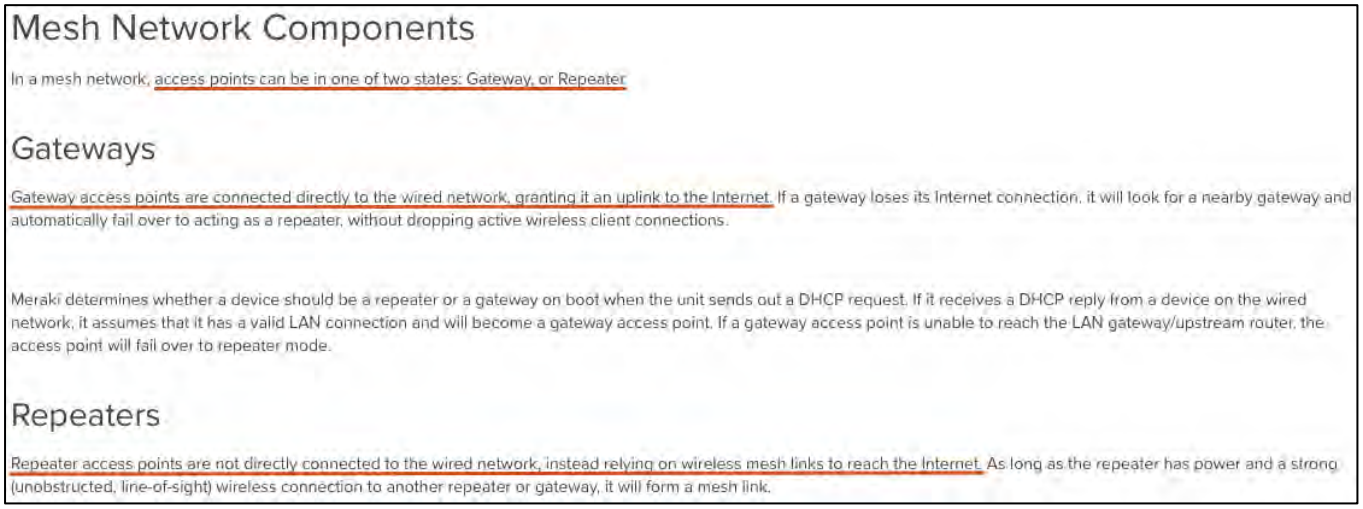
If a device is offline, it will continue to attempt to connect to the Meraki cloud until it gains connectivity. Once the device comes online, it automatically receives the most recent configuration settings from the Meraki cloud. If changes are made to the device configuration while the device is online, the device receives and updates these changes automatically. These changes are generally available on the device in a matter of seconds. However, large quantities of changes may take noticeably longer to reach their devices. If no configuration changes are made by the user, the device continues to periodically check for updates to its configuration on its own.

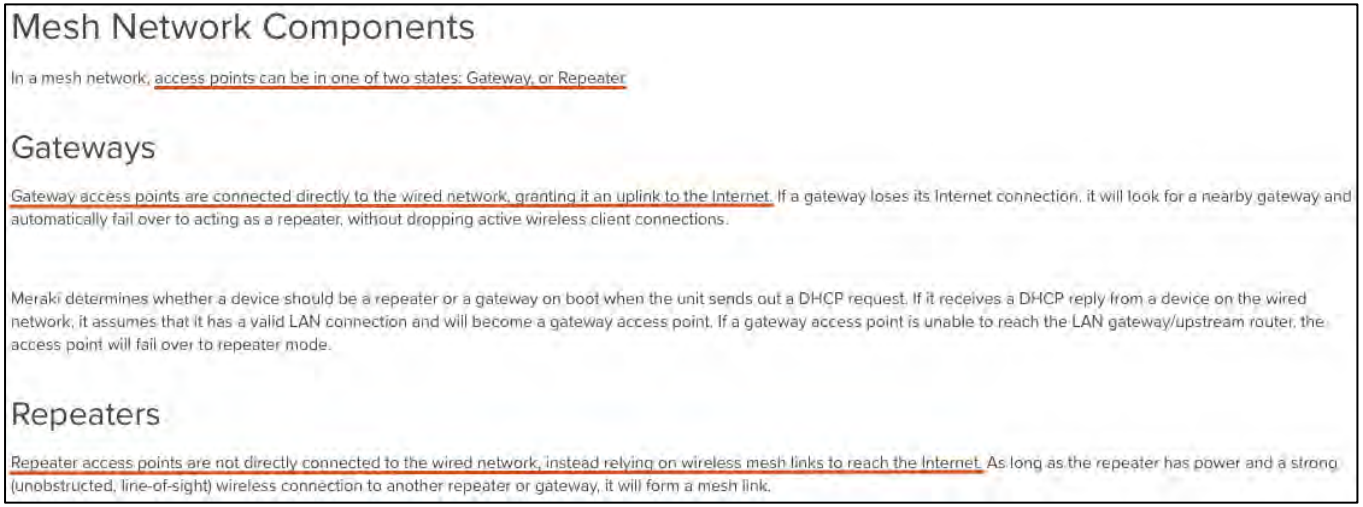
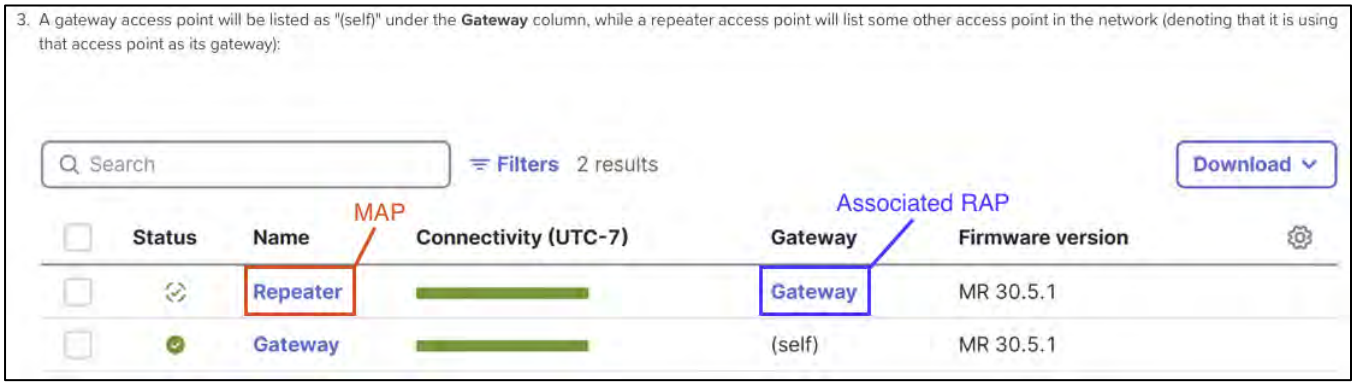
As the device runs on the network, it will communicate device and network usage analytics back to the Meraki cloud. Dashboard analytics based on this information, in the form of graphs and charts, are updated regularly in the Meraki cloud and are displayed in the dashboard of users when they are viewing this information.

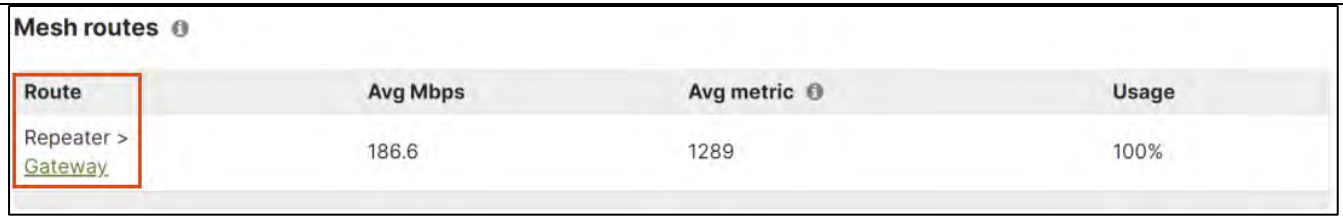
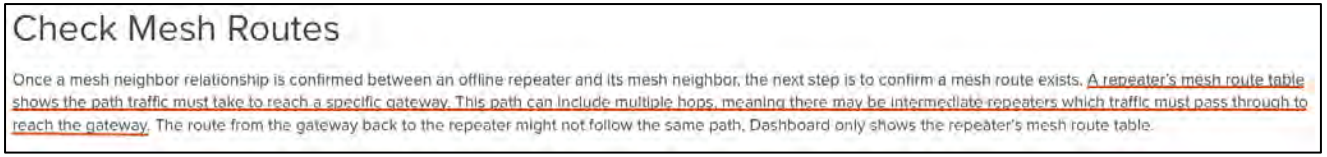
Source: Meraki Cloud Architecture

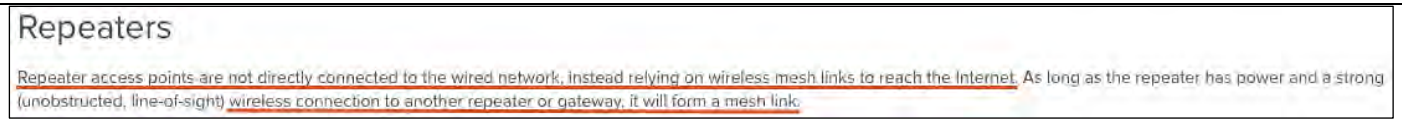
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (control of access points/node) in the same way (by deploying a server in data communication with the access points/node) and achieves the same result (control of access points/node). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses a cloud-based access server in data communication with the mesh network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1.c. wherein the nodes are comprised of: one or more root access point (RAP) nodes having a wired connection to an external network; one or more mesh access point (MAP) nodes;	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco Meraki Wireless Mesh Network are comprised of one or more gateways (*i.e.*, root access point (RAP) nodes) having a wired connection to an external network and one or more repeaters (*i.e.*, mesh access point (MAP) nodes): APs within the Cisco Meraki Wireless Mesh Network operate as either a gateway (*i.e.*, RAP node) or repeater (*i.e.*, MAP node). Gateway access points are connected directly to the wired network, granting it an uplink to the internet. Repeater access points are not connected to the wired network, and rely on wireless mesh links to reach the internet.  Mesh Network Components In a mesh network, access points can be in one of two states: Gateway, or Repeater Gateways Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode. Repeaters Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (wireless mesh network topology) in the same way (by deploying RAPs having a wired connection to an external network and MAPs) and achieves the same result (wireless mesh network). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses a cloud-based access server in data communication with the RAP and a wired connection to the external network, with one or more MAPs. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.
1.d. wherein a MAP node is in a data communication with the external network	The Meraki and Catalyst CW Access Points *(e.g.*, the MR and CW series access points) in the Cisco Meraki Wireless Mesh Network are comprised of one or more repeaters (*i.e.*, mesh access point (MAP) nodes) wherein the repeaters are in data communication with the external network through an associated gateway (i.e., RAP node).

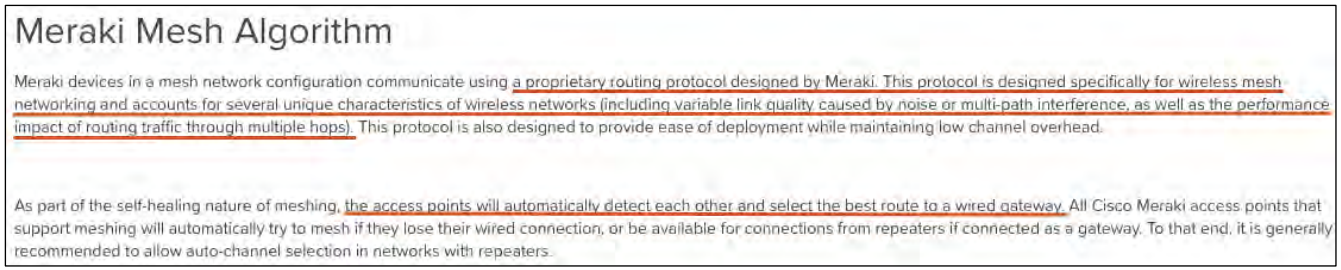
through an associated RAP node;	Repeater access points are not connected to the wired network, and rely on wireless mesh links to reach the internet.  Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide Each repeater in the Cisco Meraki Wireless Mesh Network is associated with a gateway, through which the repeater communicates with the external network:  Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide A repeater's mesh route(s) shows the path traffic that the repeater must take to reach a specific gateway, through which the repeater communicates with the external network:

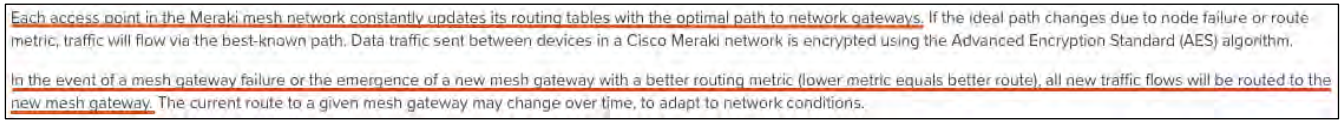
	 Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide  Source: Troubleshooting Mesh Communication between Access Points If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (data communication between MAPs and the external network) in the same way (through wireless connection with an associated RAP) and achieves the same result (data communication between MAPs and the external network). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses a RAP to connect MAPs to the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.e. wherein the MAP node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node;	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, mesh access point (MAP)) nodes wherein the repeater node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node. The repeaters (*i.e.*, MAPs) have wireless connections to the external network via an associated gateway (*i.e.*, RAP) either directly or through a wireless connection to another repeater:



Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

A repeater selects the best mesh route to a gateway based on the "Meraki Mesh Algorithm," which is a proprietary routing protocol designed by Meraki. On information and belief, the Meraki Mesh Algorithm operates the same or similar to Cisco's Adaptive Wireless Path Protocol ("AWPP") and forms parent-child relationships in the mesh, with parents selected based on criteria or characteristics of the wireless network including link quality and performance through multiple hops. Like AWPP (*see* Exhibit 1 for AWPP details), and the self-healing nature of meshing, the APs will automatically detect each other and select the best route to a wired gateway. A repeater's best route to a gateway may either be a direct wireless data communication with the gateway (*i.e.*, the associated parent node is the associated RAP node) or through wireless data communication with another repeater in data communication with the gateway (*i.e.*, the associated parent node is another MAP node in a data communication with the external network through its associated RAP node).





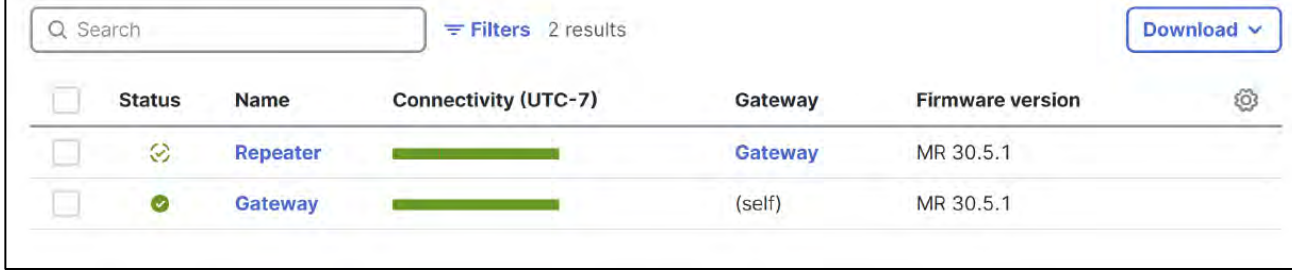
Source: Wireless Mesh Networking; *see also* Mesh Deployment Guide

Repeaters will identify the other AP in the network that it is using as its gateway, and its route back to the gateway, as well identify and track neighbors. Repeaters know the distance to its neighbors as well as the link quality. This is shown at the access server using identifier values to identify the repeaters, *e.g.*, "Name."

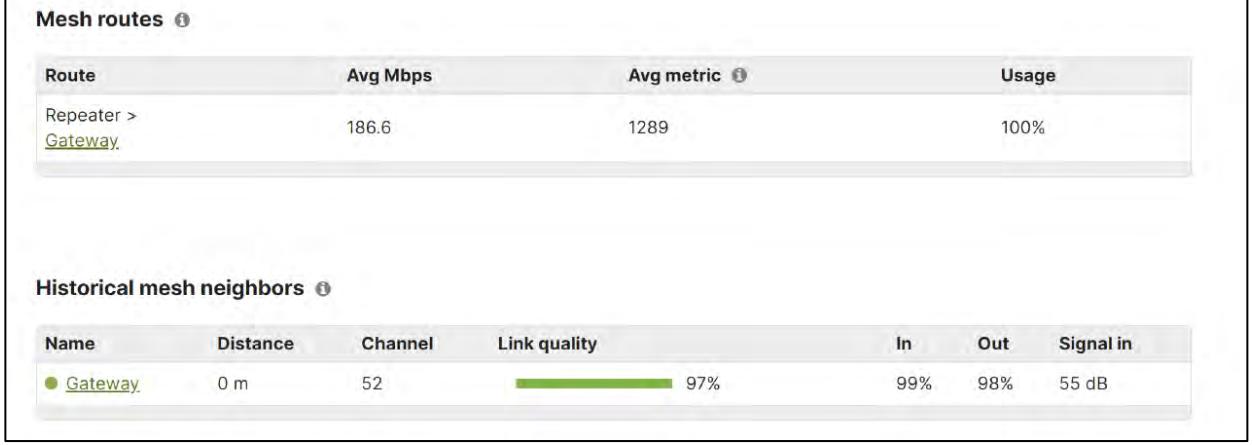
Identifying a Repeater access point vs Gateway access point

1. Navigate to **Wireless > Monitor > Access points**
2. Click the Gear icon on the right-hand side and make sure the **Gateway** option is selected:

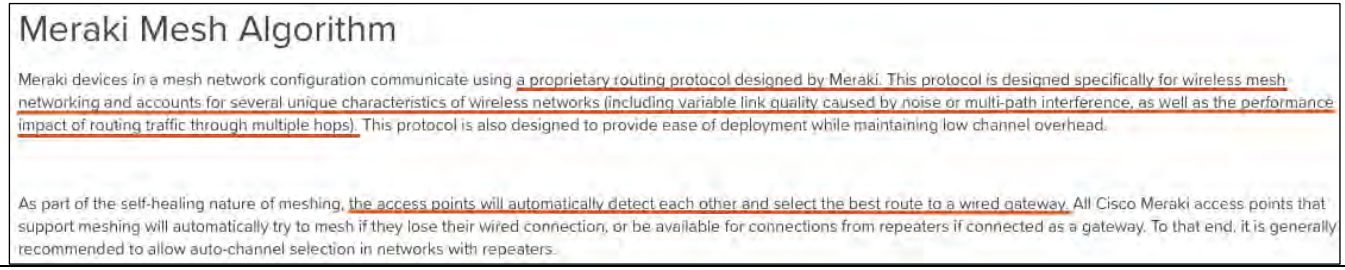
3. A gateway access point will be listed as "(self)" under the **Gateway** column, while a repeater access point will list some other access point in the network (denoting that it is using that access point as its gateway):



7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:



Source: Wireless Mesh Networking

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (data communication between MAPs and the external network) in the same way (through data communication with an associated parent node that is either the associated RAP or another MAP in data communication with the external network through its associated RAP node) and achieves the same result (data communication between MAPs and the external network). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses a gateway (*i.e.*, RAP) to connect repeaters (*i.e.*, MAPs) to the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.f. wherein the MAP node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node, process for the automatic connection comprising:	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes wherein the repeater node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the repeater node and establishing a parent-child relationship with the associated parent node. As discussed for claim element 1.e above, a repeater selects the best mesh route to a gateway based on the "Meraki Mesh Algorithm," which is a proprietary routing protocol designed by Meraki. On information and belief, the Meraki Mesh Algorithm operates the same or similar to Cisco's Adaptive Wireless Path Protocol ("AWPP") and forms parent-child relationships in the mesh, with parents selected based on criteria or characteristics of the wireless network including link quality and performance through multiple hops. Like AWPP (*see* Exhibit 1 for AWPP details), and the self-healing nature of meshing, the APs will automatically detect each other and select the best route to a wired gateway. A repeater's best route to a gateway may either be a direct wireless data communication with the gateway (*i.e.*, the associated parent node is the associated RAP node) or through wireless data communication with another repeater in data communication with the gateway (*i.e.*, the associated parent node is another MAP node in a data communication with the external network through its associated RAP node).  Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

Source: Wireless Mesh Networking

A repeater selects a gateway through which to communicate with the external network by scanning all channels to collect info from all neighbors. Based on the scan results, the repeater develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. The repeater selects the best mesh route from amongst all available routes based on the link quality and the number of hops provided by each neighbor. The neighbor which provides the best mesh route to a gateway (*i.e.*, parent node) is selected by the repeater.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

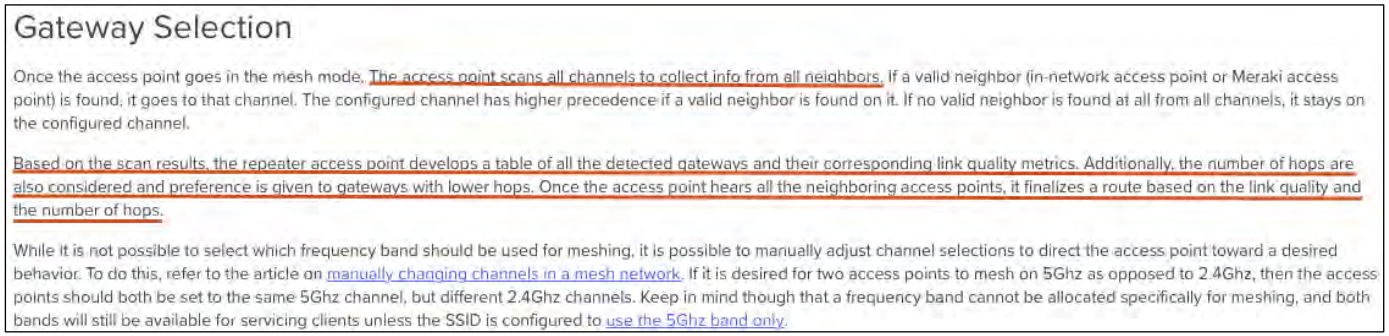
Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

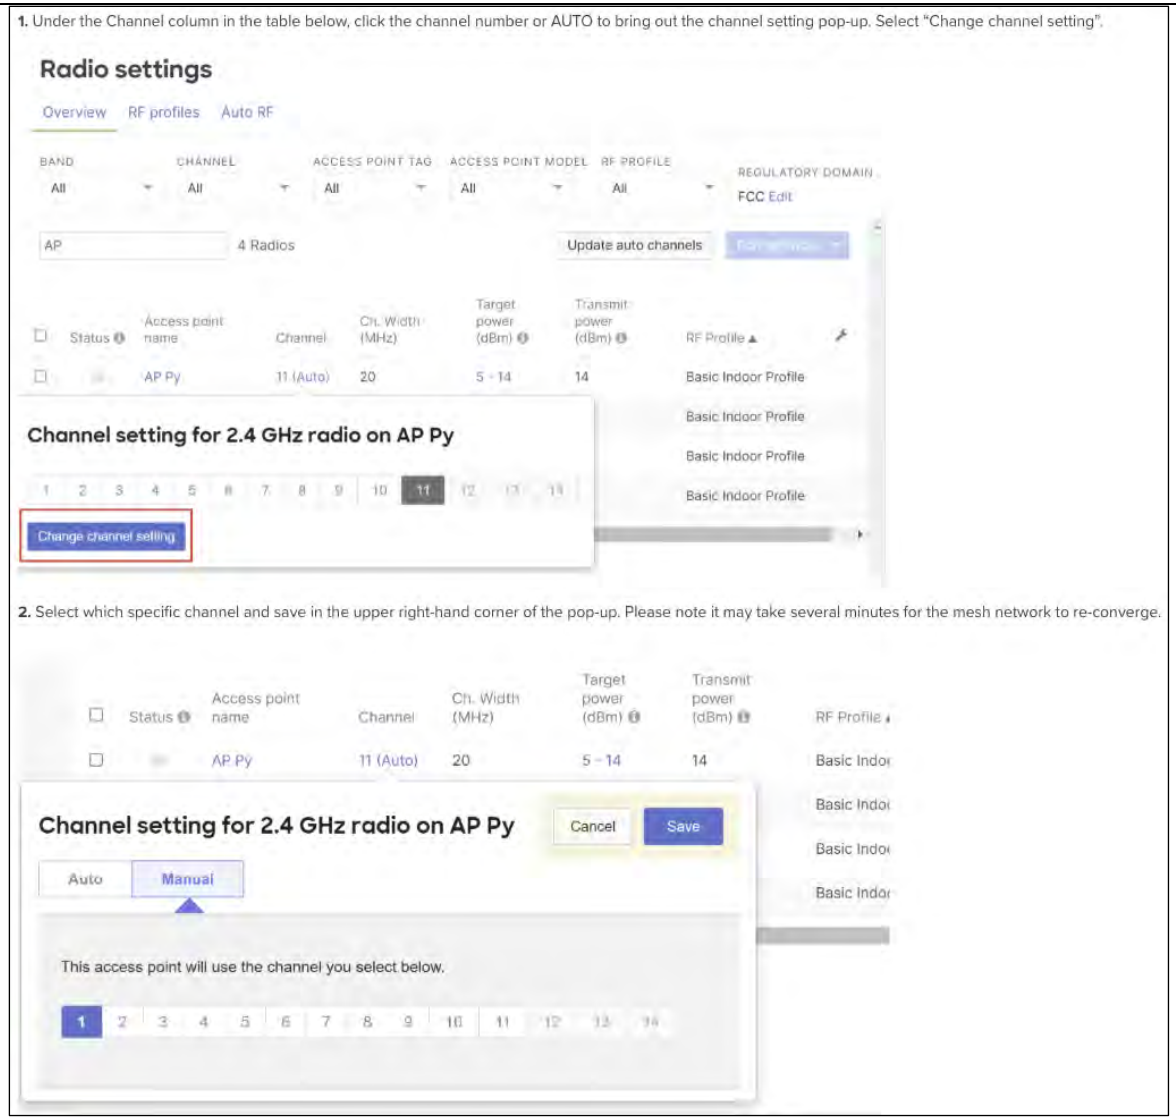
While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only.

Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (establishing wireless data communication between an MAP with the external network) in the same way (by selecting the best neighbor node based on certain criteria) and achieves the same result (communication with the external network via the best neighbor node (*i.e.*, parent node)). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses certain criteria to select a neighbor that provides the best route back to the gateway and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

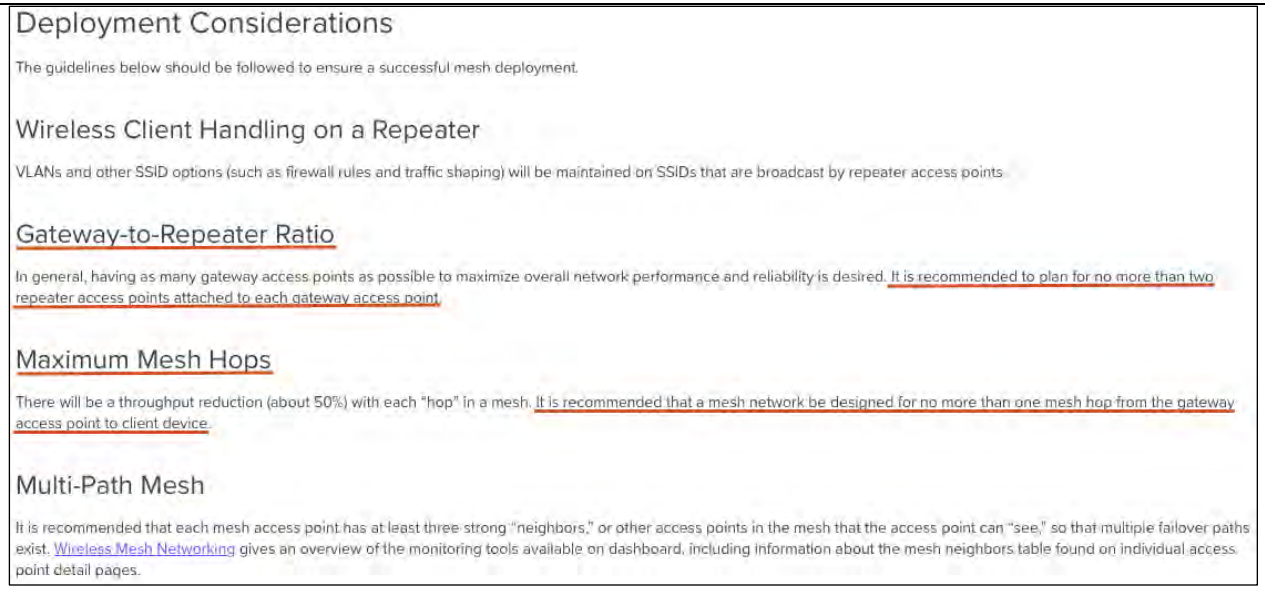
	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.g. selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node; and	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes wherein automatic connection to the associated parent node comprises selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node. The repeater selects the neighbor node providing the best mesh route to a gateway using the Meraki Mesh Algorithm based on, *inter alia*, channel configurations, link quality metrics, and number of hops: **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only. Source: Wireless Mesh Networking

	As discussed for claim element 1.b, the Cisco Meraki Wireless Mesh Network includes a Meraki cloud server (*i.e.*, access server), which is in data communication with the access points/nodes, *e.g.*, the MR and CW series Access Points, and is used to manage, configure, and monitor the access points/nodes. The Meraki cloud server provides the framework for repeaters to select the best mesh paths, including defining criteria for selecting the best neighbor (*i.e.*, parent). For example, the Meraki cloud can be used to configure an access point's channel, which is given higher precedence if a valid neighbor is found on it during gateway selection.  Gateway Selection Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only. Source: Wireless Mesh Networking



Source: Manually Changing Channels in a Mesh Network

Cisco also sets or recommends the maximum number of mesh hops, the gateway-to-repeater ratio, and link quality.



Deployment Considerations

The guidelines below should be followed to ensure a successful mesh deployment.

Wireless Client Handling on a Repeater

VLANs and other SSID options (such as firewall rules and traffic shaping) will be maintained on SSIDs that are broadcast by repeater access points.

Gateway-to-Repeater Ratio

In general, having as many gateway access points as possible to maximize overall network performance and reliability is desired. It is recommended to plan for no more than two repeater access points attached to each gateway access point.

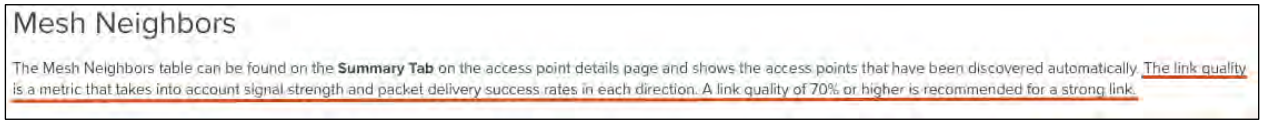
Maximum Mesh Hops

There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device.

Multi-Path Mesh

It is recommended that each mesh access point has at least three strong "neighbors," or other access points in the mesh that the access point can "see," so that multiple failover paths exist. Wireless Mesh Networking gives an overview of the monitoring tools available on dashboard, including information about the mesh neighbors table found on individual access point detail pages.

Source: Mesh Deployment Guide



Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

The Meraki cloud server can further be used to monitor various mesh metrics, including mesh throughput, mesh routes, and mesh neighbors.

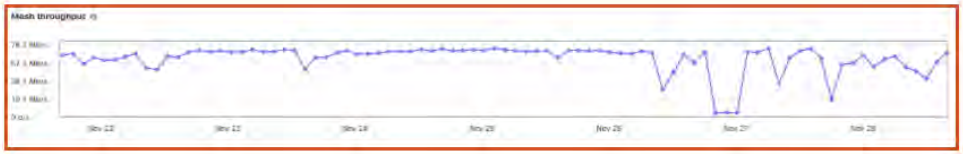
Monitoring Mesh

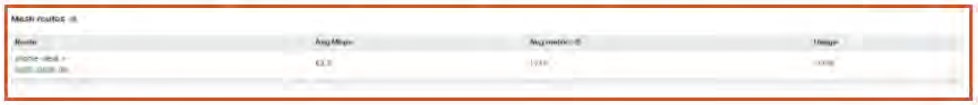
Mesh monitoring tools are located at the bottom of every access point detail page, which can be accessed by navigating to **Wireless > Monitor > Access Points**, then clicking on an Access Point.

The image below shows an example access point acting as a repeater. The time selector at the top right-hand corner will adjust the timeframe of all of the UI components in the mesh monitoring section of the UI.

The time selector may select data from:

- 2 Hours
- 1 day
- 1 week
- 1 month





Mesh Routes

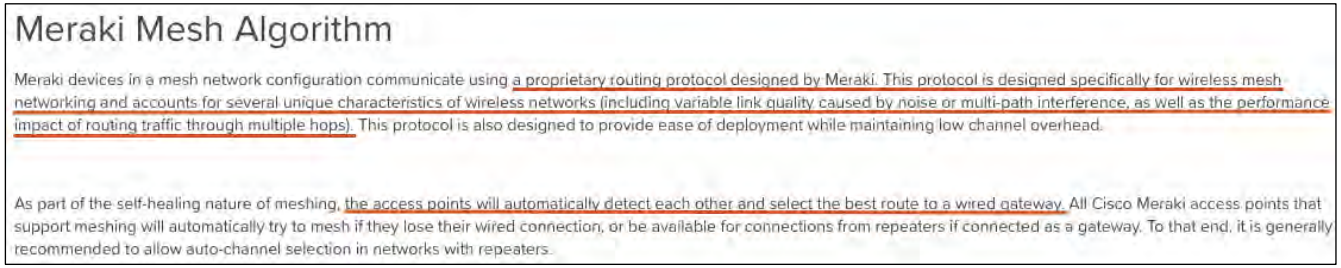
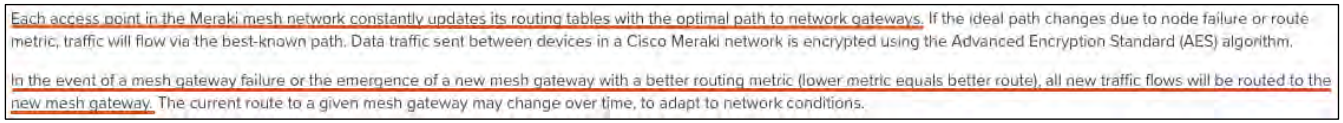
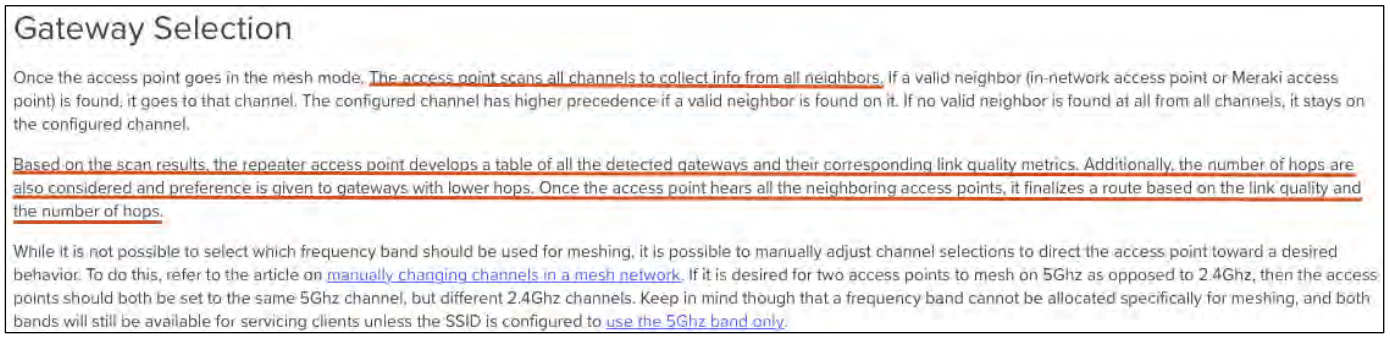
The Routes Table shows the routes used by different flows over time. As new routes are selected, they are added to the routes table. The overall amount of traffic per-route over the time period selected is shown in the Usage column. The Metric is also displayed in this table, representing a combination of loss and packet delivery times. Avg. Mbps throughput values are also provided for customers to gauge the capacity of that particular mesh route.

Mesh Neighbors

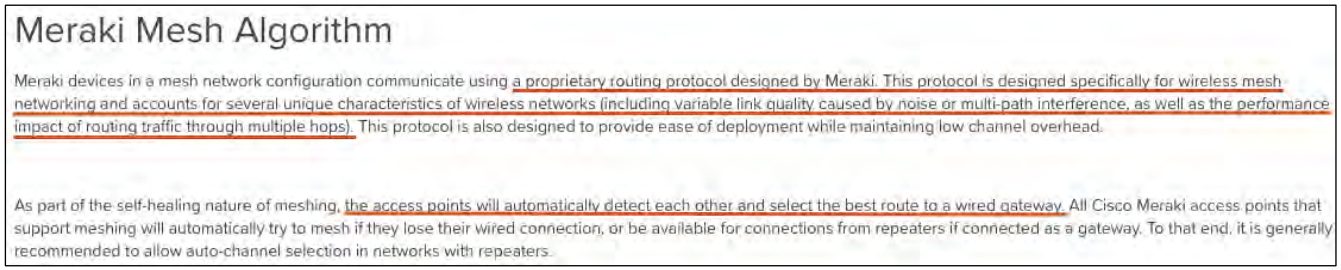
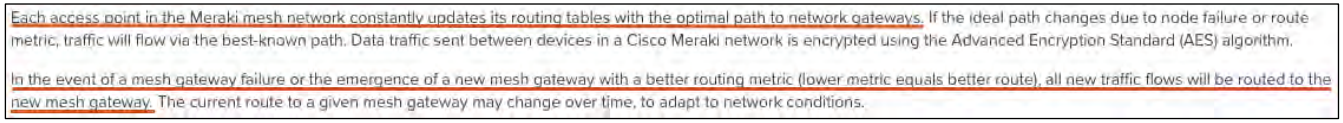
The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

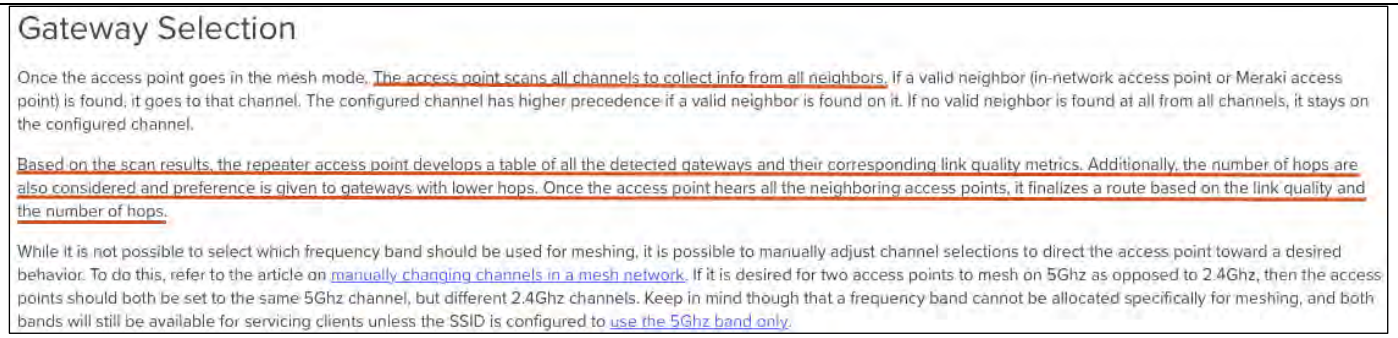
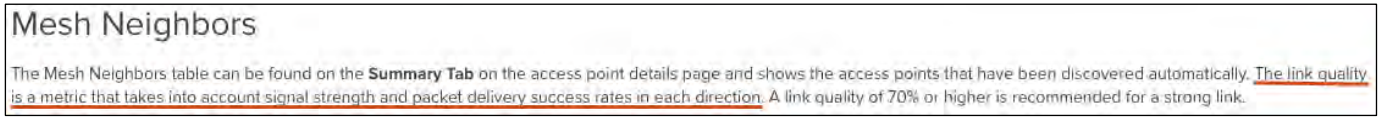
Source: Wireless Mesh Networking

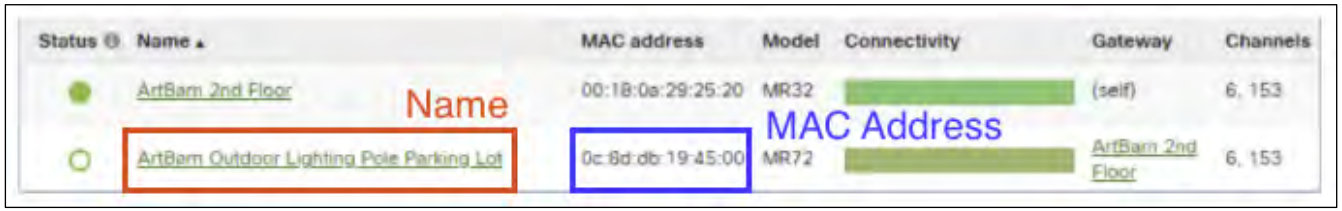
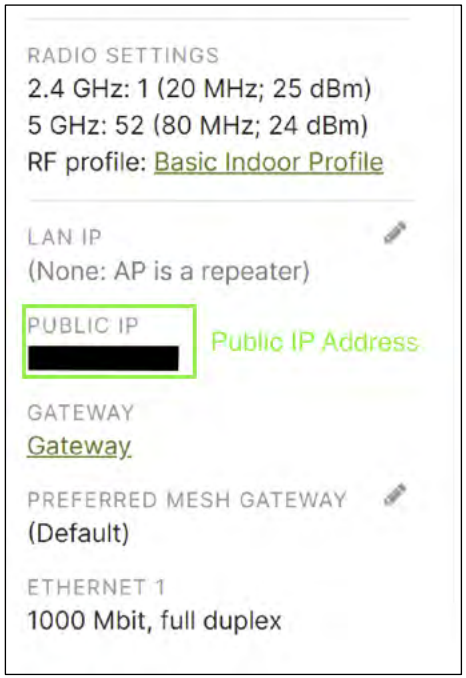
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (establishing wireless data communication between a repeater and the external network) in the same way (by selecting the best neighbor based on certain criteria defined by the server) and achieves the same result (selecting the best neighbor that offers the best route to the external network). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses certain criteria to select a parent that is the

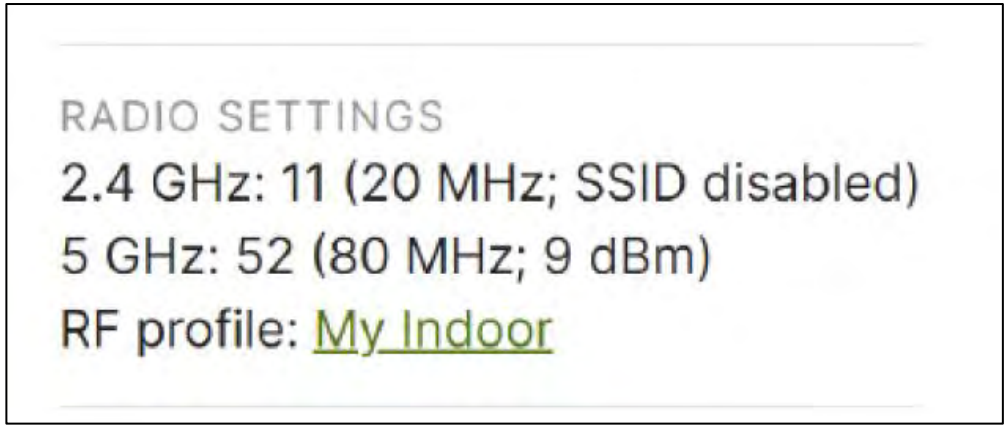
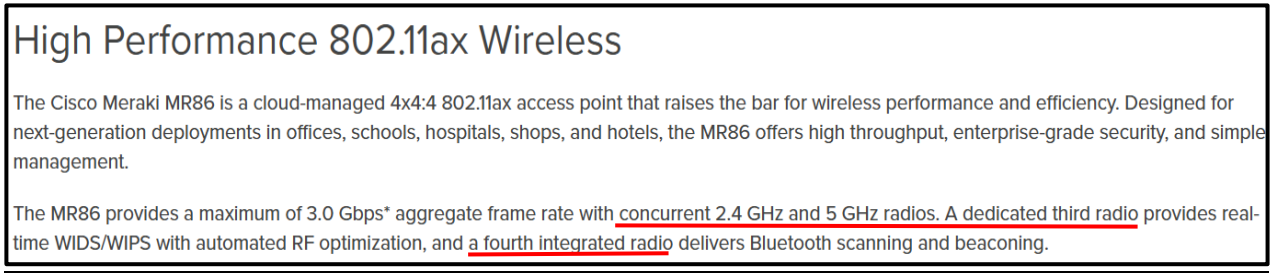
	best route back to the RAP and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.h. establishing a direct wireless data connection between the MAP node and the associated parent node;	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes wherein automatic connection to the associated parent node comprises establishing a direct wireless data connection between the MAP node and the associated parent node. The repeater establishes a direct wireless data connection with the best neighbor selected based on the Meraki Mesh Algorithm.  Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.  Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.  Gateway Selection Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only. Source: Wireless Mesh Networking

	MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.i. wherein the automatic connection process establishes the MAP node's routing path to the external network;	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more mesh access point (MAP) nodes wherein automatic connection process to the associated parent node establishes the MAP node's routing path to the external network. A repeater's selection of the best neighbor establishes the best mesh routing path to the external network. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only. Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (establishing the MAP node's routing path to the external network) in the same way (by selecting a parent node based on certain criteria) and achieves the same result

	(selecting a parent node and associated routing path that offers the best route to the server and external network). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses certain criteria to select a parent that is the best route back to the RAP and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.j. the one or more parent selection criteria comprising: a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path;	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes wherein the one or more parent selection criteria comprises a throughput value representing signal strength between two nodes and a latency value representing number of hops in a routing path. The repeater selects the neighbor node providing the best mesh route to a gateway using the Meraki Mesh Algorithm based on, *inter alia*, link quality metrics and number of hops:  Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.  Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

	 Source: Wireless Mesh Networking Link quality metric takes into account, *inter alia*, signal strength.  Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting the best neighbor and mesh route for a repeater) in the same way (based on certain selection criteria, *e.g.*, signal strength and number of hops) and achieves the same result (selecting a neighbor and associated routing path that offers the best route to the external network). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses certain criteria including signal strength and number of hops in a routing path to select a neighbor that provides the best route back to a gateway and the external network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.k. the MAP node comprising:	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) comprising one or more identifier values representing a unique identity of the MAP node in the wireless mesh network.

one or more identifier values representing a unique identity of the MAP node in the wireless mesh network;	Each repeater in the Cisco's Meraki Wireless Mesh Network comprises a unique name, MAC address, and public IP address that represents the unique identity of the repeater.  Source: Manually Changing Channels in a Mesh Network  Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (identifying a repeater) in the same way (by assigning a unique identifier, *e.g.*, name, MAC, and/or IP address to the repeater) and achieves the same result (data flow amongst a plurality of repeaters). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network uses MAC addresses to identify

	and add access points to the mesh network. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1.l. two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes comprising two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network. Access points in the Cisco Meraki Wireless Mesh Network have multiple radio interfaces, including 2.4 and 5 GHz radios.  Source: Wireless Mesh Networking  Source: MR86 Datasheet

Product options: Wi-Fi 6E indoor models

	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: Meraki Family Datasheet

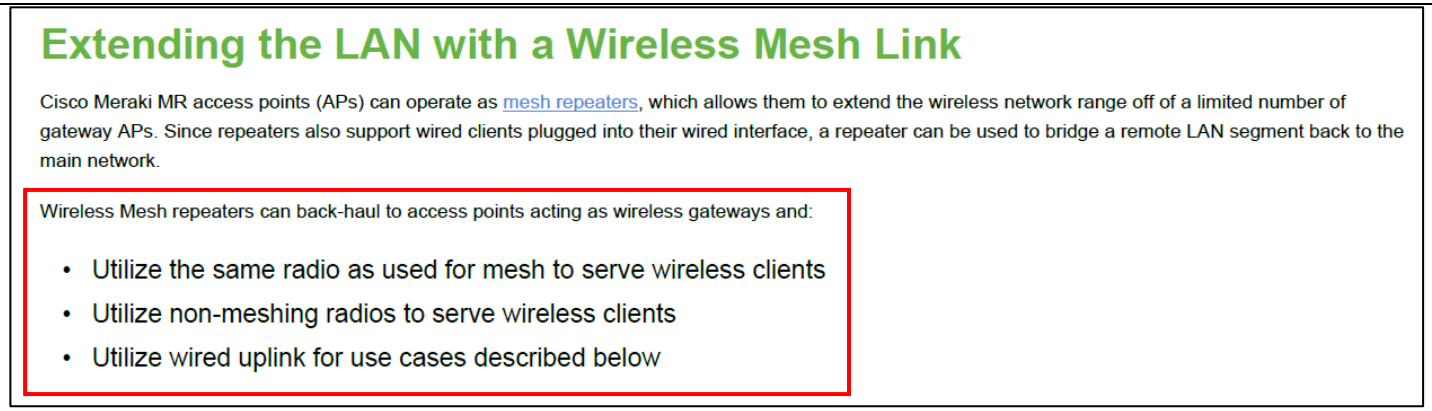
Radios may be used concurrently for meshing while also serving wireless clients (*i.e.*, for wirelessly communicating with another node or a client device in the wireless mesh network).

Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: Mesh Deployment Guide

<table>
<tr><td></td><td> 

Source: Extending the LAN with a Wireless Mesh Link

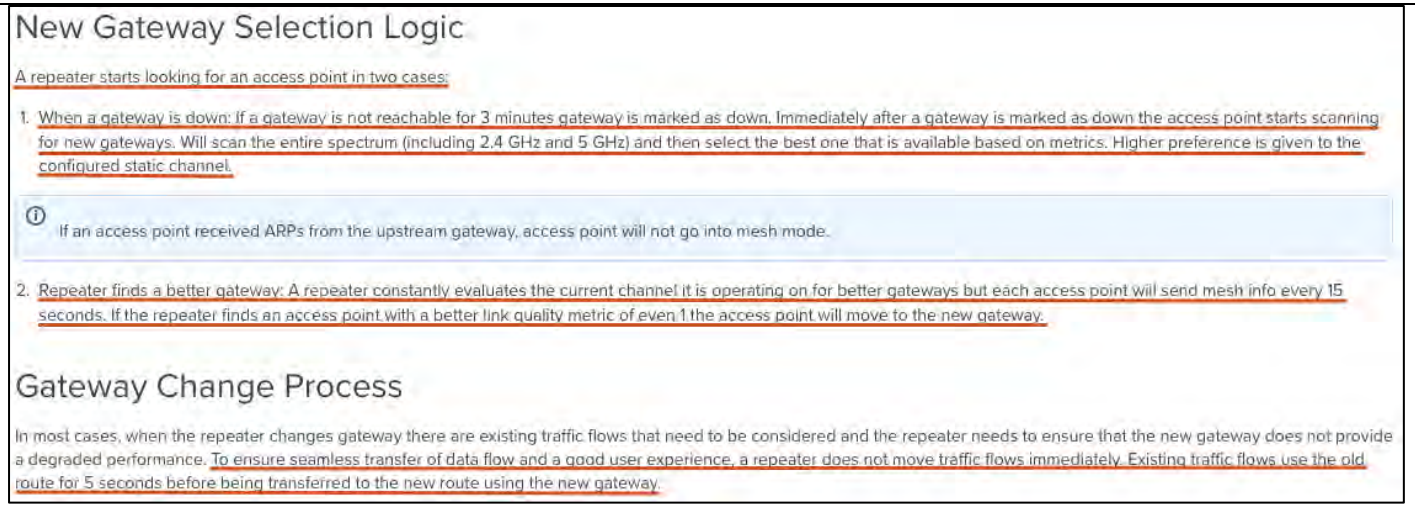
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Meraki and Catalyst CW Access Points (e.g., the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network performs the same function (wireless communication) in the same way (with two or more radios) and achieves the same result (wireless communication in two or more frequency bands and/or with two or more AP/client devices). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's access points use two or more radios to communicate in the mesh network and to communicate with client devices. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.</td></tr>
<tr><th>Claim 2</th><th></th></tr>
<tr><td>2.a. The wireless mesh network of claim 1 further comprising:

one or more network health monitoring criteria defined by the access server, wherein the one or more network health monitoring criteria comprises one or more of:</td><td>The Meraki and Catalyst CW Access Points (e.g., the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more repeater (i.e., MAP) nodes wherein one or more network health monitoring criteria is defined by the access server, e.g., the Meraki cloud server, and comprises one or more of: whether or not the associated parent node is offline or can be found in the wireless mesh network; whether or not data can be routed to the external network through the MAP node's routing path; whether or not the one or more parent selection criteria is satisfied.

As discussed for claim 1 above, Cisco's Meraki Wireless Mesh Networks use the Meraki Mesh Algorithm to configure mesh network communication. The algorithm enables access points to automatically detect each other and select the best route to a wired gateway. The mesh neighbor that provides the best mesh route to a gateway is selected</td></tr>
</table>

whether or not the associated parent node is offline or can be found in the wireless mesh network; whether or not data can be routed to the external network through the MAP node's routing path; whether or not the one or more parent selection criteria is satisfied; and	based on, *inter alia*, link quality metrics and number of hops. The Meraki cloud server provides the framework for access points to communicate and choose optimal mesh route, including defining criteria for selecting the best neighbor and other criteria. *See supra* discussion for claim elements 1.e-1.i. The Meraki Mesh Algorithm further comprises self-healing capability, wherein access points will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. Moreover, the algorithm enables monitoring of a change in the ideal mesh routing path due to, for example, a mesh gateway failure (*i.e.*, whether or not the associated parent node is offline or can be found in the wireless mesh network and/or whether or not data can be routed to the external network through the MAP node's routing path) or the emergence of a new mesh gateway with a better routing metric (*i.e.*, whether or not the one or more parent selection criteria is satisfied). **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking If a change in the ideal mesh routing path is detected, the Algorithm provides a new gateway selection logic for selecting a new mesh route.

	 Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (monitor network health, *e.g.*, connection between repeater and gateway) in the same way (by monitoring whether a gateway is offline) and achieves the same result (roaming to a different routing path). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network monitors the connection between the parent and child and will select a new parent as needed. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.b. testing whether the MAP node's routing path fails to satisfy the one or more network health monitoring criteria;	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes wherein the one or more repeater node tests whether its routing path fails to satisfy the one or more network health monitoring criteria. The Meraki Mesh Algorithm enables repeaters to test whether a gateway is reachable. If a gateway is not reachable for 3 minutes, the repeater starts scanning for new gateways.

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

Source: Wireless Mesh Networking

The Meraki Mesh Algorithm also enables repeaters to test whether there is a better gateway every 15 seconds based on link quality metric and number of hops. If a gateway with better metrics is found, the repeater will move to the new gateway.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

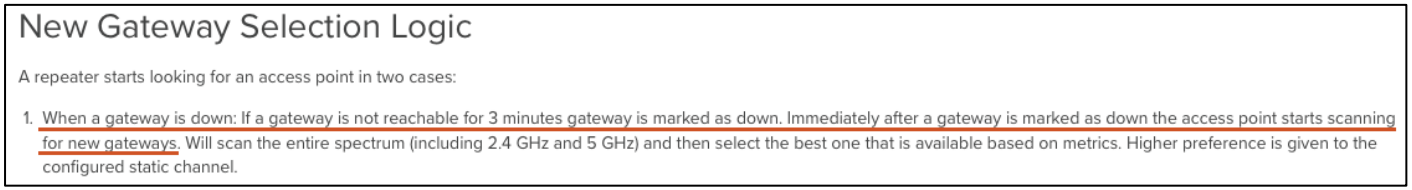
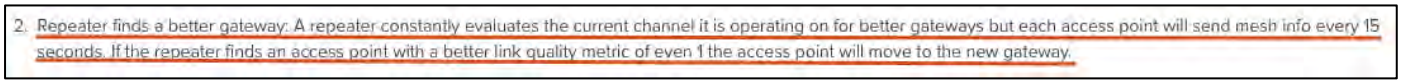
In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

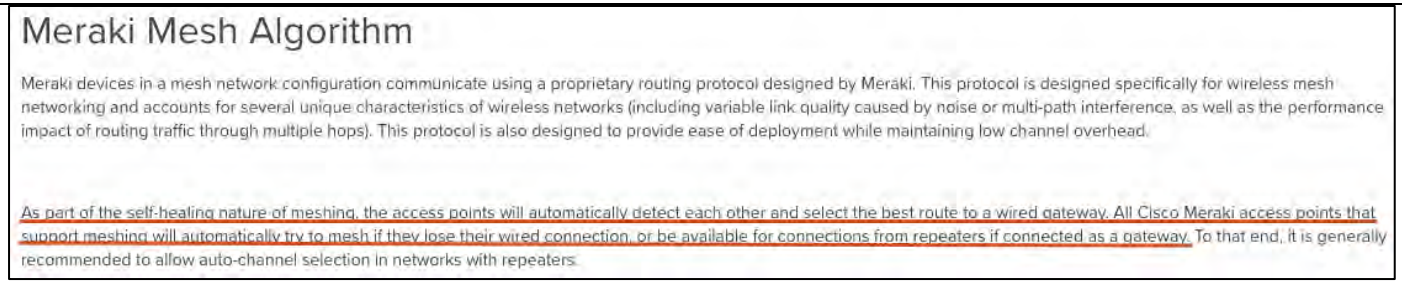
Gateway Selection

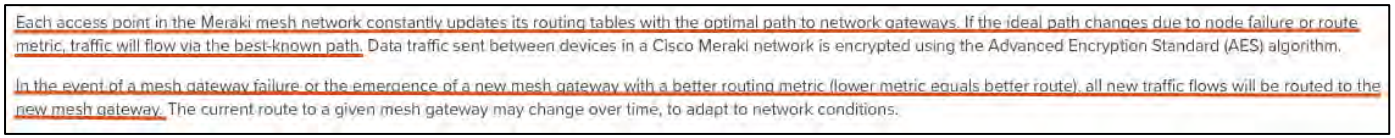
Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

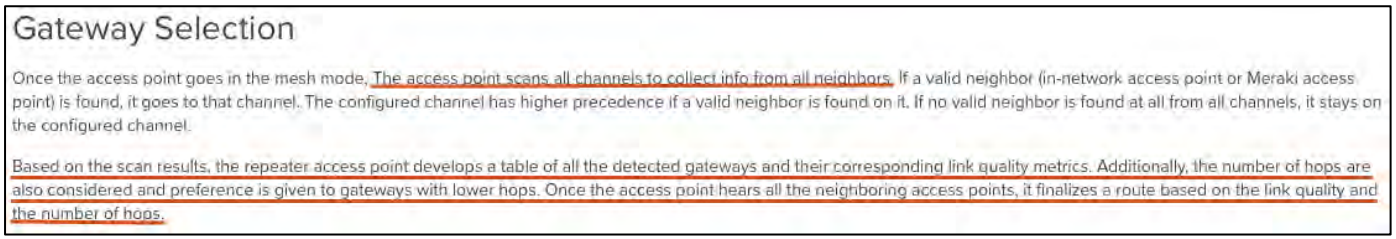
Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (test the network health of a repeater's routing path) in the same way (by monitoring whether a gateway is unreachable and/or whether a better mesh route with better link quality metrics and number of hops exists) and achieves the same result (roaming to a different routing path). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will monitor and test the repeater's routing path back to the gateway. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
2.c. when it is determined that the MAP node's routing path does not meet one or more of the one or more network health monitoring criteria, establishing a new routing path from the MAP node to the external network by connecting the MAP node to a new associated parent node through the automatic connection process.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network are comprised of one or more repeater (*i.e.*, MAP) nodes wherein when it is determined that the MAP node's routing path does not meet one or more of the one or more network health monitoring criteria, establishing a new routing path from the MAP node to the external network by connecting the MAP node to a new associated parent node through the automatic connection process. When a repeater detects that a gateway is down (*e.g.*, if the gateway is not reachable for 3 minutes), the repeater starts scanning for a new gateway. A new gateway and mesh route thereto is selected based on link quality metrics and the number of hops of available neighbors.  New Gateway Selection Logic A repeater starts looking for an access point in two cases: 1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel. Source: Wireless Mesh Networking A repeater further detects whether a better neighbor/gateway with better link quality metrics and number of hops is available every 15 seconds. If a better neighbor/gateway and mesh route is found, the repeater will move to the new gateway by connecting to the associated neighbor.  2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.







Source: Wireless Mesh Networking

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (establishing a new routing path from the repeater to the external network upon detection of a downed/better gateway) in the same way (by connecting the repeater to a new associated neighbor) and achieves the same result (data communication between the repeater and the external network via the new routing path). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will monitor and test the MAPs routing path back to the RAP and automatically connect the MAP to the best path. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 3	
3. The wireless mesh network of claim 2 wherein	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 2 as shown above, comprises causing an update to the

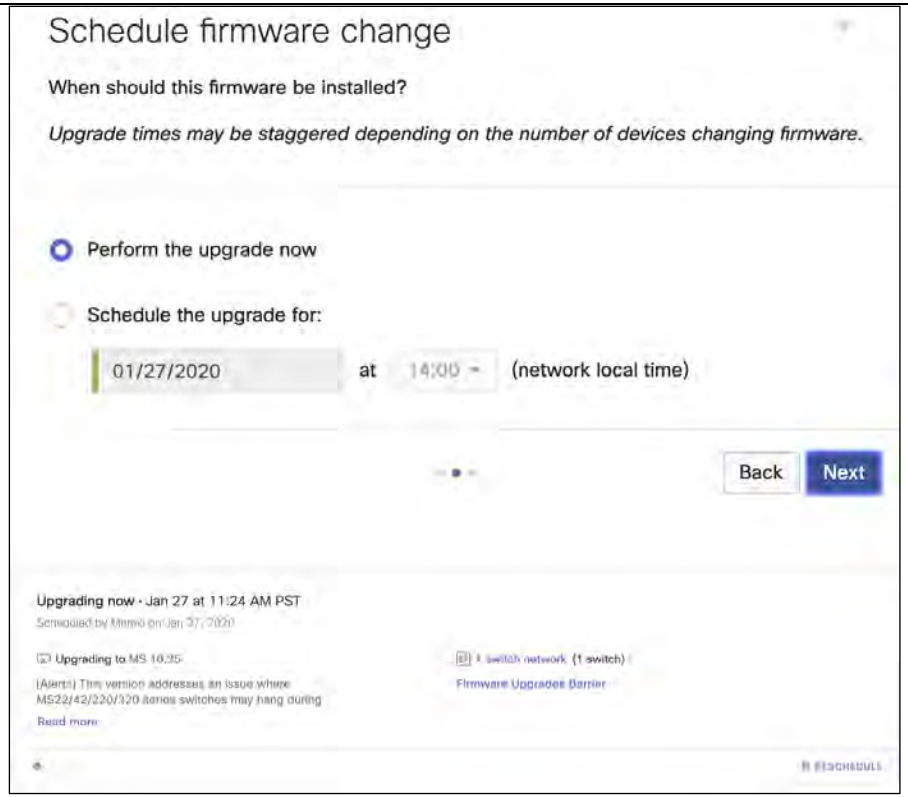
the access server further comprises causing an update to the wireless mesh network, the update comprising one or more of: a change to the one or more parent selection criteria, a change to the one or more network health monitoring criteria, and an upgrade of software or firmware residing on at least some of the nodes.	wireless mesh network, the update comprising one or more of: a change to the one or more parent selection criteria, a change to the one or more network health monitoring criteria, and an upgrade of software or firmware residing on at least some of the nodes. The Meraki Cloud Server (*i.e.*, access server) provides dashboard tool for firmware upgrades to the access points in the Cisco Meraki Wireless Mesh Network. **Overview** Firmware upgrades allow network administrators to utilize the latest features and security enhancements on their Meraki devices. The Cisco Meraki dashboard allows admins to easily schedule and reschedule firmware upgrades on their networks, opt in to beta firmware releases, view firmware change log notes, and set maintenance windows. This article outlines the functions of the firmware upgrades tool in the dashboard, as well as additional considerations for using and managing firmware. **Managing Firmware as an Organization Admin** The firmware upgrades tool in the dashboard allows organization admins to quickly and easily manage firmware versions on a per-network and per-device type basis. Additionally, the firmware upgrades tool can be used to schedule, reschedule, and cancel bulk upgrades of networks, view firmware change log notes, view firmware version numbers, and roll back the firmware on a recently upgraded network. **Scheduling Firmware Upgrades** Keeping up-to-date on firmware allows administrators to utilize the latest features and ensures that the latest security enhancements are running on their hardware. Admins can upgrade to the latest stable or beta firmware. Follow the steps below to schedule a firmware upgrade. 1. Navigate to **Organization > Monitor > Firmware upgrades**. 2. Click the **Schedule Upgrades** tab in the upper-left. 3. Select the devices or networks to be upgraded by clicking the checkboxes beside the network names. Admins can specify upgrades on a per-network or per-device type basis by using the **Device type, Current version,** and/or **Status** drop-down selectors. 4. Click the **Schedule upgrades** button. 5. Select the firmware version for upgrade using the **Target firmware version** selector. 6. Select either **Perform the upgrade now** or **Schedule the upgrade for**, specifying a specific date and time for the upgrade. 7. Review the **Change Summary** and select **Schedule change for network**.

1. Select **Organization > Monitor > Firmware upgrades > Schedule Upgrades**. From here you can select an individual or many networks within your organization



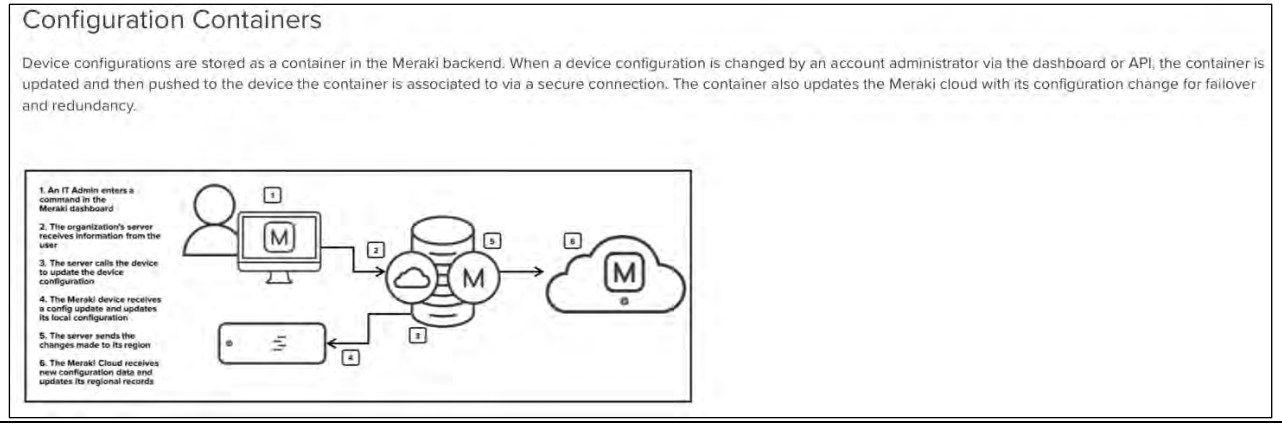
2. Select the **Target firmware version** you are trying to upgrade to from the drop-down menu. In this case, because we are upgrading from MS 9.27 we must upgrade to MS 10.35 first. Note that the latest stable is not a selectable option and won't be until you upgrade to the required intermediary firmware version. This is a prime example of a firmware barrier



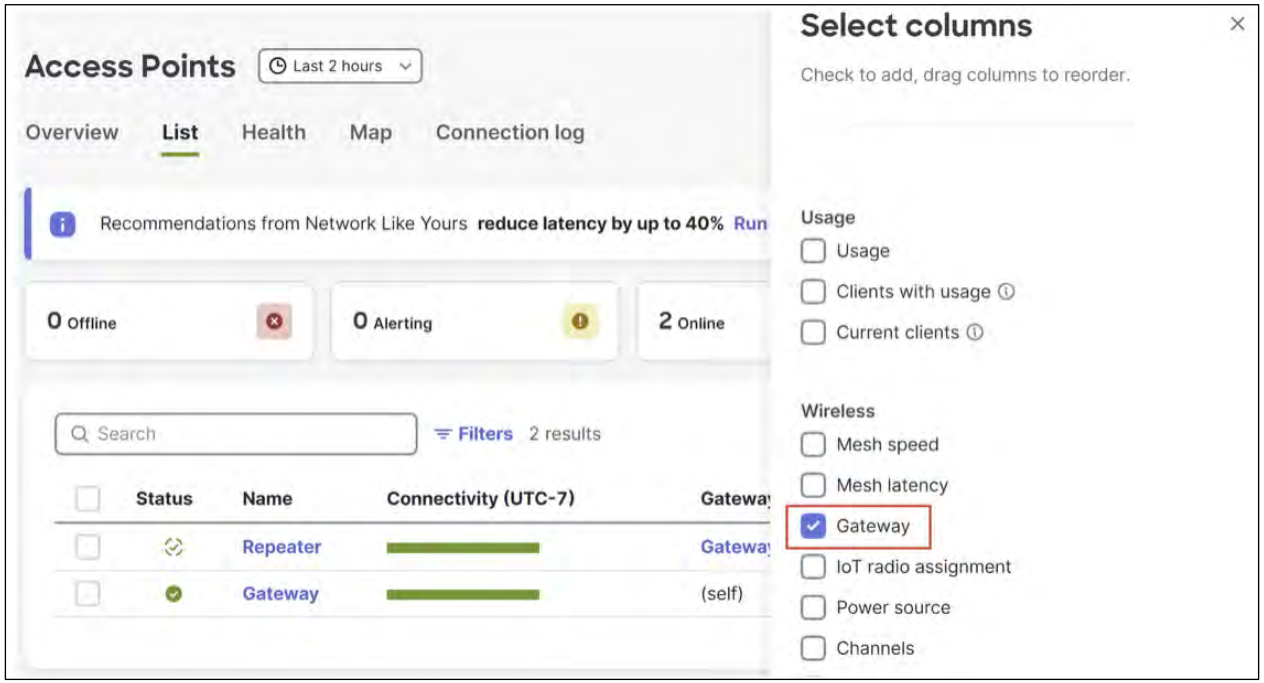


Source: Managing Firmware Upgrades

Moreover, the Meraki Cloud Server can be used to update access point configurations.



Source: Meraki Cloud Architecture

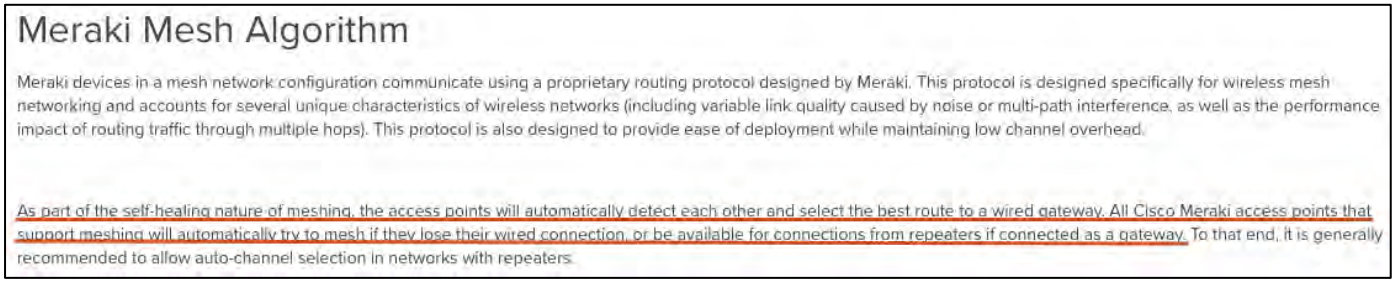
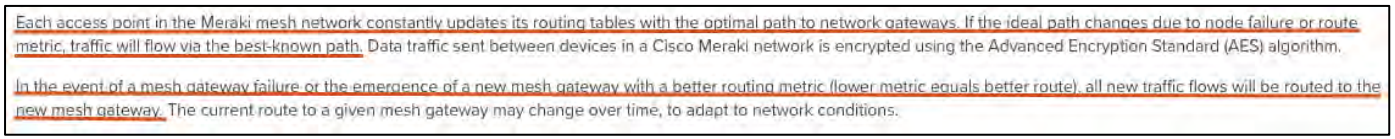


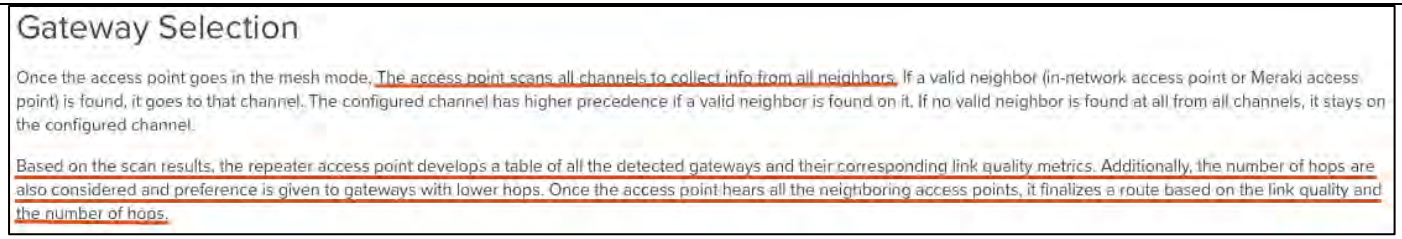


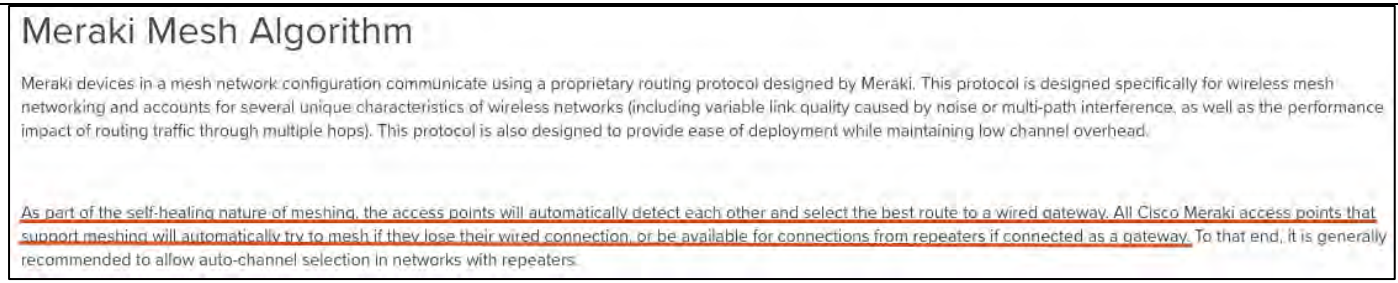
Source: Wireless Mesh Networking

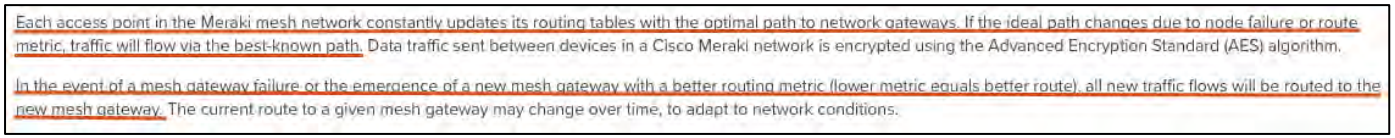
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. The Meraki Cloud Server performs the same function (updating the wireless mesh network) in the same way (by rolling out the update using the Meraki Cloud Server) and achieves the same result (wireless mesh network with updated settings, software, or firmware). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will provide updates, such as images, to the access points. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

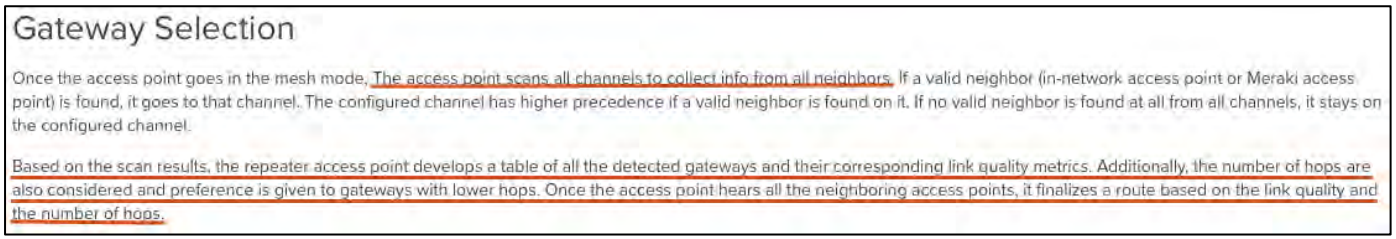
MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to..

Claim 4	
4. The wireless mesh network of claim 2 wherein the at least one of the one or more network health monitoring criteria overlap with at least one of the one or more parent selection criteria.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 2 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the at least one of the one or more network health monitoring criteria overlap with at least one of the one or more parent selection criteria. As discussed for claim 2 above, a repeater monitors all of its neighbors' link quality metrics and number of hops to determine whether they are better than its associated gateway and whether to switch to a better neighbor. *See supra* discussion for claim 2. The neighbors' link quality metrics and number of hops are criteria that overlaps with those for gateway selection as discussed for claim 1. *See supra* discussion for claim 1. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 5	
5. The wireless mesh network of claim 1 wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their latency values.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their latency values. As discussed for claim 1 above, Cisco's Meraki Wireless Mesh Network use the Meraki Mesh Algorithm to automatically detect and select the best mesh route to a gateway. The algorithm does so by comparing, *inter alia*, the number of hops (*i.e.*, latency value) for each mesh route and preference is given to gateways with lower hops.  Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.  Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.

	 Source: Wireless Mesh Networking If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (using criteria such as number of hops) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will select the parent node based at least in part on a comparison of their latency values. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 6	
6. The wireless mesh network of claim 5 wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a lower latency value than the another potential parent node.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 5 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a lower latency value than the another potential parent node. As discussed for claims 1 and 5 above, Cisco's Meraki Wireless Mesh Network use the Meraki Mesh Algorithm to automatically detect and select the best mesh route to a gateway. The Algorithm does so by comparing, *inter alia*, the number of hops (*i.e.*, latency value) for each mesh route and preference is given to gateways with lower hops (*i.e.*, lower latency value).







Source: Wireless Mesh Networking

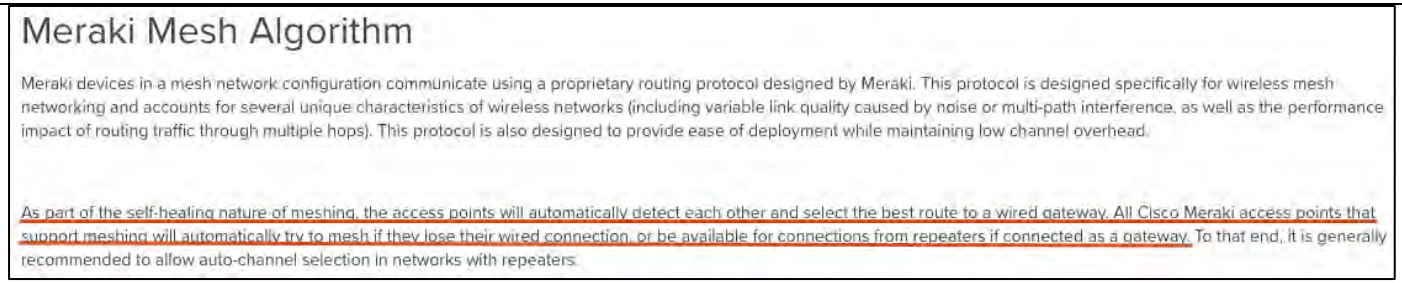
If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (using criteria such as number of hops) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will select the parent node based at least in part on a comparison of their latency values. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 7	
7. The wireless mesh network of claim 1 wherein the selection of the parent	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh

node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their throughput values.	access point (MAP) nodes wherein the selection of the parent node further comprises selecting a potential parent node over another potential parent node based at least in part on a comparison of their throughput values. As discussed for claim 1 above, Cisco's Meraki Wireless Mesh Network use the Meraki Mesh Algorithm to automatically detect and select the best mesh route to a gateway. The algorithm does so by comparing, *inter alia*, the link quality metrics and number of hops for each mesh route and preference is given to gateways with higher link quality and lower hops. **Meraki Mesh Algorithm** Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking Link quality metrics accounts for signal strength. Thus, comparing mesh routes' link quality metrics is a comparison of throughput values. **Mesh Neighbors** The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

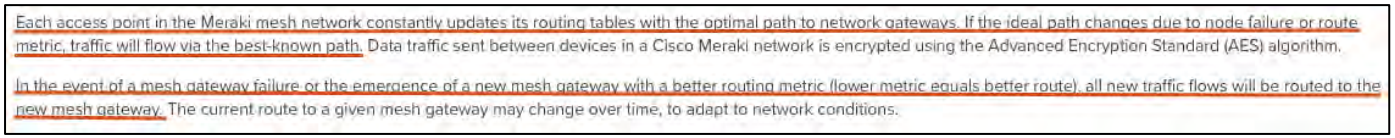
	Source: Wireless Mesh Networking Each hop reduces throughput by approximately 50%. Thus, comparing mesh routes' number of hops is also a comparison of throughput values. Maximum Mesh Hops There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device. Source: Mesh Deployment Guide If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (using criteria such as throughput or signal strength or number of hops) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will select the parent node based at least in part on a comparison of their throughput values. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 8	
8. The wireless mesh network of claim 7 wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a higher throughput value than the another potential parent node.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 7 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises selecting the potential parent node over the another potential parent node if the potential parent node has a higher throughput value than the another potential parent node. As discussed for claims 1 and 7 above, Cisco's Meraki Wireless Mesh Network use the Meraki Mesh Algorithm to automatically detect and select the best mesh route to a gateway. The algorithm does so by comparing, *inter alia*, the link quality metrics (which accounts for signal strength) and number of hops (which impacts throughput) for each mesh route and preference is given to gateways with higher link quality and lower hops.



Meraki Mesh Algorithm

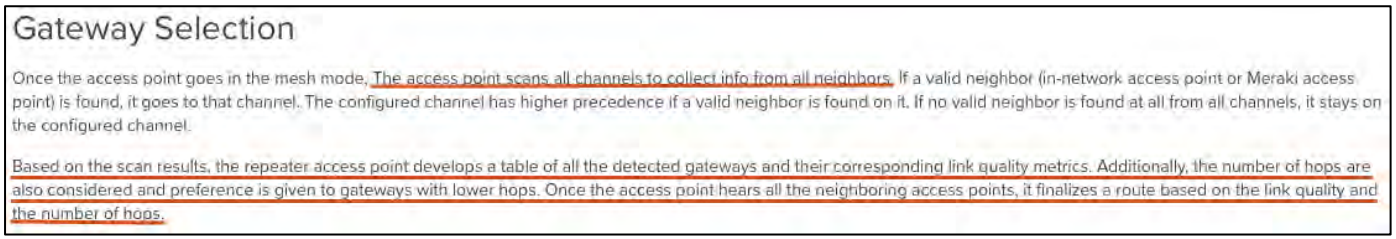
Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.



Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm.

In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions.



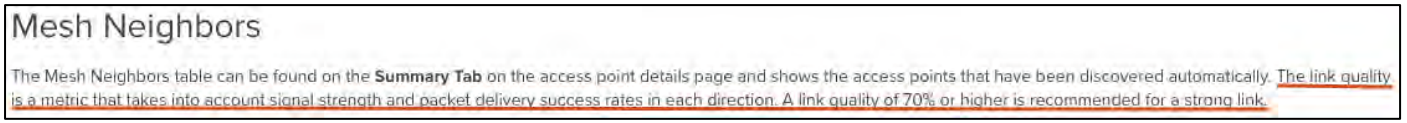
Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

Source: Wireless Mesh Networking

Link quality metrics accounts for signal strength. Thus, comparing mesh routes' link quality metrics and giving preference to mesh routes with higher link quality metrics is selecting a mesh route to a gateway that has a higher throughput value.

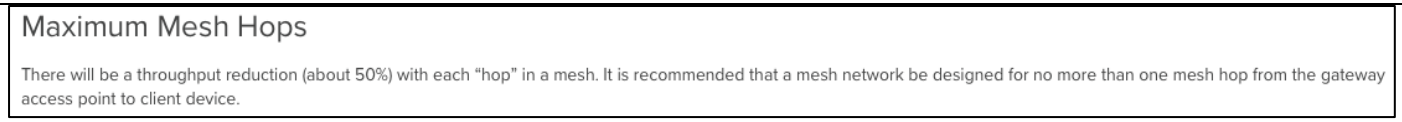
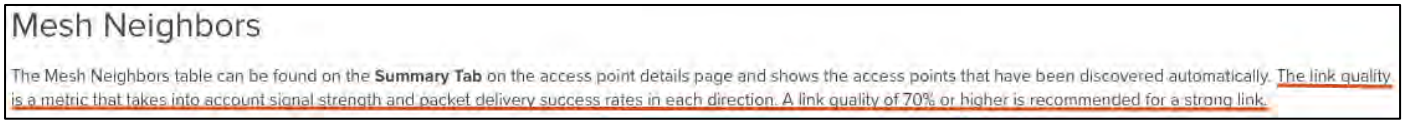


Mesh Neighbors

The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link.

Source: Wireless Mesh Networking

Each hop reduces throughput by approximately 50%. Thus, comparing mesh routes' number of hops and giving preference to mesh routes with lower hops is also selecting a mesh route to a gateway that has a higher throughput value.

	 Source: Mesh Deployment Guide If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (using criteria such as throughput or signal strength or number of hops) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will select the parent node based at least in part on a comparison of their throughput values. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor. MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
Claim 9	
9. The wireless mesh network of claim 1 wherein the selection of the parent node further comprises not selecting a potential parent node if the potential parent node's wireless signal strength with the MAP node does not meet a threshold throughput value defined by the access server.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises not selecting a potential parent node if the potential parent node's wireless signal strength with the MAP node does not meet a threshold throughput value defined by the access server. As discussed for claims 1, 7 and 8 above, Cisco's Meraki Wireless Mesh Network use the Meraki Mesh Algorithm to automatically detect and select the best mesh route to a gateway. The Algorithm does so by comparing, *inter alia*, the link quality metrics (which accounts for signal strength) and number of hops (which impacts throughput) for each mesh route and preference is given to gateways with higher link quality and lower hops. Cisco recommends a link quality of 70% and no more than one mesh hop from the gateway to the client device.  Source: Wireless Mesh Networking

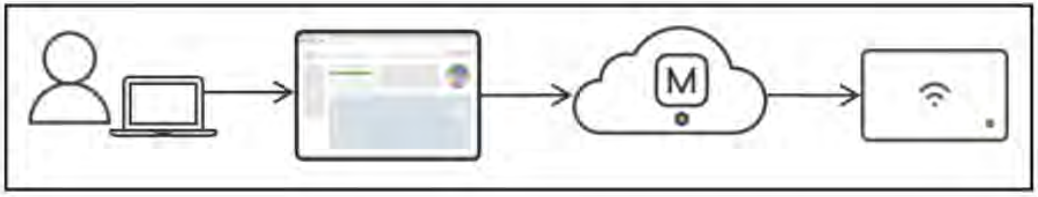
Maximum Mesh Hops

There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device.

Source: Mesh Deployment Guide

The Meraki Cloud Server can be used to manage and configure the access points. On information and belief, the Meraki Cloud Server can be used to configure the minimum link quality and/or the maximum number of hops for a mesh route to be selected by a repeater.

The Meraki cloud solution is a centralized management service that allows users to manage all of their Meraki network devices via a single, simple and secure platform.



Users are able to deploy, monitor, and configure their Meraki devices via the Meraki dashboard web interface or via APIs. Once a user makes a configuration change, the change request is sent to the Meraki cloud and is then pushed to the relevant device(s).

Management data: The data (configuration, statistics, monitoring, etc.) that flows from Meraki devices (wireless access points, switches, security appliances) to the Meraki cloud over a secure internet connection.

Management Data

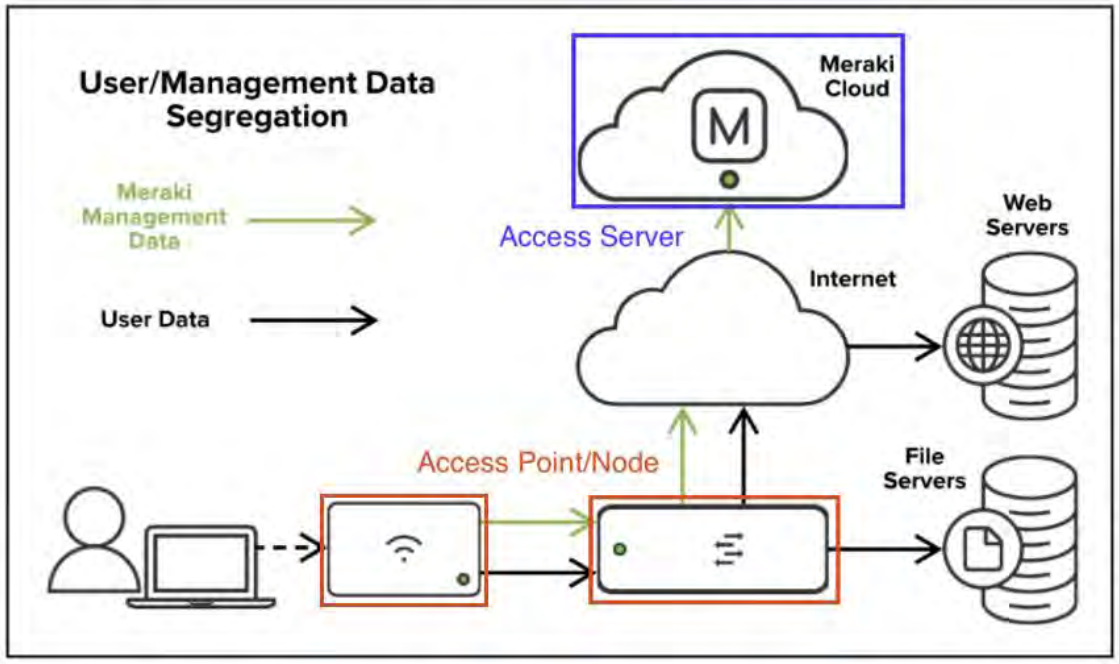
The Meraki cloud gathers and stores certain types of "management" data to enable its solutions. All forms of data are encrypted in transit to and from Meraki servers. There are four major types of data stored in the Meraki cloud:

User records: Includes account email and company name or other optional information such as user name and address.

Configuration data: Includes network settings and configurations made by customers in the Meraki dashboard.

Analytics data: Includes client, traffic, and location analytics data, providing visualizations and network insights into traffic patterns across customer sites.

Customer-uploaded assets: Includes custom floor plans and splash logos.

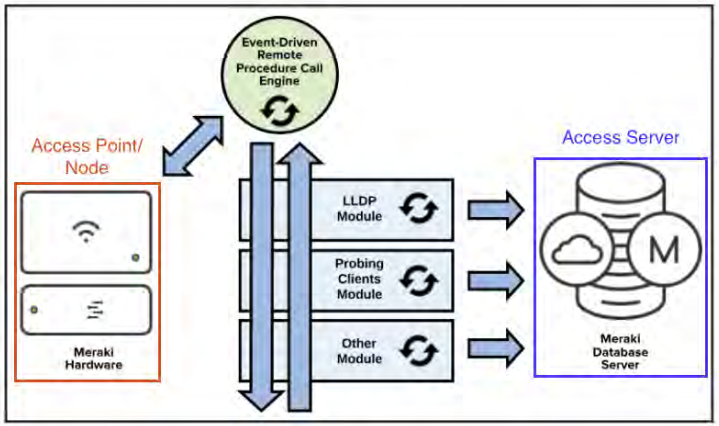


Meraki Cloud Architecture

The Meraki cloud is the backbone of the Meraki management solution. This "cloud" is a collection of highly reliable multi-tenant servers strategically distributed around the world at Meraki data centers and select public cloud service providers. For the remainder of this document we will refer to both the Meraki data centers and the public cloud service providers' locations we use as "data centers". The servers at these data centers are powerful hosting computers comprised of many separate user accounts. They are called multi-tenant servers because the accounts share (equal) computing resources on their host (the server). However, even though these accounts share resources, Meraki ensures that customer information is kept secure by restricting organization access based on account authentication, as well as hashing authentication information such as user passwords or API keys.

Meraki Device-to-Cloud Communications

Meraki uses an event-driven remote procedure call (RPC) engine for Meraki devices to communicate to the dashboard and for Meraki servers to send and receive data. Meraki hardware devices act as the server/receiver as the Meraki cloud initiates calls to the devices for data collection and configuration deployment. The cloud infrastructure is the initiator, so configurations can be executed in the cloud before the devices are actually online or even physically deployed.



Communication Process

If a device is offline, it will continue to attempt to connect to the Meraki cloud until it gains connectivity. Once the device comes online, it automatically receives the most recent configuration settings from the Meraki cloud. If changes are made to the device configuration while the device is online, the device receives and updates these changes automatically. These changes are generally available on the device in a matter of seconds. However, large quantities of changes may take noticeably longer to reach their devices. If no configuration changes are made by the user, the device continues to periodically check for updates to its configuration on its own.

As the device runs on the network, it will communicate device and network usage analytics back to the Meraki cloud. Dashboard analytics based on this information, in the form of graphs and charts, are updated regularly in the Meraki cloud and are displayed in the dashboard of users when they are viewing this information.

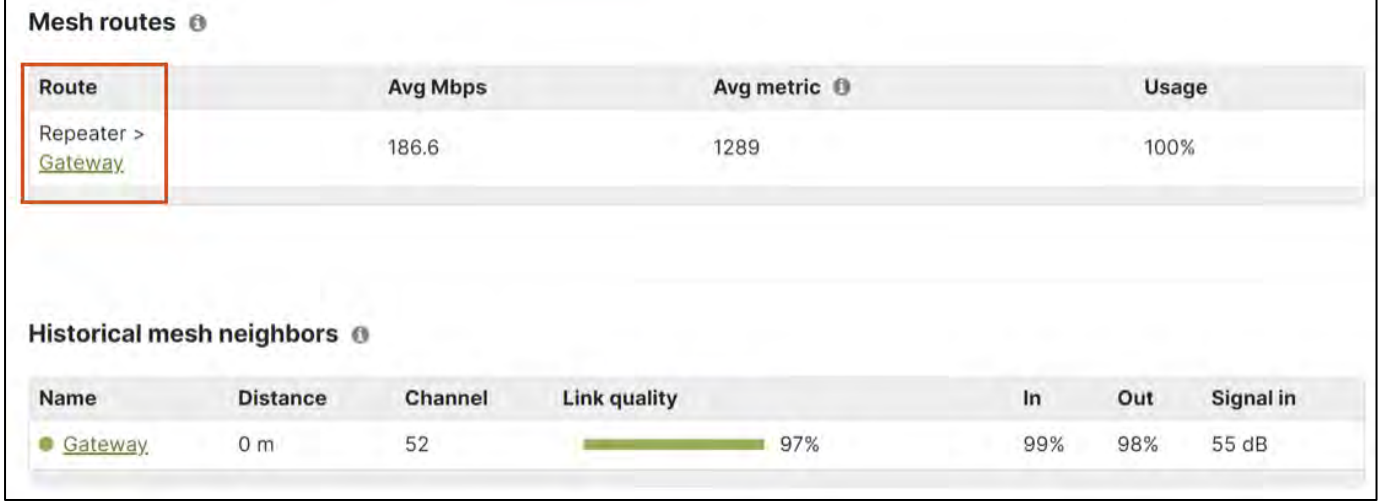
Source: Meraki Cloud Architecture

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (using criteria such as throughput or signal strength) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will select the parent node based at least in part on its throughput value. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 10	
10. The wireless mesh network of claim 1 wherein the MAP node further comprises a data store having data representing the MAP node's routing path to the external network.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) node which comprises a data store having data representing the MAP node's routing path to the external network. The access points in Cisco's Meraki Wireless Mesh Network develop and maintain a routing table comprising the optimal path to neighboring gateways for selecting the best mesh route. Each access point in the Meraki mesh network constantly updates its routing tables with the optimal path to network gateways. If the ideal path changes due to node failure or route metric, traffic will flow via the best-known path. Data traffic sent between devices in a Cisco Meraki network is encrypted using the Advanced Encryption Standard (AES) algorithm. In the event of a mesh gateway failure or the emergence of a new mesh gateway with a better routing metric (lower metric equals better route), all new traffic flows will be routed to the new mesh gateway. The current route to a given mesh gateway may change over time, to adapt to network conditions. Gateway Selection Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. Source: Wireless Mesh Networking Moreover, the Meraki Cloud Server provides live and historical mesh route data (*i.e.*, data representing the MAP node's routing path to the external network) for each access points in the Cisco Meraki Wireless Mesh Network.

7. A repeater's details page also shows information about its route to the mesh gateway. This information is displayed on the **Summary** tab under the **Live Data** and **Historical device data**:



Mesh routes

Route	Avg Mbps	Avg metric	Usage
Repeater > Gateway	186.6	1289	100%

Historical mesh neighbors

Name	Distance	Channel	Link quality	In	Out	Signal in
Gateway	0 m	52	97%	99%	98%	55 dB

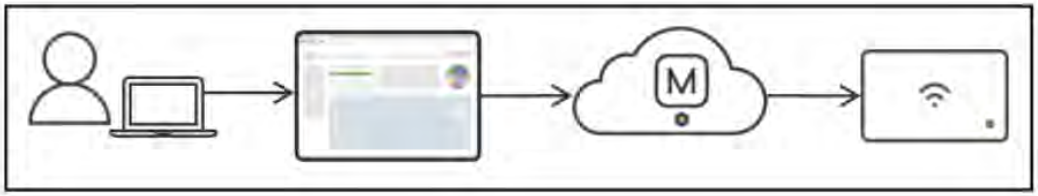
Mesh Routes

The Routes Table shows the routes used by different flows over time. As new routes are selected, they are added to the routes table. The overall amount of traffic per-route over the time period selected is shown in the Usage column. The Metric is also displayed in this table, representing a combination of loss and packet delivery times. Avg. Mbps throughput values are also provided for customers to gauge the capacity of that particular mesh route.

Source: Wireless Mesh Networking

The access points necessarily include a data store to store the routing table and/or mesh route data to be transmitted to the Meraki Cloud Server.

The Meraki cloud solution is a centralized management service that allows users to manage all of their Meraki network devices via a single, simple and secure platform.



Users are able to deploy, monitor, and configure their Meraki devices via the Meraki dashboard web interface or via APIs. Once a user makes a configuration change, the change request is sent to the Meraki cloud and is then pushed to the relevant device(s).

Management data: The data (configuration, statistics, monitoring, etc.) that flows from Meraki devices (wireless access points, switches, security appliances) to the Meraki cloud over a secure internet connection.

Management Data

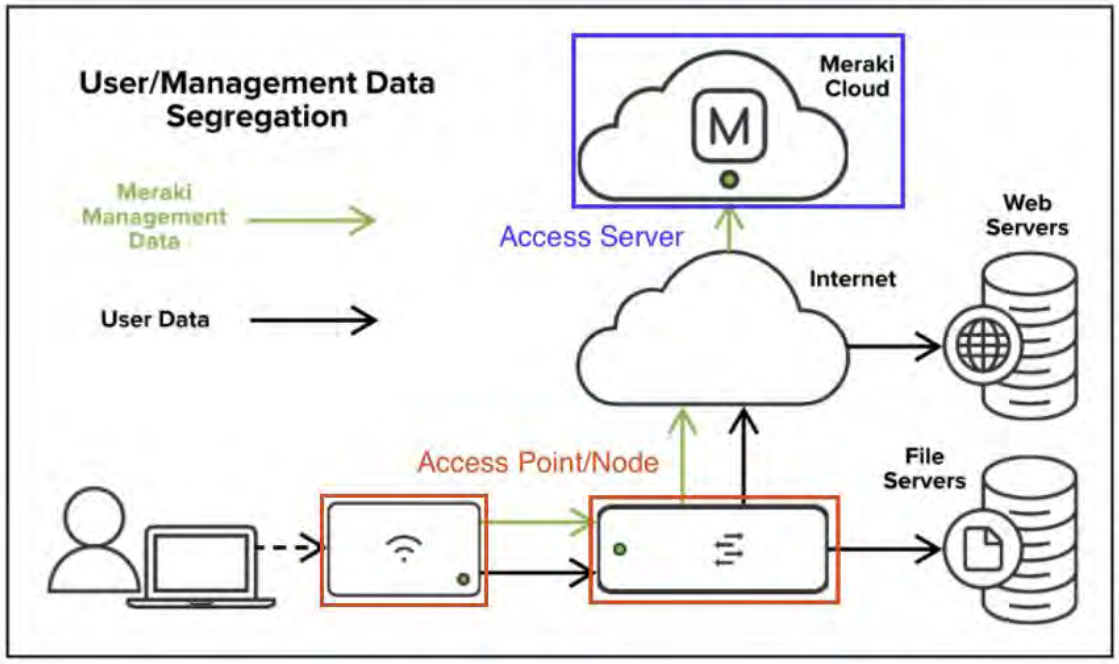
The Meraki cloud gathers and stores certain types of "management" data to enable its solutions. All forms of data are encrypted in transit to and from Meraki servers. There are four major types of data stored in the Meraki cloud:

User records: Includes account email and company name or other optional information such as user name and address.

Configuration data: Includes network settings and configurations made by customers in the Meraki dashboard.

Analytics data: Includes client, traffic, and location analytics data, providing visualizations and network insights into traffic patterns across customer sites.

Customer-uploaded assets: Includes custom floor plans and splash logos.

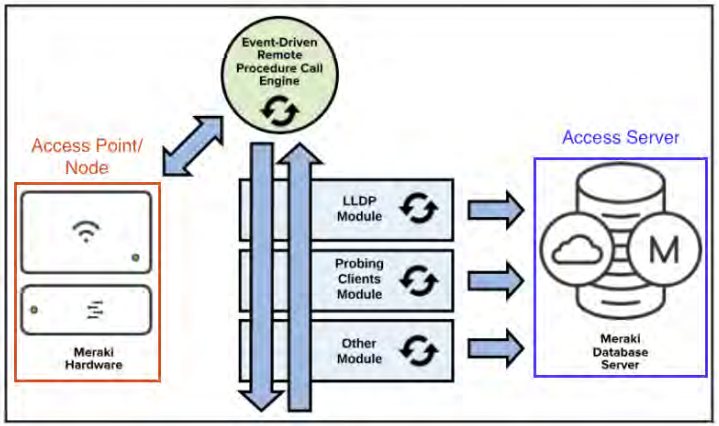


Meraki Cloud Architecture

The Meraki cloud is the backbone of the Meraki management solution. This "cloud" is a collection of highly reliable multi-tenant servers strategically distributed around the world at Meraki data centers and select public cloud service providers. For the remainder of this document we will refer to both the Meraki data centers and the public cloud service providers' locations we use as "data centers". The servers at these data centers are powerful hosting computers comprised of many separate user accounts. They are called multi-tenant servers because the accounts share (equal) computing resources on their host (the server). However, even though these accounts share resources, Meraki ensures that customer information is kept secure by restricting organization access based on account authentication, as well as hashing authentication information such as user passwords or API keys.

Meraki Device-to-Cloud Communications

Meraki uses an event-driven remote procedure call (RPC) engine for Meraki devices to communicate to the dashboard and for Meraki servers to send and receive data. Meraki hardware devices act as the server/receiver as the Meraki cloud initiates calls to the devices for data collection and configuration deployment. The cloud infrastructure is the initiator, so configurations can be executed in the cloud before the devices are actually online or even physically deployed.



Communication Process

If a device is offline, it will continue to attempt to connect to the Meraki cloud until it gains connectivity. Once the device comes online, it automatically receives the most recent configuration settings from the Meraki cloud. If changes are made to the device configuration while the device is online, the device receives and updates these changes automatically. These changes are generally available on the device in a matter of seconds. However, large quantities of changes may take noticeably longer to reach their devices. If no configuration changes are made by the user, the device continues to periodically check for updates to its configuration on its own.

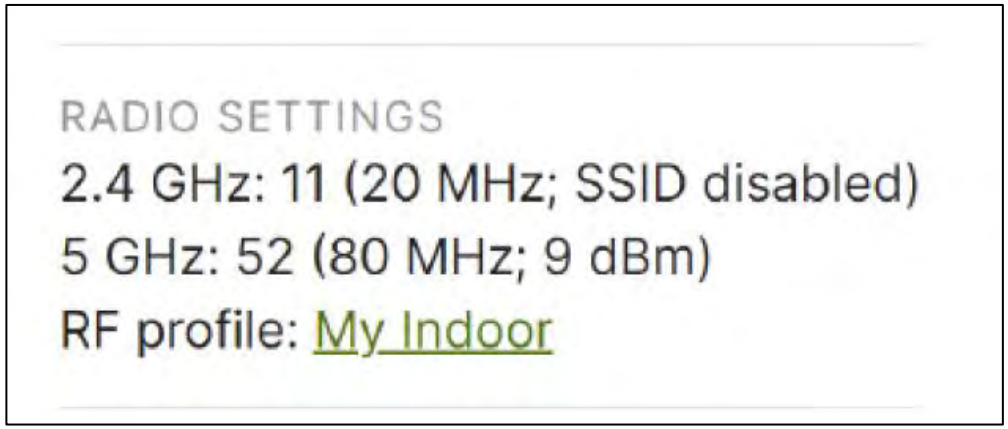
As the device runs on the network, it will communicate device and network usage analytics back to the Meraki cloud. Dashboard analytics based on this information, in the form of graphs and charts, are updated regularly in the Meraki cloud and are displayed in the dashboard of users when they are viewing this information.

Source: Meraki Cloud Architecture

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (maintaining a list of neighbors and routing information) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network maintains routing information and a list of neighbors to use to determine the best route to a gateway. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 11	
11. The wireless mesh network of claim 1 wherein the MAP node further comprises: an ethernet port through which a direct wired connection can be established with another device.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) node which comprises an ethernet port through which a direct wired connection can be established with another device: **High availability/resilience** The flagship MR57 features dual configurable Ethernet ports that can support 1, 2.5, or 5 Gbps, with link aggregation between both network ports for redundancy and high availability.  MR57 **MR28** **MR36** **Usage** Entry-level high-peformance Wi-Fi 6 AP designed for basic, medium-density deployments with Bluetooth® Low Energy included. / Next-generation entry-level high-performance Wi-Fi 6 AP with security-scanning radio and Bluetooth® Low Energy included **Radio specification** 1× 2.4 GHz 802.11b/g/n/ax, 1× 5 GHz 802.11a/n/ac/ax, 1× Bluetooth® radio, 1.5 Gbps max rate, 2× 2:2 MU-MIMO with beamforming / 1× 2.4 GHz 802.11b/g/n/ax, 1× 5 GHz 802.11a/n/ac/ax, 1× WIDS/WIPS, 1× Bluetooth® radio, 1.5 Gbps max rate, 2× 2:2 MU-MIMO with beamforming **Interface** 1× Gigabit Ethernet port / 1× Gigabit Ethernet port Source: Meraki Family Datasheet

Claim 12	
12. The wireless mesh network of claim 1 wherein the one or more identifier values of the MAP comprises a MAC address or an IP address.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) wherein the one or more identifier values of the MAP comprises a MAC address or an IP address. As discussed and shown above for claim element 1.j, each repeater (*i.e.*, MAP) in the Cisco Meraki Wireless Mesh Network comprises a unique MAC address and an IP address. *See supra* discussion for claim element 1.j.
Claim 13	
13. The wireless mesh network of claim 1 wherein a first radio of the MAP is adapted for communicating with the client device and wherein a second radio is adapted for communicating with another node.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, are comprised of one or more mesh access point (MAP) nodes comprising two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network, wherein a first radio of the MAP is adapted for communicating with the client device and wherein a second radio is adapted for communicating with another node. As discussed for claim element 1.k above, Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) have two or more radios for wirelessly communicating in a Cisco Meraki Wireless Mesh Network with another node or a client device. *See supra* discussion for claim element 1.k. APs in the Cisco Meraki Wireless Mesh Network have multiple radio interfaces, including 2.4 and 5 GHz radios.  Source: Wireless Mesh Networking

High Performance 802.11ax Wireless

The Cisco Meraki MR86 is a cloud-managed 4x4:4 802.11ax access point that raises the bar for wireless performance and efficiency. Designed for next-generation deployments in offices, schools, hospitals, shops, and hotels, the MR86 offers high throughput, enterprise-grade security, and simple management.

The MR86 provides a maximum of 3.0 Gbps* aggregate frame rate with concurrent 2.4 GHz and 5 GHz radios. A dedicated third radio provides real-time WIDS/WIPS with automated RF optimization, and a fourth integrated radio delivers Bluetooth scanning and beaconing.

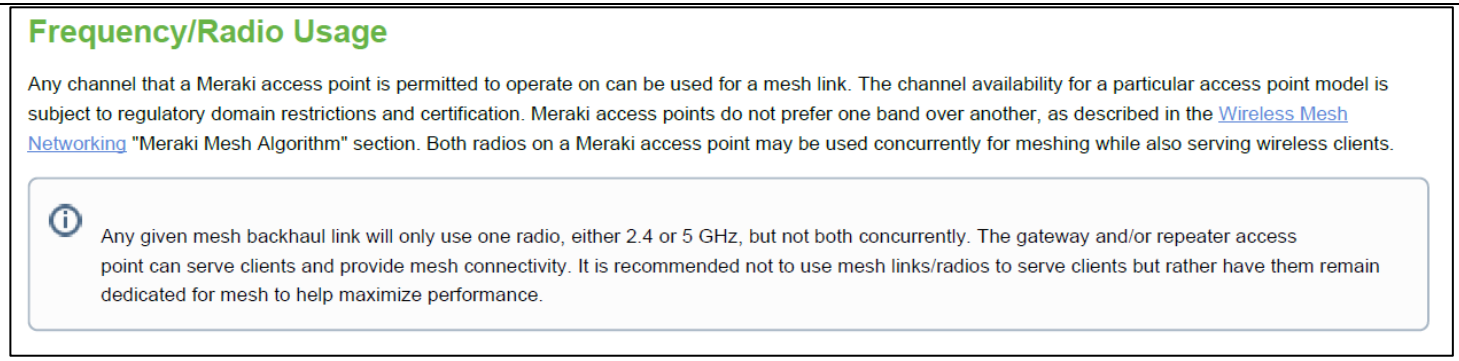
Source: MR86 Datasheet

Product options: Wi-Fi 6E indoor models

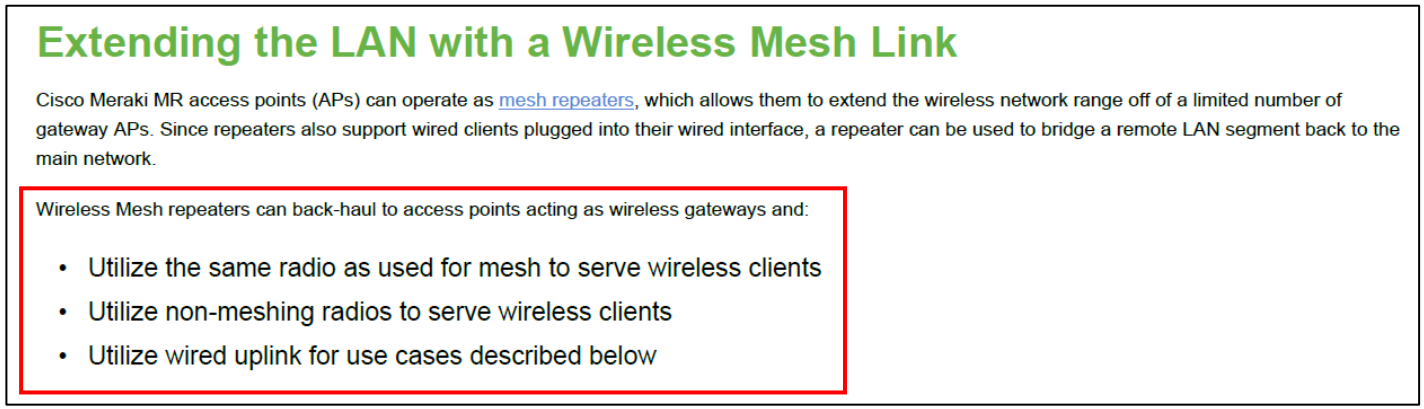
	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: Meraki Family Datasheet

Both radios may be used either for meshing (*i.e.*, adapted for communicating with another node) or serving wireless clients (*i.e.*, adapted for communicating with the client device) or both.



Source: Mesh Deployment Guide

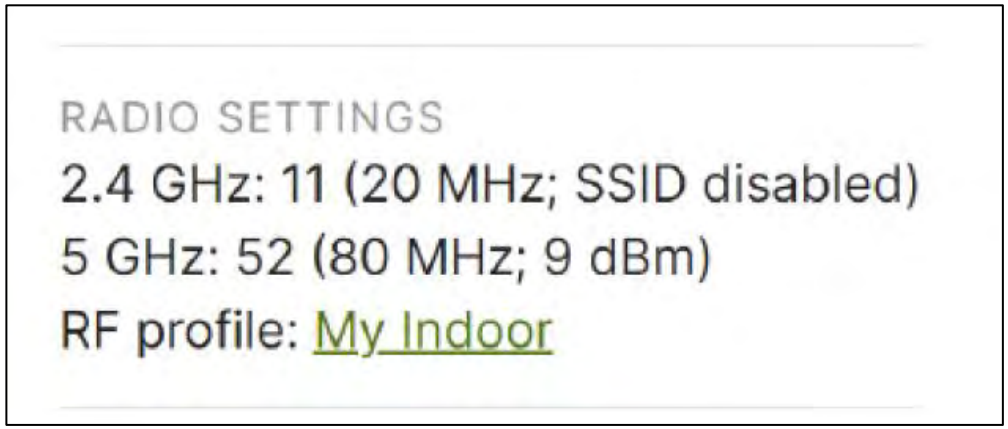


Source: Extending the LAN with a Wireless Mesh Link

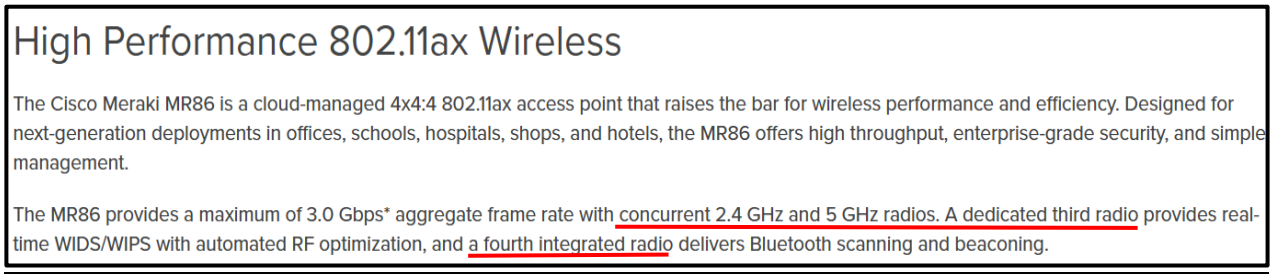
MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 14	
14. The wireless mesh network of claim 13, wherein the second radio is configurable to be used as a dual-purpose radio for communicating with both the client and another node.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 13 as shown above, are comprised of one or more mesh access point (MAP) nodes comprising two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network, wherein a first radio of the MAP is adapted for communicating with the client device and wherein a second radio is configurable to be used as a dual-purpose radio for communicating with both the client and another node.

As discussed for claim elements 1.k and 13 above, Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) have two or more radios for wirelessly communicating in a Cisco Meraki Wireless Mesh Network with another node or a client device. *See* supra discussion for claim element 1.k and 13. APs in the Cisco Meraki Wireless Mesh Network have multiple radio interfaces, including 2.4 and 5 GHz radios.



Source: Wireless Mesh Networking



Source: MR86 Datasheet

Product options: Wi-Fi 6E indoor models

	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: Meraki Family Datasheet

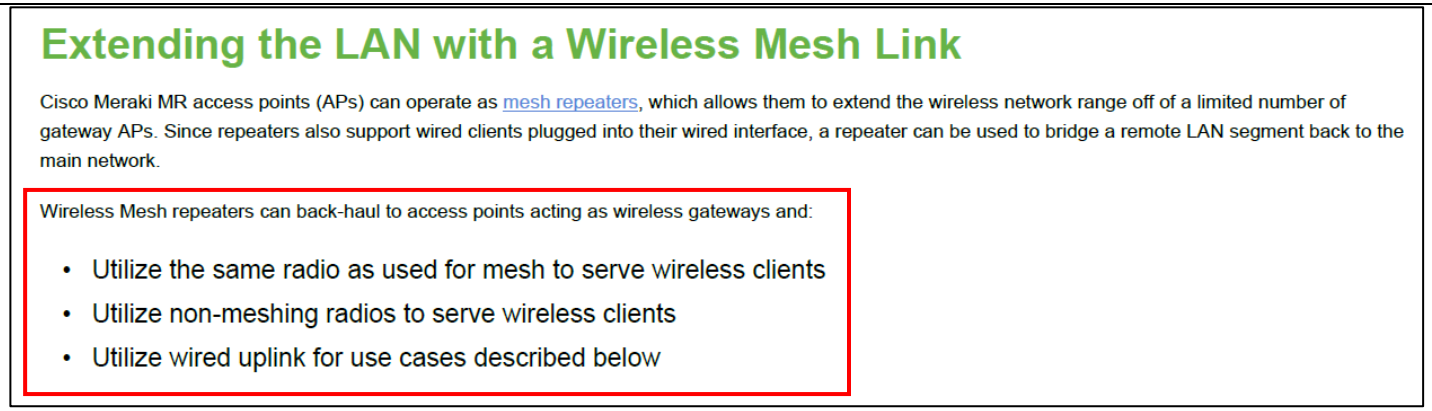
Both radios may be used either for meshing (*i.e.*, adapted for communicating with another node) or serving wireless clients (*i.e.*, adapted for communicating with the client device) or both (*i.e.*, configurable to be used as a dual-purpose radio for communicating with both the client and another node).

Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

ⓘ Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

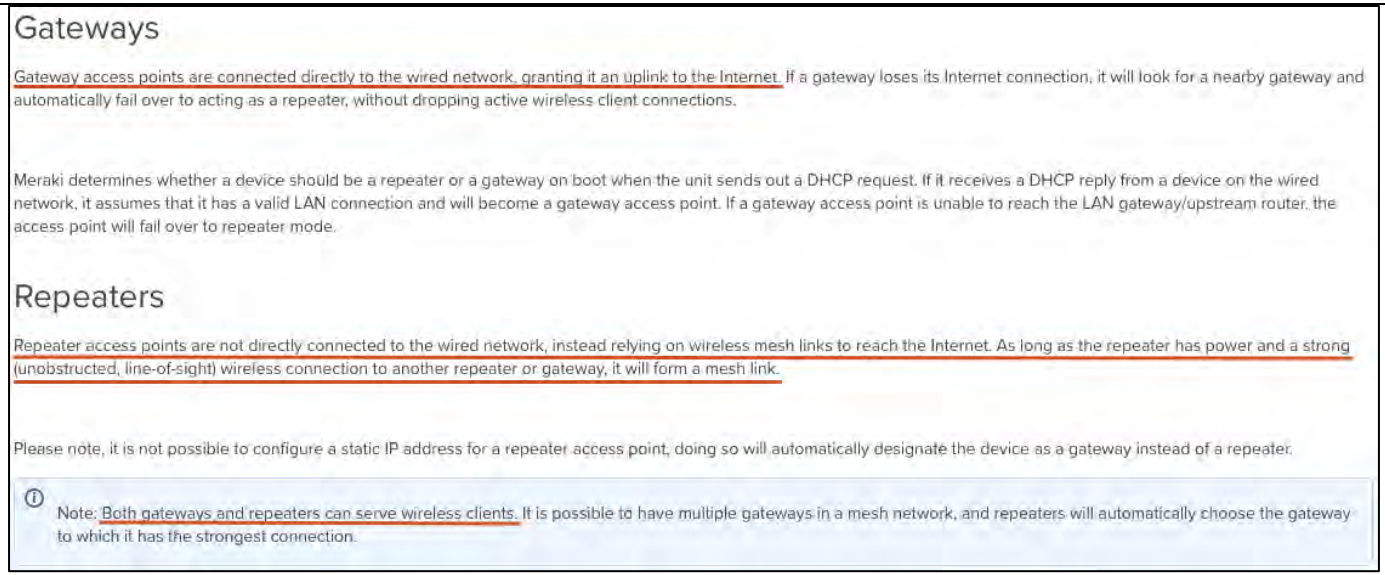
Source: Mesh Deployment Guide

<table>
<tr><td></td><td> 

Source: Extending the LAN with a Wireless Mesh Link

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.</td></tr>
<tr><th>Claim 15</th><th></th></tr>
<tr><td>15. The wireless mesh network of claim 1 further comprising a client device wirelessly connected to the MAP node, wherein the client device exchanges data with the external network through the MAP's routing path to the external network.</td><td>The Cisco Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises a client device wirelessly connected to the MAP node, wherein the client device exchanges data with the external network through the MAP's routing path to the external network.

As discussed for claim 1 above, a repeater selects the best mesh route to a gateway for data communication with the external network. See supra discussion for claim 1. Such a repeater can serve wireless client and provide data exchange between the client and the external network through the repeater's mesh route.</td></tr>
</table>

<table>
<tr><td></td><td> 

Source: Wireless Mesh Networking</td></tr>
<tr><td>Claim 16</td><td></td></tr>
<tr><td>16.a. The wireless mesh network of claim 1 further comprising one or more data types being routed through the routing path, wherein the one or more data types include at least two of: voice data, video data, and other data type that is not the voice or video data;</td><td>The Cisco Meraki Wireless Mesh Network, in addition to the elements of claim 1 as shown above, comprises one or more data types being routed through the routing path, wherein the one or more data types include at least two of: voice data, video data, and other data type that is not the voice or video data.

The Cisco Meraki Wireless Mesh Network supports Quality of Service for prioritization of traffic based on data types, which include VoIP (i.e., voice) data, enterprise application data, and other data. Therefore, the mesh network data types include at least two of voice, video and other data type.

 
</td></tr>
</table>

How does QoS work?

Administratively defined, QoS network rules in Dashboard inform the switch how DSCP tags should be assigned to incoming packets. QoS on a MS Switch will tag specific types of traffic that is incoming on switch ports and will prioritize traffic based on the assigned tag. These DSCP bits are mapped to specific layer-2 CoS queues on the switch. Once a packet is assigned a DSCP tag, it is placed into the corresponding layer-2 CoS queue for forwarding. The CoS values are queues, not bits in outgoing frames.

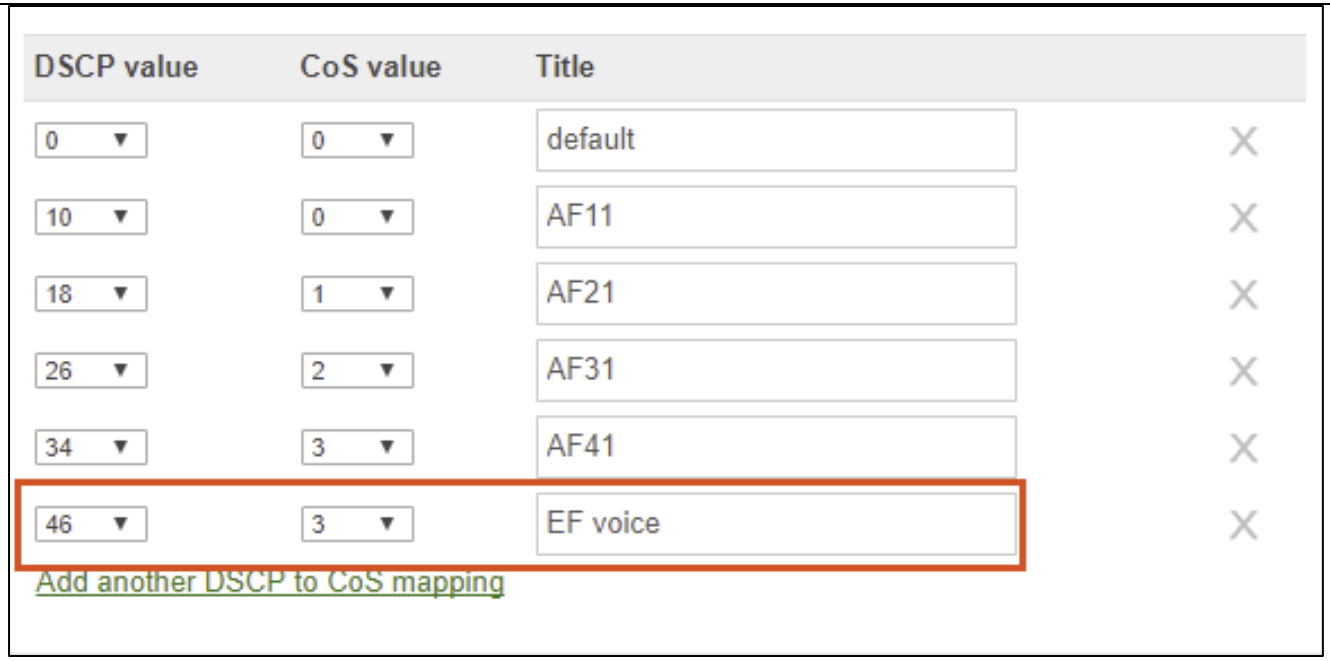
QoS guarantees a certain fraction of the uplink to each configured priority level when the link is congested. Frames in higher priority queues receive a bigger slice of bandwidth than those in a lower priority queue. If a queue is not fully utilized, the bandwidth will be used by other queues so that bandwidth is not wasted.

An MS network has 6 configurable CoS queues labeled 0-5. Each queue is serviced using FIFO. Without QoS enabled, all traffic is serviced in queue 0 (default class) using a FIFO model. The queues are weighted as follows:

CoS	Weight
0 (default class)	1
1	2
2	4
3	8
4	16
5	32

Suppose a switched environment where VoIP traffic should be in CoS queue 3, an enterprise application in CoS queue 2, and all of other traffic is unclassified. The percentage of bandwidth allocation can be calculated using the weight of the individual CoS queue as the numerator and the sum of all configured CoS queues as the denominator (in this example 8+4+1=13):

- VoIP would be guaranteed 8/13 or ~62% percent of the bandwidth. The switch would forward 8 frames from the CoS queue 3 and move to CoS queue 2.
- The enterprise application would be guaranteed 4/13 or ~30% bandwidth.The switch would forward 4 frames from the CoS queue 2 and move to the default queue.
- All other traffic would receive 1/13 or ~8% of the bandwidth. The switch would forward 1 frame from the default queue, then cycle back to CoS queue 3.

	 Source: MS Switch Quality of Service Defined MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
16.b. wherein routing a priority data type with a higher priority than another type of data, wherein the priority data type may be one or both of: the voice data or the video data.	The Cisco Meraki Wireless Mesh Network comprises one or more data types being routed through the routing path, wherein routing a priority data type with a higher priority than another type of data, wherein the priority data type may be one or both of: the voice data or the video data. The Cisco Meraki Wireless Mesh Network supports Quality of Service for prioritization of traffic based on data types, which include VoIP (*i.e.*, voice) data, enterprise application data, and other data. 

How does QoS work?

Administratively defined, QoS network rules in Dashboard inform the switch how DSCP tags should be assigned to incoming packets. QoS on a MS Switch will tag specific types of traffic that is incoming on switch ports and will prioritize traffic based on the assigned tag. These DSCP bits are mapped to specific layer-2 CoS queues on the switch. Once a packet is assigned a DSCP tag, it is placed into the corresponding layer-2 CoS queue for forwarding. The CoS values are queues, not bits in outgoing frames.

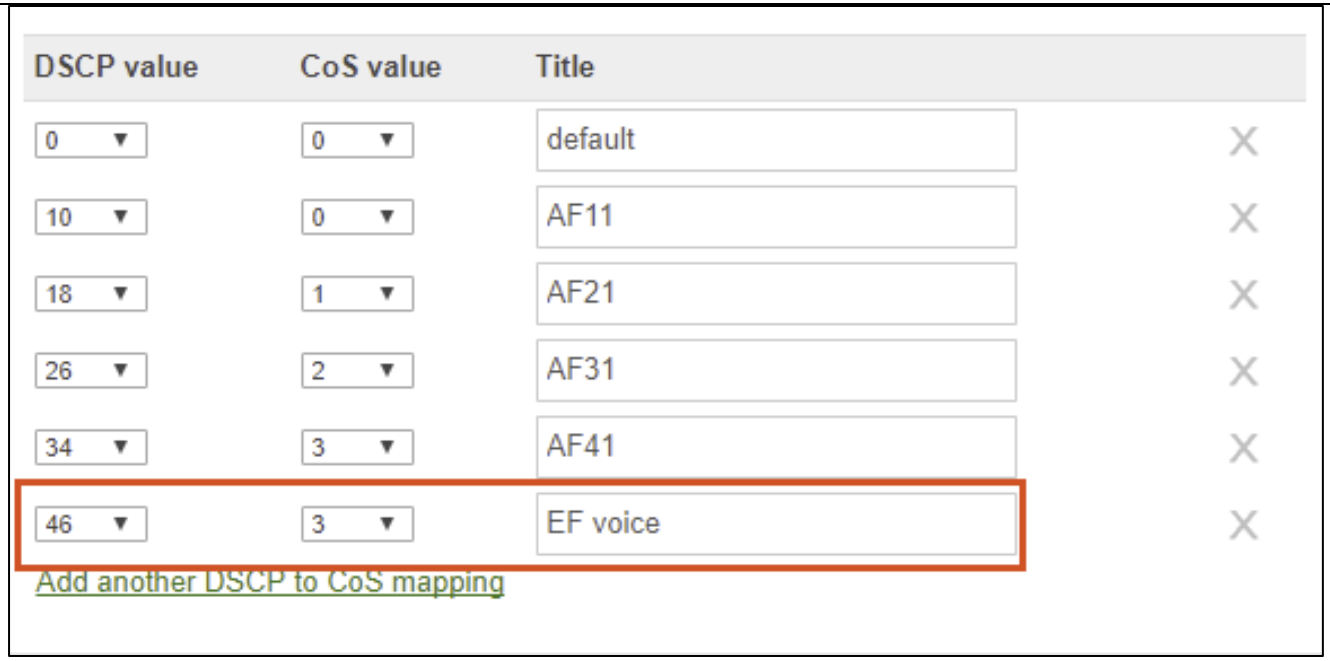
QoS guarantees a certain fraction of the uplink to each configured priority level when the link is congested. Frames in higher priority queues receive a bigger slice of bandwidth than those in a lower priority queue. If a queue is not fully utilized, the bandwidth will be used by other queues so that bandwidth is not wasted.

An MS network has 6 configurable CoS queues labeled 0-5. Each queue is serviced using FIFO. Without QoS enabled, all traffic is serviced in queue 0 (default class) using a FIFO model. The queues are weighted as follows:

CoS	Weight
0 (default class)	1
1	2
2	4
3	8
4	16
5	32

Suppose a switched environment where VoIP traffic should be in CoS queue 3, an enterprise application in CoS queue 2, and all of other traffic is unclassified. The percentage of bandwidth allocation can be calculated using the weight of the individual CoS queue as the numerator and the sum of all configured CoS queues as the denominator (in this example 8+4+1=13):

- VoIP would be guaranteed 8/13 or ~62% percent of the bandwidth. The switch would forward 8 frames from the CoS queue 3 and move to CoS queue 2.
- The enterprise application would be guaranteed 4/13 or ~30% bandwidth.The switch would forward 4 frames from the CoS queue 2 and move to the default queue.
- All other traffic would receive 1/13 or ~8% of the bandwidth. The switch would forward 1 frame from the default queue, then cycle back to CoS queue 3.



Source: MS Switch Quality of Service Defined

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 17	
17.a. A wireless mesh network comprising: a plurality of Wi-Fi nodes in the wireless mesh network organized in a tree shape;	*See supra* discussion for claim element 1.a., the contents of which are incorporated herein by reference.
17.b. an access server in data communication with the nodes,	*See supra* discussion for claim element 1.b., the contents of which are incorporated herein by reference.
17.c. wherein the nodes are comprised of:	*See supra* discussion for claim element 1.c., the contents of which are incorporated herein by reference.

one or more root access point (RAP) nodes having a wired connection to an external network; one or more mesh access point (MAP) nodes;	
17.d. wherein a MAP node is in a data communication with the external network through an associated RAP node;	*See supra* discussion for claim element 1.d., the contents of which are incorporated herein by reference.
17.e. wherein the MAP node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node;	*See supra* discussion for claim element 1.e., the contents of which are incorporated herein by reference.
17.f. wherein the MAP node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node, process for the automatic connection comprising:	*See supra* discussion for claim element 1.f., the contents of which are incorporated herein by reference.

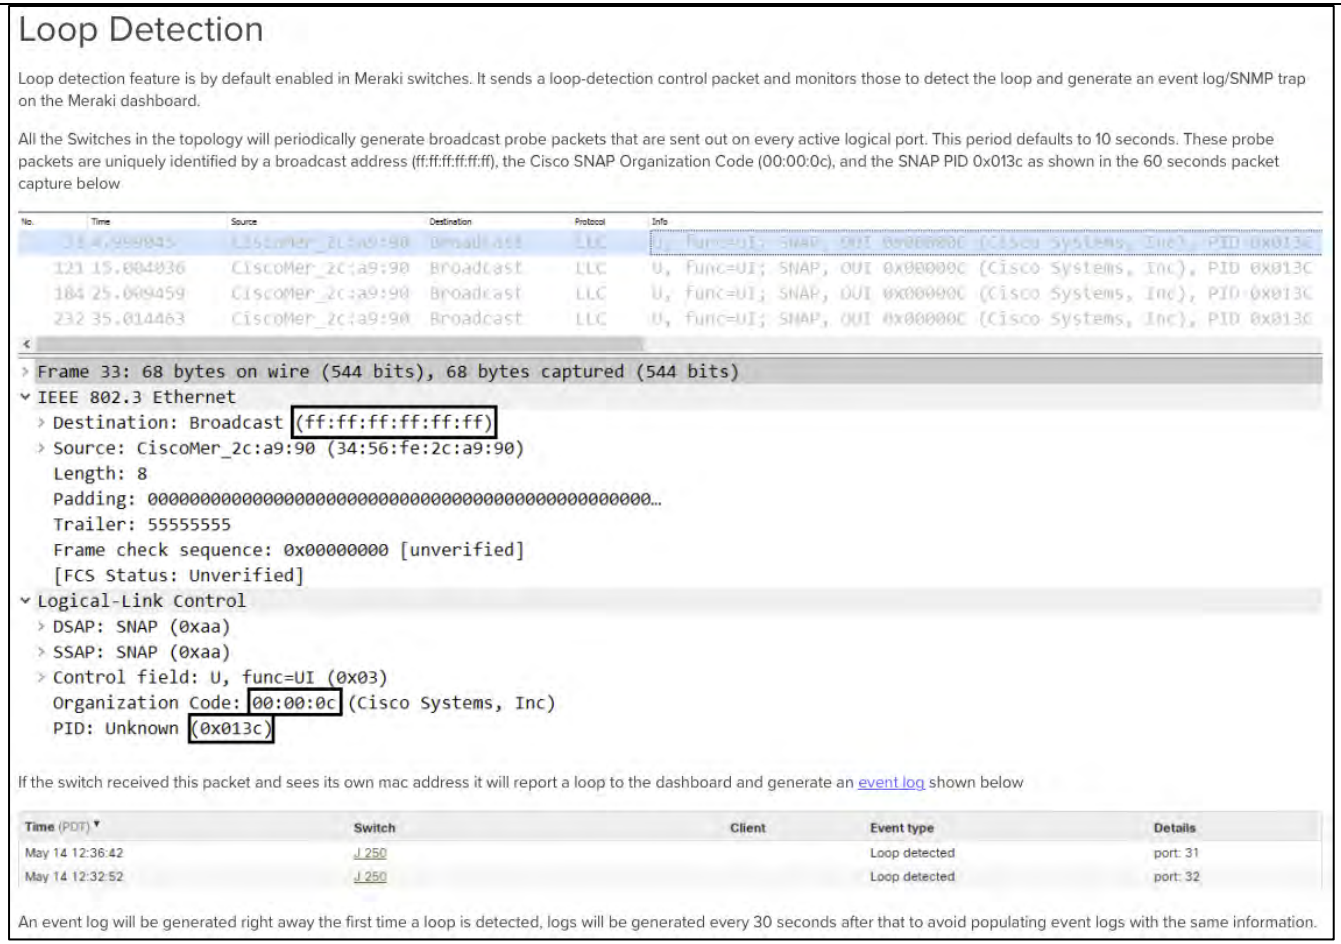
17.g. selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node; and	*See supra* discussion for claim element 1.g., the contents of which are incorporated herein by reference.
17.h. establishing a direct wireless data connection between the MAP node and the associated parent node;	*See supra* discussion for claim element 1.h., the contents of which are incorporated herein by reference.
17.i. wherein the automatic connection process establishes the MAP node's routing path to the external network;	*See supra* discussion for claim element 1.i., the contents of which are incorporated herein by reference.
17.j. the one or more parent selection criteria comprising: a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path;	*See supra* discussion for claim element 1.j., the contents of which are incorporated herein by reference.
17.k. the MAP node comprising: one or more identifier values representing a unique identity of the MAP node in the wireless mesh network;	*See supra* discussion for claim element 1.k., the contents of which are incorporated herein by reference.
17.l. two or more radios for wirelessly communicating	*See supra* discussion for claim element 1.l., the contents of which are incorporated herein by reference.

with another node or a client device in the wireless mesh network;	
17.m. one or more network health monitoring criteria defined by the access server, wherein the one or more network health monitoring criteria comprises one or more of: whether or not the associated parent node is offline or can be found in the wireless mesh network; whether or not data can be routed to the external network through the MAP node's routing path; whether or not the one or more parent selection criteria is satisfied; and whether or not the MAP node's routing path is more congested than a threshold load value representing data congestion level of a routing path;	*See supra* discussion for claim element 2.a., the contents of which are incorporated herein by reference.
17.n. testing whether the MAP node’s routing path fails to satisfy the one or more network health monitoring criteria;	*See supra* discussion for claim element 2.b., the contents of which are incorporated herein by reference.

17.o. when it is determined that the MAP node's routing path does not meet one or more of the one or more network health monitoring criteria, establishing a new routing path from the MAP node to the external network by connecting the MAP node to a new associated parent node through the automatic connection process.	*See supra* discussion for claim element 2.c., the contents of which are incorporated herein by reference.
Claim 18	
18.a. A wireless mesh network comprising: a plurality of Wi-Fi nodes in the wireless mesh network organized in a tree shape;	*See supra* discussion for claim element 1.a., the contents of which are incorporated herein by reference.
18.b. an access server in data communication with the nodes,	*See supra* discussion for claim element 1.b., the contents of which are incorporated herein by reference.
18.c. wherein the nodes are comprised of: one or more root access point (RAP) nodes having a wired connection to an external network; one or more mesh access point (MAP) nodes;	*See supra* discussion for claim element 1.c., the contents of which are incorporated herein by reference.
18.d. wherein a MAP node is in a data communication with the external network	*See supra* discussion for claim element 1.d., the contents of which are incorporated herein by reference.

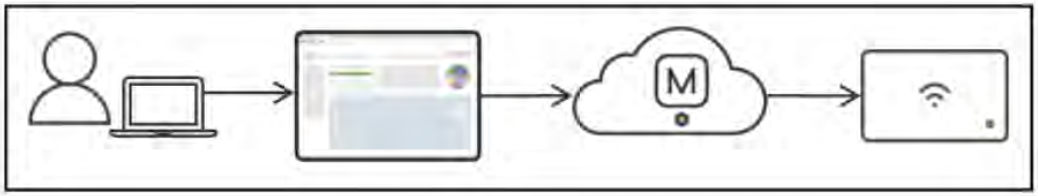
through an associated RAP node;	
18.e. wherein the MAP node is in a direct wireless data communication with a single associated parent node, wherein the associated parent node is either the associated RAP node, or another MAP node in a data communication with the external network through its associated RAP node;	*See supra* discussion for claim element 1.e., the contents of which are incorporated herein by reference.
18.f. wherein the MAP node automatically connects to the associated parent node by selecting the single associated parent node from one or more potential parent nodes that are within a direct wireless communication range of the MAP node and establishing a parent-child relationship with the associated parent node, process for the automatic connection comprising:	*See supra* discussion for claim element 1.f., the contents of which are incorporated herein by reference.
18.g. selecting a parent node from the one or more potential parent nodes based at least in part on one or more parent selection criteria defined by the access server, wherein the parent node so selected is the associated parent node; and	*See supra* discussion for claim element 1.g., the contents of which are incorporated herein by reference.

18.h. establishing a direct wireless data connection between the MAP node and the associated parent node;	*See supra* discussion for claim element 1.h., the contents of which are incorporated herein by reference.
18.i. wherein the automatic connection process establishes the MAP node's routing path to the external network;	*See supra* discussion for claim element 1.i., the contents of which are incorporated herein by reference.
18.j. the one or more parent selection criteria comprising: a loop prevention criterion which verifies that a potential parent node's routing path to the external network does not contain the MAP node; a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path;	As discussed for claim element 1.j, the one or more parent selection criteria comprising: a throughput value representing signal strength between two nodes; and a latency value representing number of hops in a routing path. *See supra* discussion for claim element 1.i. and 1.j. On information and belief, the one or more parent selection criteria further comprises a loop prevention criterion which verifies that a potential parent node's routing path to the external network does not contain the MAP node:

	 Source: Loop and MAC Flap Detection on MS MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
18.k. the MAP node comprising: one or more identifier values representing a unique	*See supra* discussion for claim element 1.k., the contents of which are incorporated herein by reference.

identity of the MAP node in the wireless mesh network;	
18.l. two or more radios for wirelessly communicating with another node or a client device in the wireless mesh network.	*See supra* discussion for claim element 1.l., the contents of which are incorporated herein by reference.
Claim 19	
19. The wireless mesh network of claim 18 wherein the selection of the parent node further comprises not selecting a potential parent node if the MAP node's potential routing path to the external network through the potential parent node exceeds a threshold latency value.	The Meraki and Catalyst CW Access Points (*e.g.*, the MR and CW series access points) in the Cisco's Meraki Wireless Mesh Network, in addition to the elements of claim 18 as shown above, are comprised of one or more mesh access point (MAP) nodes wherein the selection of the parent node further comprises not selecting a potential parent node if the MAP node's potential routing path to the external network through the potential parent node exceeds a threshold latency value. Cisco's Meraki Wireless Mesh Network use the Meraki Mesh Algorithm to automatically detect and select the best mesh route to a gateway. The algorithm does so by comparing, *inter alia*, number of hops (*i.e.*, latency value) for each mesh route and preference is given to gateways with lower hops. Cisco recommends no more than one mesh hop from the gateway to the client device. Mesh Neighbors The Mesh Neighbors table can be found on the **Summary Tab** on the access point details page and shows the access points that have been discovered automatically. The link quality is a metric that takes into account signal strength and packet delivery success rates in each direction. A link quality of 70% or higher is recommended for a strong link. Source: Wireless Mesh Networking Maximum Mesh Hops There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device. Source: Mesh Deployment Guide The Meraki Cloud Server can be used to manage and configure the access points. On information and belief, the Meraki Cloud Server can be used to configure the maximum number of hops (*i.e.*, threshold latency value) for a mesh route to be selected by a repeater.

The Meraki cloud solution is a centralized management service that allows users to manage all of their Meraki network devices via a single, simple and secure platform.



Users are able to deploy, monitor, and configure their Meraki devices via the Meraki dashboard web interface or via APIs. Once a user makes a configuration change, the change request is sent to the Meraki cloud and is then pushed to the relevant device(s).

Management data: The data (configuration, statistics, monitoring, etc.) that flows from Meraki devices (wireless access points, switches, security appliances) to the Meraki cloud over a secure internet connection.

Management Data

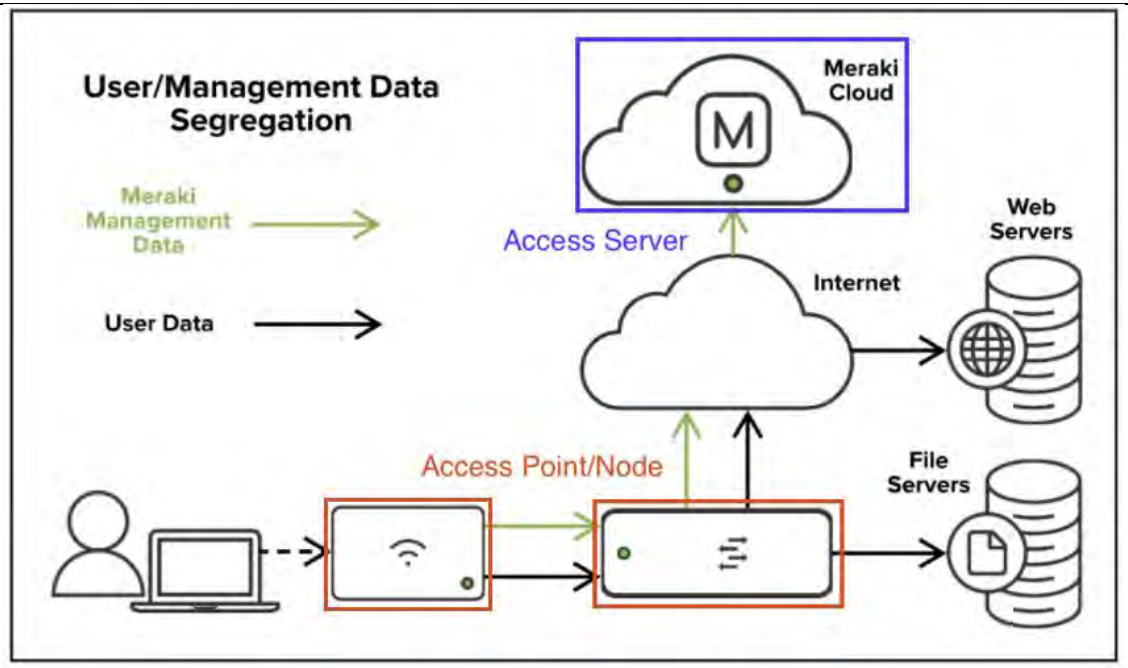
The Meraki cloud gathers and stores certain types of "management" data to enable its solutions. All forms of data are encrypted in transit to and from Meraki servers. There are four major types of data stored in the Meraki cloud:

User records: Includes account email and company name or other optional information such as user name and address.

Configuration data: Includes network settings and configurations made by customers in the Meraki dashboard.

Analytics data: Includes client, traffic, and location analytics data, providing visualizations and network insights into traffic patterns across customer sites.

Customer-uploaded assets: Includes custom floor plans and splash logos.

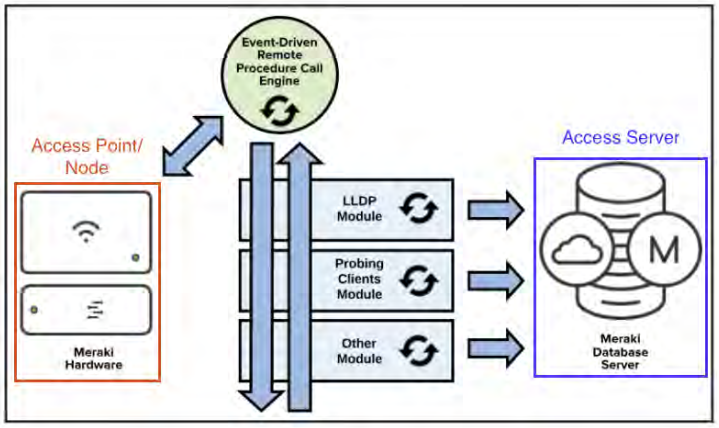


Meraki Cloud Architecture

The Meraki cloud is the backbone of the Meraki management solution. This "cloud" is a collection of highly reliable multi-tenant servers strategically distributed around the world at Meraki data centers and select public cloud service providers. For the remainder of this document we will refer to both the Meraki data centers and the public cloud service providers' locations we use as "data centers". The servers at these data centers are powerful hosting computers comprised of many separate user accounts. They are called multi-tenant servers because the accounts share (equal) computing resources on their host (the server). However, even though these accounts share resources, Meraki ensures that customer information is kept secure by restricting organization access based on account authentication, as well as hashing authentication information such as user passwords or API keys.

Meraki Device-to-Cloud Communications

Meraki uses an event-driven remote procedure call (RPC) engine for Meraki devices to communicate to the dashboard and for Meraki servers to send and receive data. Meraki hardware devices act as the server/receiver as the Meraki cloud initiates calls to the devices for data collection and configuration deployment. The cloud infrastructure is the initiator, so configurations can be executed in the cloud before the devices are actually online or even physically deployed.



Communication Process

If a device is offline, it will continue to attempt to connect to the Meraki cloud until it gains connectivity. Once the device comes online, it automatically receives the most recent configuration settings from the Meraki cloud. If changes are made to the device configuration while the device is online, the device receives and updates these changes automatically. These changes are generally available on the device in a matter of seconds. However, large quantities of changes may take noticeably longer to reach their devices. If no configuration changes are made by the user, the device continues to periodically check for updates to its configuration on its own.

As the device runs on the network, it will communicate device and network usage analytics back to the Meraki cloud. Dashboard analytics based on this information, in the form of graphs and charts, are updated regularly in the Meraki cloud and are displayed in the dashboard of users when they are viewing this information.

Source: Meraki Cloud Architecture

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (selecting a gateway) in the same way (using criteria such as throughput or signal strength) and achieves the same result (best mesh route to a gateway). Any differences between the accused Cisco's Meraki Wireless Mesh Network and the claimed element are insignificant or insubstantial. Cisco's Meraki Wireless Mesh Network will select the parent node based at least in part on its throughput value. A person skilled in the relevant art would consider any difference between the claimed invention and the accused network to be minor.

MeshDynamics asserts that these claim elements are software limitations. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Exhibit 9 – U.S. Patent No. 8,520,691 (Cisco Wireless Mesh Networks with Catalyst and Aironet Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Wireless Mesh Network including but not limited to Cisco's Catalyst and Aironet Access Points, with related controllers and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1, 9-11, and 35 (the "Asserted Claims") of U.S. Patent No. 8,502,691 (the "'691 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '691 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringed by testing, setting up, and using the Cisco Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Wireless Mesh Network in the United States.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the '691 Patent and make and use the infringing Cisco Wireless Mesh Network. Specifically, Cisco's Wireless Mesh Access Points, Design and Deployment Guides, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco's website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1, 9-11, and 35 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Wireless Mesh Network, and the infringing configuration and operation of the controller, access points, and software that comprise the infringing Cisco Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the '691 Patent and/or its predecessors

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd.*, 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g.*, https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

since 2009 when Cisco executives met with MeshDynamics to discuss partnering opportunities and discussed the application that would result in the '691 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the '691 Patent, or in the alternative was willfully blind to the '691 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringes the '691 Patent. Cisco Accused Products come enabled as a mesh access point and automatically connect to the controller to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particular suited for Cisco's Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco's selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the '691 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco's Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the '691 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends in the alternative that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics's preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts

presently known to it. To the best of MeshDynamics's knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco document productions, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
1pre: A structured mesh network capable of isolated operation, comprising:	Cisco's Wireless Mesh Network (a component of the Cisco Unified Wireless Network solution) is a structured mesh network capable of isolated operation, (*e.g.*, fault tolerant resilient mode), which comprises a plurality of Wi-Fi nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Series Access Points). Cisco Catalyst 9100 and Aironet Access Points support Wi-Fi capabilities. The Catalyst series is a follow-up from the Cisco Aironet series access points, and the Aironet and Catalyst Access Points contain the same mesh networking capabilities, as shown below in Cisco's documentation and guides that apply to both product lines: Q. Which Cisco access points support Wi-Fi 6? A. Cisco Catalyst 9100 enterprise-class access points support Wi-Fi 6 capabilities, extending the power of intent-based networking with hardware and software innovations. The Cisco Catalyst 9100 APs offer resiliency and security while enabling intelligence at the network edge. Q. What is the portfolio transition from the Cisco Aironet Series to the Cisco Catalyst 9100 platform? A. The Cisco Catalyst 9100 Access Points consist of the Cisco Catalyst 9105, 9115, 9120, 9130, 9136, 9162, 9164, 9166 and 9166D1, and outdoor 9163E Series. These access points are the follow-ups to the Cisco Aironet 1815, 1850, 2800, and 3800 Series Access Points, respectively. The Cisco Catalyst 9100 Access Points come equipped with Wi-Fi 6/6E capabilities, a better industrial design, and improved RF performance, and deliver reliability, security, and intelligence at scale. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

Cisco Aironet 1560 Series Outdoor Access Points

Product Type	Outdoor and Industrial Wireless
Status	End of Sale **EOL Details**
Series Release Date	05-JUL-2016
End-of-Sale Date	31-JAN-2023
End-of-Support Date	31-JAN-2028

This product is supported by Cisco, but is no longer being sold. Ready for an upgrade? The **Cisco Catalyst 9124AX Series Access Points** offer greater performance and functionality.

Source: https://www.cisco.com/c/en/us/support/wireless/aironet-1560-series/series.html

Mesh deployments are supported on the Cisco Catalyst 9100 Access Points:

Q. Will Mesh deployment be supported on the Catalyst 9124 Access Points?

A. Yes. Mesh deployments will be supported on the Catalyst 9124 Access Points. Embedded Wireless Controller deployment is also supported with the Catalyst 9124 Access Point acting as a controller for up to 100 APs.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-802-11ax-faq-cte-en.html

In addition to the Wave 1 and Wave 2 access points, mesh deployments are supported on the Cisco Catalyst 9100 Access Points, including, *e.g.*, Catalyst 9124 Access Points, Catalyst 9130 Access Points and "AX" access points:

Families of mesh APs

The two families of mesh APs available are the Cisco Wave 1 APs and the Cisco Wave 2 APs + the AX APs. These two families of mesh APs have different capabilities and features, and they are not fully compatible with each other. The Cisco Wave 1 APs do not support the same feature set as the Cisco Wave 2 + AX APs.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-15/config-guide/b_wl_17_15_cg/m_mesh_ewlc.html; see also https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg/mesh-access-points.html

From this release, mesh support is included in the Cisco Catalyst 9130AX Series Access Points. All traditional capabilities of mesh are included in the Cisco Catalyst 9130AX Series APs operating in Cisco IOS XE Dublin 17.12.1.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/ewc/17-13/config-guide/ewc_cg_17_13/mesh_access_points.html

Cisco Catalyst 9100 and Aironet Access Points can be deployed in a Cisco Wireless Mesh Network with a controller and software, such as the Catalyst 9800 Controller and Cisco DNA and Prime Infrastructure:

Mesh Network Components

This section describes the mesh network components.

The Cisco wireless mesh network has four core components:

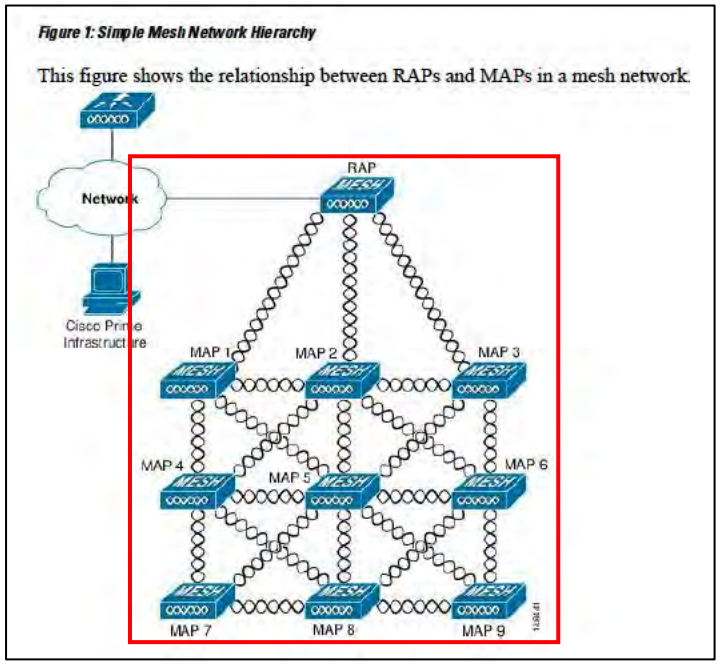
- Cisco C9800 and IOS-XE 17.1
- Cisco Aironet and Catalyst series access points
- Cisco PI and Cisco DNA Center
- Mesh software architecture

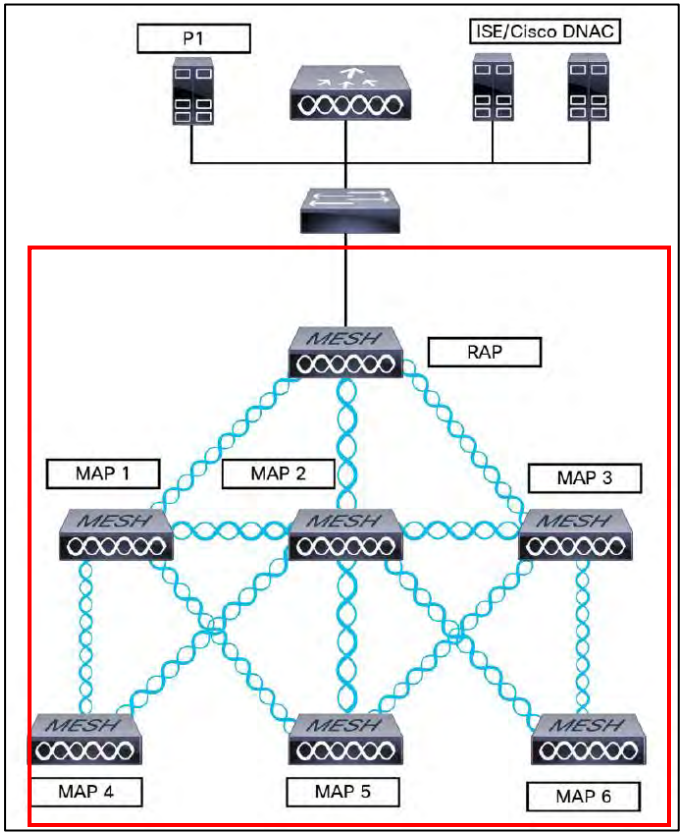
Mesh networking employs Cisco Aironet 1540, 1560, or 1570 Series access points and the latest Catalyst 9124AX Series outdoor mesh access points; Catalyst 9800 Series wireless controllers; and Cisco DNA Center to provide scalable, central management and mobility between indoor and outdoor deployments. The CAPWAP protocol manages the connection of mesh access points to the network.

- Cisco Aironet Wave 1 indoor APs: 1700, 2700, and 3700 series.
- Cisco Aironet Wave 2 indoor APs: 1815i, 1815m, 1830, 1850, 2800, 3800, and 4800 Series.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh; *see also* https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

Cisco Wireless Mesh Networks organize the wireless access points in a structured mesh network and tree-like configuration with root access point(s) ("RAP") connected to the external network and mesh access points ("MAPs") connected to the RAP, as shown in red below:





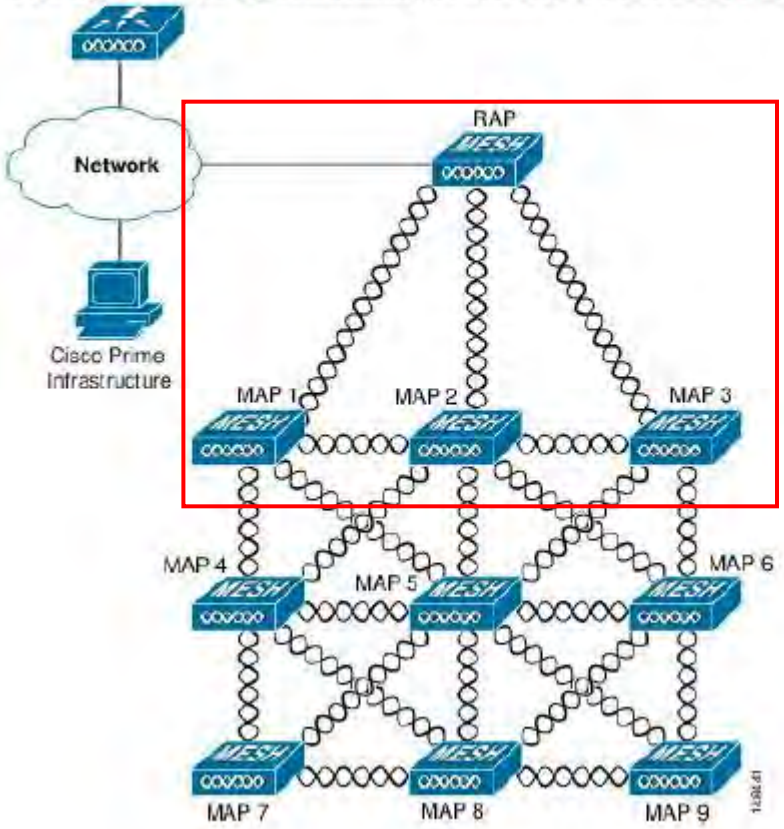
Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

If the mesh network is isolated from the controller (and the external network), the mesh network is capable of continued operation in different modes, *e.g.*, Flex + Bridge mode and/or fault tolerant resilient mode:

	A Control and Provisioning of Wireless Access Points protocol (CAPWAP) Access Point (AP) can be configured to operate in two different modes: • FlexConnect mode • Bridge/Mesh mode The following are the bridging features for Flex+Bridge mode: • The Flex+Bridge mode supports the centrally switched 802.11 WLAN. Traffic for this tunneled WLAN is forwarded to and from a CAPWAP controller over an IP tunnel. • The Flex+Bridge mode supports the Root Ethernet VLAN Bridging. A root AP bridges the traffic for bridged 802.11 WLANs and secondary Ethernet LANs to a local Ethernet LAN over its root Ethernet port. • The Flex+Bridge mode bridging is supported on Secondary Ethernet Access Ports and Secondary Ethernet VLAN Trunk Ports. • Fault Tolerant Resilient Mode enables an AP to continue bridging traffic when the connection to the CAPWAP controller is lost. Both mesh and non-mesh root APs continue to bridge traffic. A child mesh AP (MAP) maintains its link to a parent AP and continues to bridge traffic till the parent link is lost. A child mesh AP cannot establish a new parent or child link till it reconnects to the CAPWAP controller. Existing wireless clients on the locally switching WLAN can stay connected with their AP in this mode. Their traffic will continue to flow through the Mesh and wired network. No new or disconnected wireless client can associate to the Mesh AP in this mode. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/flexconnect.html.
1(a): at least two structured mesh nodes; wherein each structured mesh node comprises at least a connectivity logic;	Cisco's Wireless Mesh Network comprises at least two structured mesh nodes (*e.g.*, Cisco Catalyst 9100 and Aironet Series Access Points) wherein each structured mesh node comprises at least a connectivity logic. Catalyst 9100 and Aironet Access Points (two or more) in Cisco Wireless Mesh Networks use connectivity logic, for example, in the Adaptive Wireless Path Protocol (AWPP) to determine the best path through the mesh access points to the controller. This includes selection of the best path for each MAP to the RAP. The AWPP is specifically designed for wireless mesh networks and to maintain the tree-like topology.

Figure 1: Simple Mesh Network Hierarchy

This figure shows the relationship between RAPs and MAPs in a mesh network.



Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

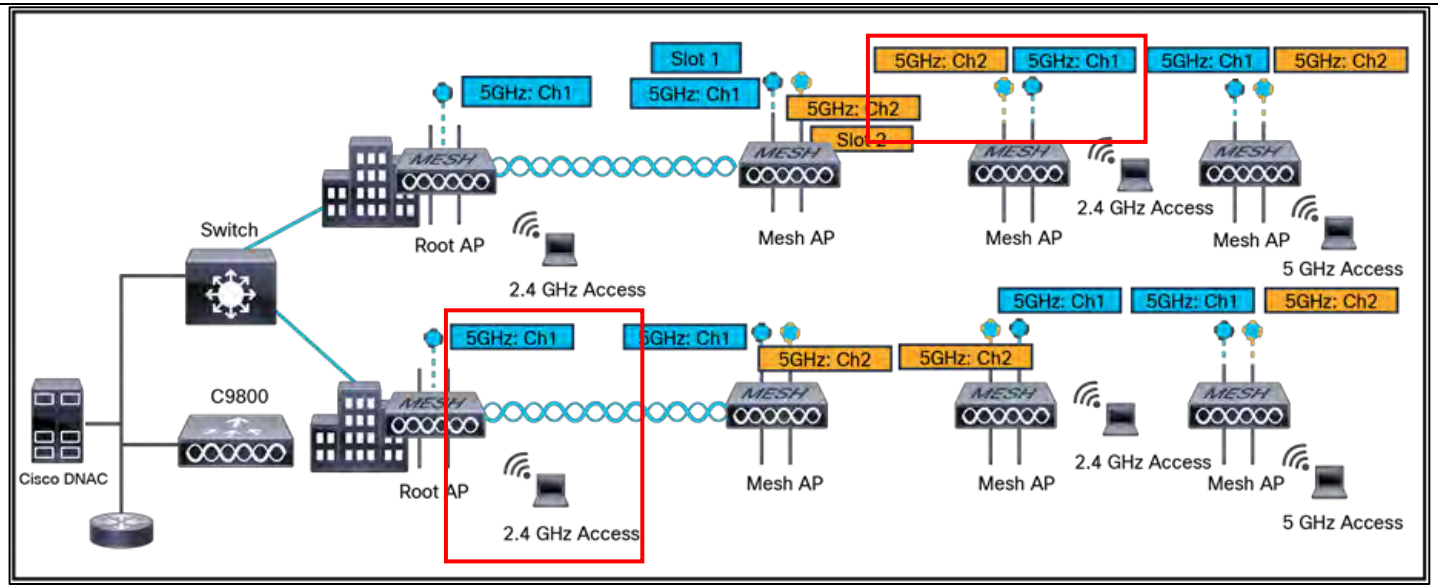
AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Each structured mesh node comprises at least a conectivity logic (*e.g.*, runs the Adaptive Wireless Path Protocol).

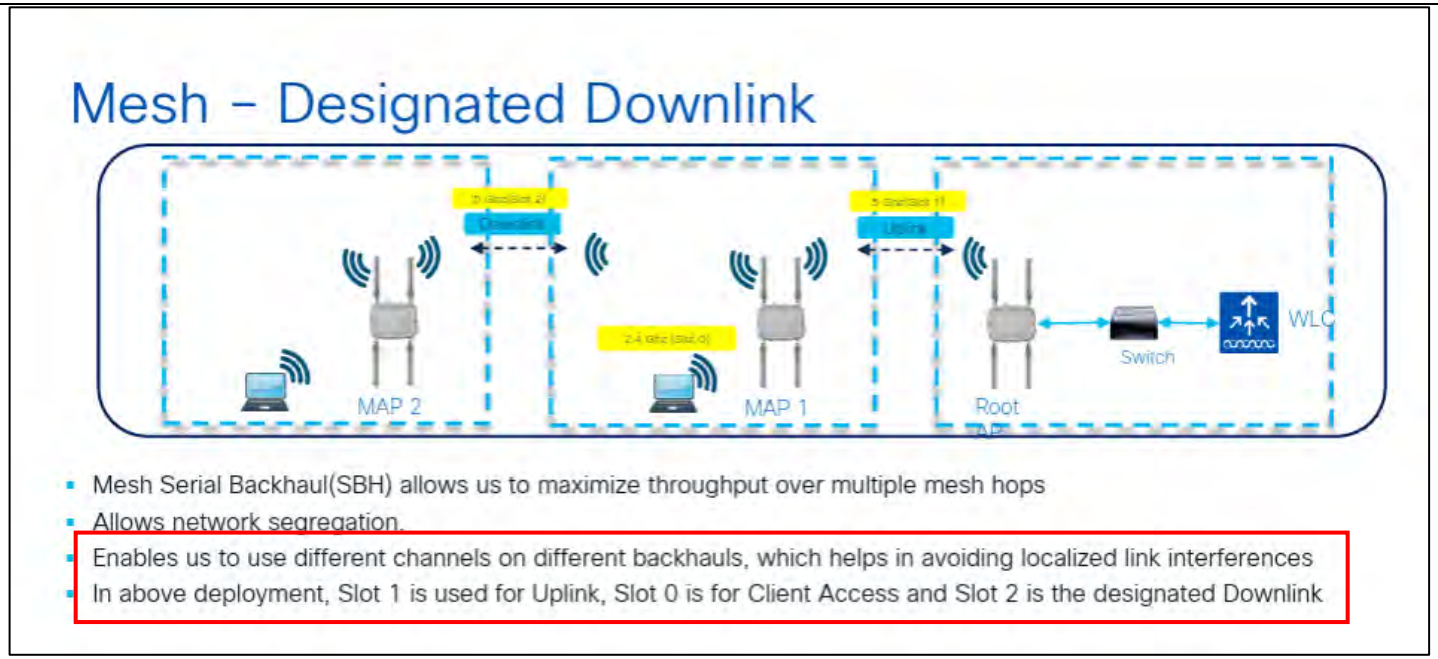
	AWPP Each AP runs the Adaptive Wireless Path Protocol (AWPP)-a new protocol designed from the ground up by Cisco specifically for the wireless environment. This protocol allows RAPs to communicate with each other to determine the best path back to the wired network via the RAP. Once the optimal path is established, AWPP continues to run in the background to establish alternative routes back to the RAP just in case the topology changes or conditions cause the link strength to weaken. This protocol takes into consideration things like interference and characteristics of the specific radio so that the mesh can be self-configuring and self-healing. AWPP actually has the ability to consider all relevant elements of the wireless environment so that the mesh network's functionality isn't disrupted and can provide consistent coverage. Source: https://www.learncisco.net/courses/cisco-ccna/part-13-wireless-networks/ciscos-wireless-solutions.html. On information and belief, the connectivity logic may also encompass firmware, software, and/or user interface component—whether automated or administrator-controlled—on the access point that facilitates connectivity decisions within the mesh network, to client devices, or to the external network, as will be shown in Cisco's proprietary code and information. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's solution achieves the same function (*e.g.*, determining whether a node connects to an external network or another mesh node), by substantially the same way (*e.g.*, through the AWPP algorithm executing either locally or as part of a centralized control scheme, as part of an overlayed SDN system such as Catalyst Center, or as part of the underlying architecture of each node), and yields the same result (*e.g.*, establishing uplink/downlink relationships and maintaining tree-based topology). This overarching behavior is functionally equivalent, and any implementation difference is insubstantial. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1(b): an uplink radio operating on an uplink frequency and a downlink	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise an uplink radio operating on an uplink frequency and a downlink radio operating on a distinct downlink frequency.

radio operating on a distinct downlink frequency;	Each access point contains at least two radios supporting distinct 5GHz and 2.4 GHz frequencies for backhaul (*e.g.*, uplink or downlink) and client access (*e.g.*, downlink). If Universal Backhaul Access is enabled, both 5 GHz and 2.4 GHz can be used for client access as well as backhaul. Enterprise 11n/ac mesh is added to the Catalyst 9800 controller features to work with the 802.11n/ac APs. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but add higher backhaul and client access speeds. The 802.11ac indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP, and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4-GHz radios in Release 17.1 can be used for local (client) access as well as backhaul. Enterprise 11ac mesh supports point-to-point, point-to-multipoint, and mesh types of architectures. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh .The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios: • From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable • 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands. Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/technotes/8-5/cisco-enterprise-mobility-design-guide-8-5.pdf Mesh Serial Backhaul and Universal Backhaul Access allow for the radios either operate at 5 GHz or 2.4 GHz or at different channels, *e.g.*, distinct uplink and downlink frequencies as shown below:



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Furthermore, Mesh Serial Backhaul and Universal Backhaul Access allow for distinct frequencies or channels to be assigned to different slots designated for Uplink, Downlink, or Client Access.



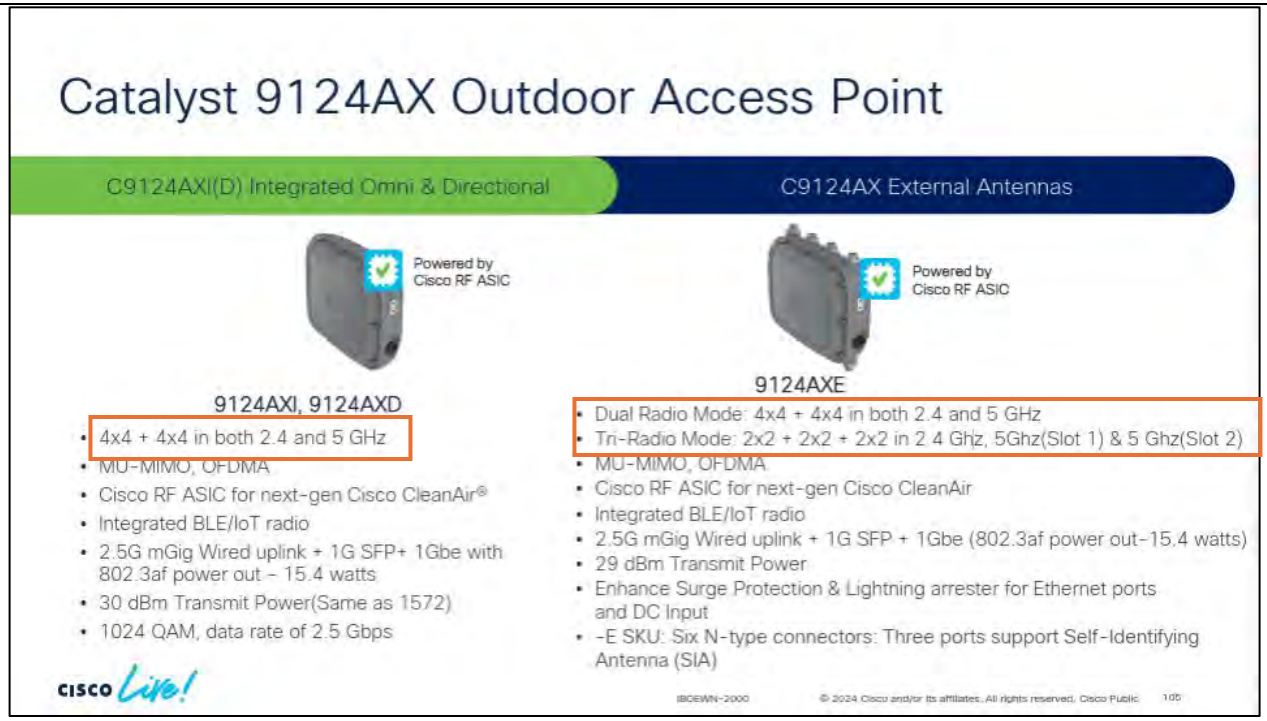
Source: https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/IBOEWN-2000.pdf.

Information About Mesh Serial Backhaul

The Mesh Serial Backhaul feature in a mesh access point (MAP), allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. This allows the in-bound and out-bound traffic to flow through exclusive communication channels, thereby improving performance and avoiding problems associated with a shared access medium.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html.

The accused products support different radio modes (*e.g.*, dual-radio, tri-radio, and quad-radio modes), in which an uplink radio can be assigned to operate on one frequency and a downlink radio can be assigned to a distinct frequency.



Source: https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/IBOEWN-2000.pdf.

C9124AXE – Radio Mode 17.7+

Mode	2.4 Ghz (Slot 0)	5 Ghz (Slot 1)	5 Ghz (Slot 2)
Dual Radio	4x4	4x4	Down
Tri-Radio	2x2	2x2	2x2



Note: To enable Tri-Radio Mode, C9124AXE should be powered up by 802.3bt/UPOE/DC

Source: https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/IBOEWN-2000.pdf.

Flexible Radio Assignment (FRA) with quad-radio mode (with future software SW upgrade)

FRA allows the access points to intelligently determine the operating mode of serving radios based on the RF environment and traffic demands. The access points can operate in the following modes:

- Tri-radio mode: One 4x4 6-GHz radio, one 8x8 5-GHz radio, and one 4x4 2.4-GHz radio.
- Quad-radio mode: One 4x4 6-GHz radio, Dual 4x4 5-GHz radios, and one 4x4 2.4-GHz radio. With four total 4x4 radios (quad-radio) inside the access point, client device capacity can be increased on demand.

The access point's default mode is tri-radio with 8x8 5-GHz and 4x4 for 6-GHz and 2.4-GHz. It has the ability to split the 8x8 radio into two separate 4x4 5-GHz radios through a future software upgrade, thereby enabling the benefits of FRA while allowing the 2.4- and 6-GHz radios to remain active.

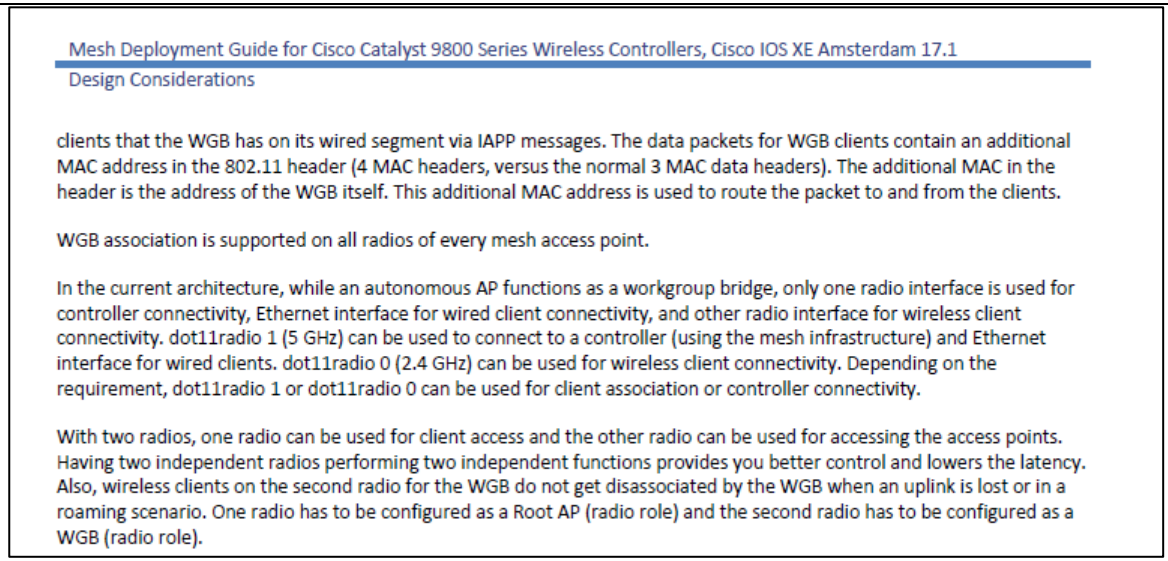
Source: https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9100ax-access-points/nb-06-cat9136-access-point-ds-cte-en.html?oid=dstwls028547.

AWPP will scan and switch to a parent off-channel to find the best path back to the RAP; thus, AWPP will look to connect (*i.e.*, uplink) with nodes operating on other frequencies.

	Parent Loss Detection/ Keep Alive Timers	**Channel Scan/Seek**	**DHCP/CAPWAP Information**	**Time per hop (sec)**
Standard	21 / 3 sec	Scan/Seek all 2.4 and 5 Ghz Channels	Renew / Restart CAPWAP	48.6*
Fast	7 / 3 sec	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	20.5*
Very Fast	4 / 1.5 sec	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	15.9*
CCN/BG Scan Fast/VF	4 / 3 sec for 50ms	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	8-10 sec

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

Cisco documentation confirms that support for using separate radios for uplink (mesh/controller connectivity) and downlink (client access) existed at least as early as IOS-XE Amsterdam 17.1.



Mesh Deployment Guide for Cisco Catalyst 9800 Series Wireless Controllers, Cisco IOS XE Amsterdam 17.1
Design Considerations

clients that the WGB has on its wired segment via IAPP messages. The data packets for WGB clients contain an additional MAC address in the 802.11 header (4 MAC headers, versus the normal 3 MAC data headers). The additional MAC in the header is the address of the WGB itself. This additional MAC address is used to route the packet to and from the clients.

WGB association is supported on all radios of every mesh access point.

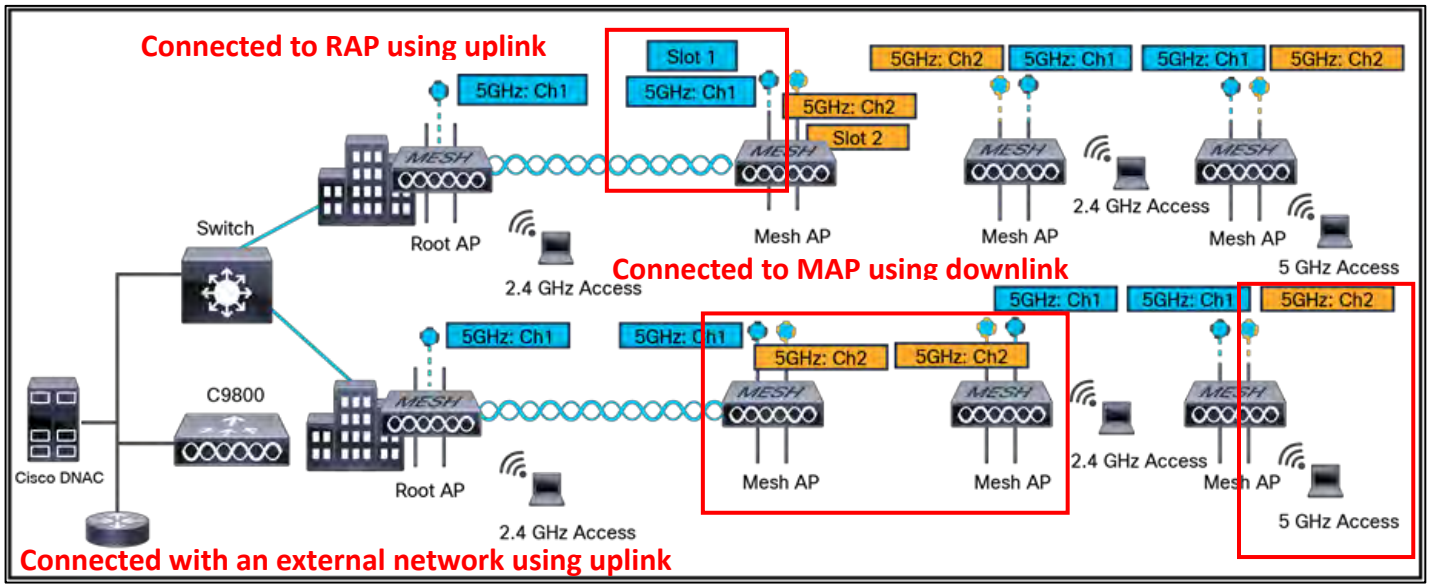
In the current architecture, while an autonomous AP functions as a workgroup bridge, only one radio interface is used for controller connectivity, Ethernet interface for wired client connectivity, and other radio interface for wireless client connectivity. dot11radio 1 (5 GHz) can be used to connect to a controller (using the mesh infrastructure) and Ethernet interface for wired clients. dot11radio 0 (2.4 GHz) can be used for wireless client connectivity. Depending on the requirement, dot11radio 1 or dot11radio 0 can be used for client association or controller connectivity.

With two radios, one radio can be used for client access and the other radio can be used for accessing the access points. Having two independent radios performing two independent functions provides you better control and lowers the latency. Also, wireless clients on the second radio for the WGB do not get disassociated by the WGB when an uplink is lost or in a roaming scenario. One radio has to be configured as a Root AP (radio role) and the second radio has to be configured as a WGB (radio role).

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's mesh solution achieves the same function (*e.g.*, separating uplink and downlink communication paths), by substantially the same means (*e.g.*, assigning non-overlapping 5GHz sub-bands such as 5.15 GHz and 5.8 GHz for uplink and downlink use), and yields the same result (*e.g.*, enabling duplex communication and reducing interference). Even if both radios operate within a broader 2.4GHz or 5GHz band, this is an insubstantial difference and functionally equivalent to the claim requirement. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1(c): wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node using its uplink radio and client devices or other mesh nodes connect to each node using each node's downlink radio;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a connectivity logic, wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node using its uplink radio and client devices or other mesh nodes connect to each node using each node's downlink radio. For example, as shown below, and in the Cisco documentation, the connectivity logic in Cisco's Catalyst 9100 and Aironet Access Points determine whether each access point connects with the external network (*e.g.*, is a MAP connecting to an external network via a RAP and controller) or another access point (*e.g.*, is a MAP connecting to a RAP or another MAP) using the uplink radio. Similarly, the logic determines when and whether the downlink radio is connected to another mesh node (*e.g.*, MAP) or a client device (*e.g.*, a smartphone or laptop).  Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh. The connectivity logic (*e.g.*, AWPP) determines the best path back to the controller for the mesh network, and therefore determines whether a mesh node connects to the controller or a RAP or another MAP using its uplink or downlink radio:

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

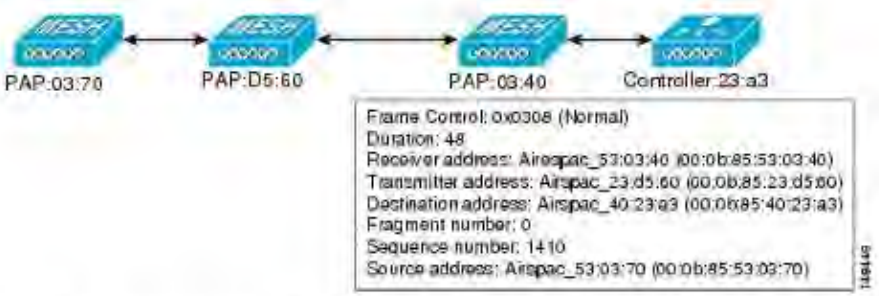
To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Additionally, the connectivity logic (*e.g.*, AWPP) determines whether each structured mesh node connects with an external network or another node on a hop-by-hop basis. The connectivity logic further determines whether an access point, through its uplink radio, connects to a Root Access Point (RAP) thereby providing connectivity to the external network (*e.g.*, through a WAN link), or instead connects to another MAP node serving as its parent.

Figure 4: Wireless Mesh Frame



As this frame is sent, the transmitter and receiver addresses change on a hop-by-hop basis. AWPP is used to determine the receiver address at each hop. The transmitter address is known because it is the current mesh access point. The source and destination addresses are the same over the entire path.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

The connectivity logic (*e.g.*, AWPP) may also determine whether each node connects with an external network to exit the isolated configuration (*e.g.*, move from standalone mode to connected mode after a connection is lost and restablished):

- After moving back to connected mode, from standalone mode, best neighbor selection timer and BGN timer will be restarted, so allow the child mesh AP to roam to the best possible neighbor.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

The mesh system supports various user-defined and preset configurations for backhaul and client communications (*e.g.*, Mesh Serial Backhaul, Universal Backhaul Access) enabling the connectivity logic to determine whether each structured mesh node (*e.g.*, RAP or MAP) connects to an external network or another mesh node using its uplink radio, and whether client devices or other mesh nodes connect to each node using its downlink radio.

Information About Mesh Serial Backhaul

The Mesh Serial Backhaul feature in a mesh access point (MAP), allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. This allows the in-bound and out-bound traffic to flow through exclusive communication channels, thereby improving performance and avoiding problems associated with a shared access medium.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html.

Enterprise 11n/ac mesh is added to the Catalyst 9800 controller features to work with the 802.11n/ac APs. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but add higher backhaul and client access speeds. The 802.11ac indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP, and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4-GHz radios in Release 17.1 can be used for local (client) access as well as backhaul. Enterprise 11ac mesh supports point-to-point, point-to-multipoint, and mesh types of architectures.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

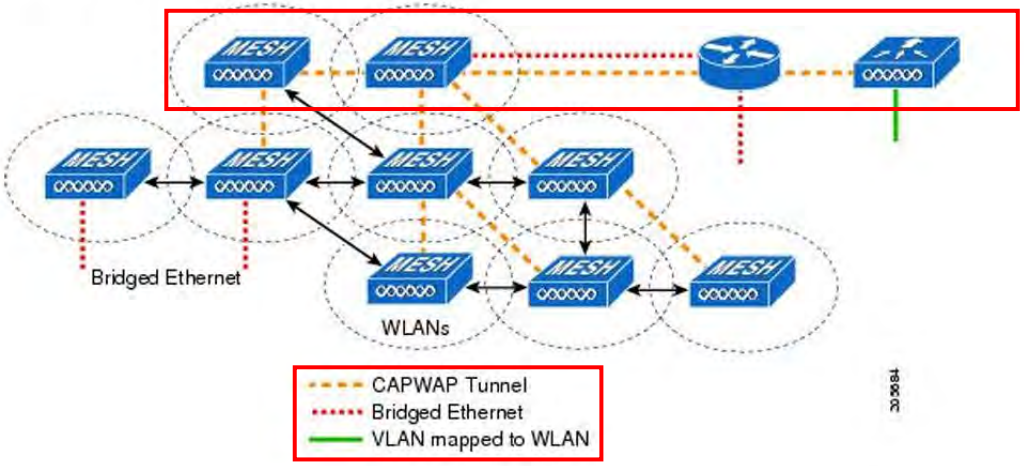
.The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios:

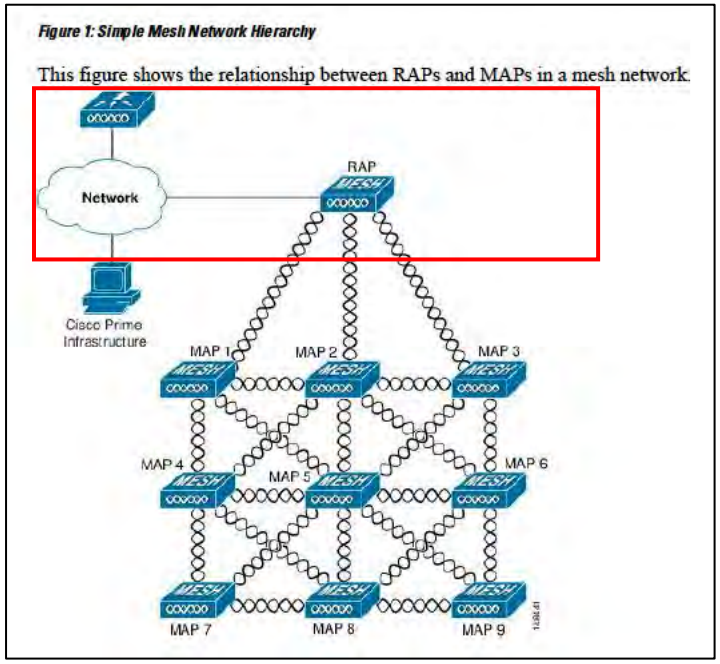
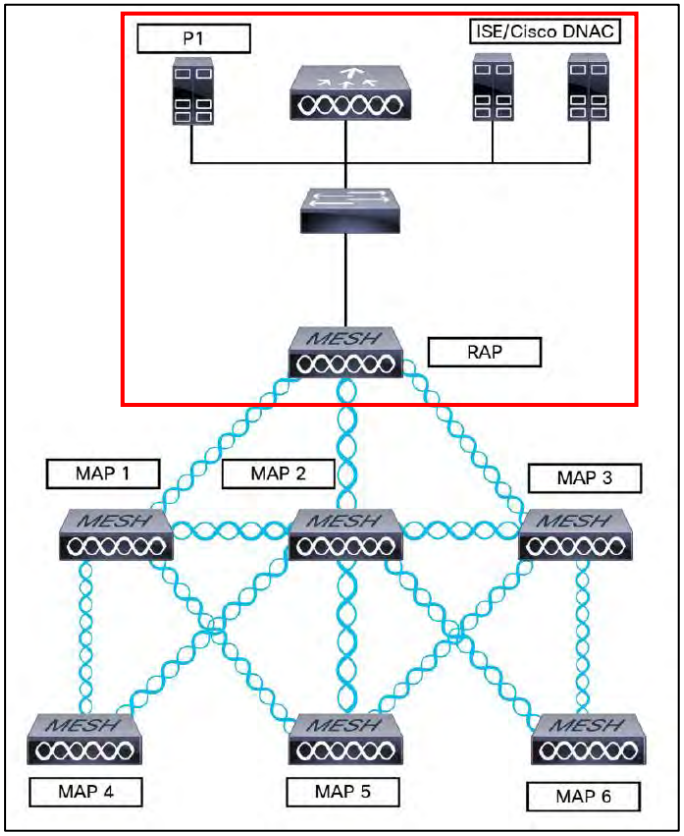
- From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable
- 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled

The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/technotes/8-5/cisco-enterprise-mobility-design-guide-8-5.pdf.

Access points can maintain operation and locally switch client traffic when the connection to the controller is lost, and that they are configured and controlled from the corporate network over a WAN link. This demonstrates that connection from an AP (including a RAP in a mesh deployment) can provide access to the corporate network through the WAN, i.e., to an external network beyond just another access point.

	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's implementation using AWPP and CAPWAP with VLAN tunneling performs the same function (*e.g.*, routing traffic between mesh nodes and the external network), in substantially the same way (*e.g.*, by dynamically selecting the next hop toward a RAP whose traffic is tunneled to a controller mapped to an external VLAN), and achieves the same result (*e.g.*, establishing and managing a connection between mesh nodes and external networks). Connection to an external network via VLAN tunneling is an insubstantial difference from determining external and mesh connections via internal logic. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. Each mesh access point within the mesh forms an CAPWAP session with a controller. WLAN traffic is encapsulated inside CAPWAP and is mapped to a VLAN interface on the controller. Bridged Ethernet traffic can be passed from each Ethernet interface on the mesh network and does not have to be mapped to an interface on the controller (see Figure 5: Logical Bridge and WLAN Mapping, on page 12). ***Figure 5: Logical Bridge and WLAN Mapping***  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1(d): wherein the structured mesh network functions in two configurations selected	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function in two configurations (*e.g.*, a "first connected configuration" wherein the connection to the external network via the

depending on whether a connection to an external network is present; in the first connected configuration the structured mesh network includes at least one structured mesh node's uplink radio comprising a connection to an external network;	controller is not lost, and an "isolated configuration" wherein the connection to the external network via the controller is lost) selected depending on whether a connection to an external network is present. In normal operation, the Cisco Wireless Mesh Network is connected to an external network (to make/facilitate additional connections to outside networks and destinations, for example to connect clients to an external network). When so connected, the system is configured to pass information from inside the network via the uplink radio of at least one structured mesh node (*i.e.*, RAP), as shown below and in Cisco's documentation:  *Figure 1: Simple Mesh Network Hierarchy* This figure shows the relationship between RAPs and MAPs in a mesh network.  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh. In some cases, the typical operating mode is explicitly referred to as "connected mode."

Connected mode

A Wave 2 Flex+Mesh AP (RAP or child MAP) is in connected mode when it can access and join the Catalyst 9800 and can exchange periodic keep-alive messages with the Catalyst 9800. In this mode, a Flex+Mesh AP will be able to support locally and centrally switched WLANs. It will allow regular client and child mesh APs to join.

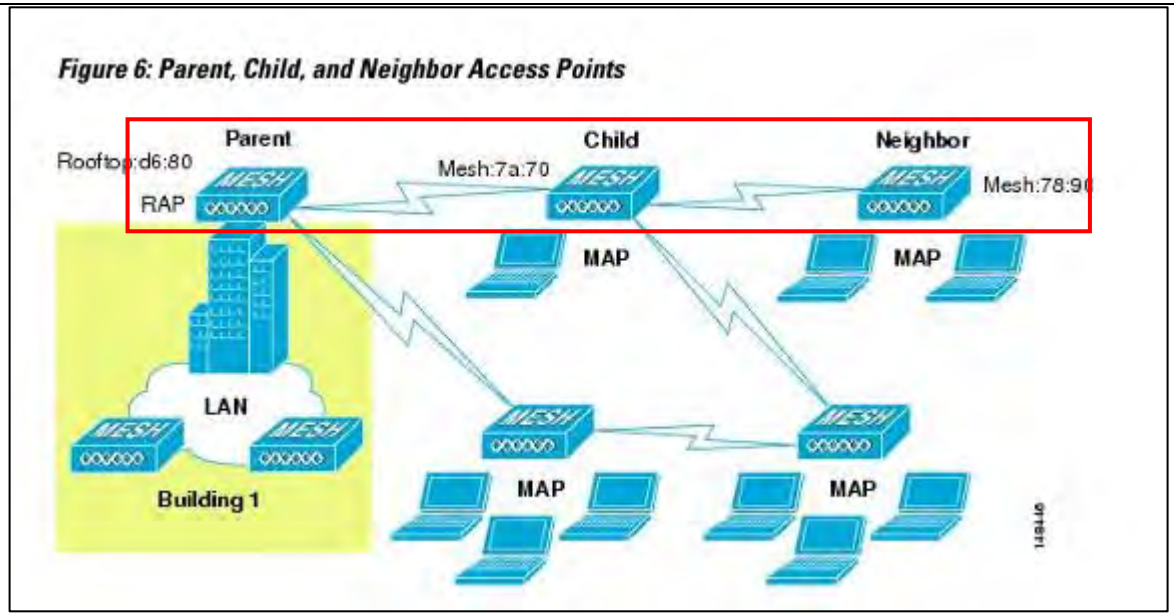
Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

The mesh nodes have a wireless connection to the controller (and the external network) via the RAP:

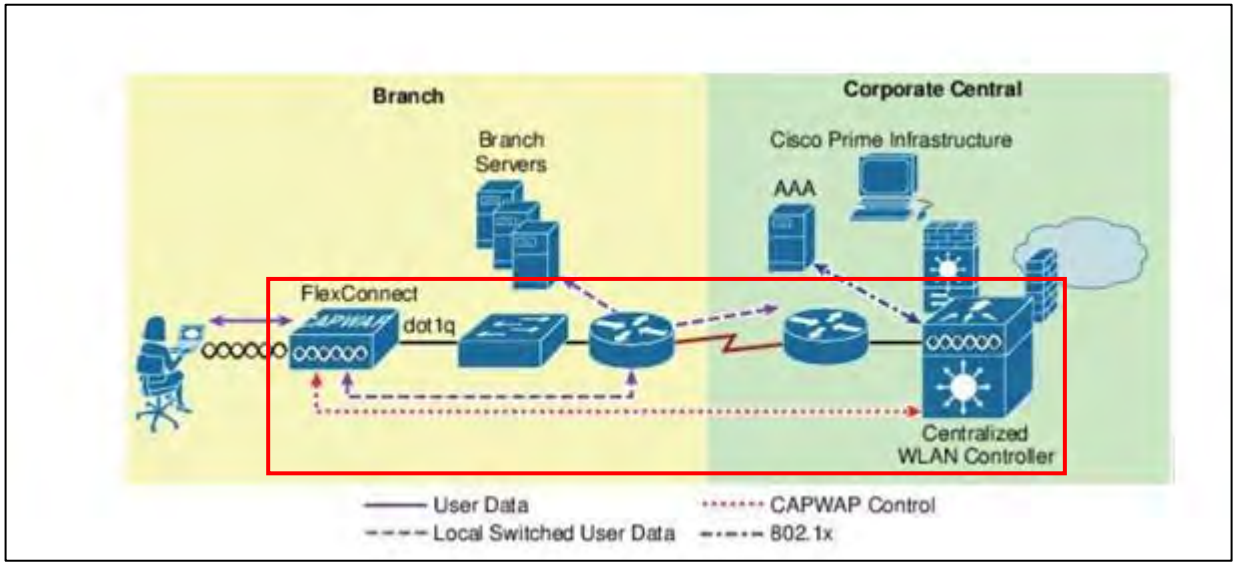
While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's system achieves the same function (external network connectivity), in substantially the same way (via the mesh node's uplink radio establishing a path to a controller), to yield the same result (connectivity beyond the local mesh network). Whether or not the controller itself is considered part of the "external network," the resulting connectivity functionally extends beyond the mesh. Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (radio-based upstream communication and/or tunneling) are insubstantially different, and the manner of establishing that connectivity (through AWPP-determined routing and/or controller tunneling) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.



Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.



Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1(e): and in the second isolated configuration none of the structured mesh nodes' uplink radio comprises a connection to an external network, and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node as isolated root children nodes forming a tree configuration;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function in the second isolated configuration (*e.g.*, fault tolerant resilient mode) where none of the structured mesh nodes' uplink radio comprises a connection to an external network, (*e.g.*, when Cisco's Catalyst 9100 and Aironet Access Points fails to connect to an external network via the controller) and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node (*e.g.*, a Cisco's Catalyst 9100 and Aironet Access Points operating as a RAP prior to a loss of connection to the controller) as isolated root children nodes (*e.g.*, a Cisco's Catalyst 9100 and Aironet Access Points operating as MAPs) forming a tree configuration. When the Cisco Wireless Mesh Network is not connected to an external network, the network is capable of functioning in a variety of localized modes. Cisco's Catalyst 9100 and Aironet Access Points are access points that can operate in several different modes, including, for example, Flex + Mesh mode, fault tolerate resistance mode, and standalone mode. Each of these modes maintain the connectivity of mesh access points when the connection to the controller (and thus to the external network) is lost. These modes can switch client data traffic locally when a connection to the controller is lost. As shown below and in Cisco's documentation, the tree structure is maintained, as a child access point (MAP) remains connected to a parent access point (another MAP or a RAP) in the second isolated configuration. **Cisco outdoor mesh access points** In addition to the mesh mode, the mesh APs can operate in the following modes: • **Local mode:** In this mode, the AP can handle clients on its assigned channel or while monitoring all channels on the band over a 180-second period. During this time, the AP listens on each channel for 50 ms for rogue client beacons, noise floor measurements, interference, and IDS events. The AP also scans for CleanAir interference on the channel. • **FlexConnect mode:** FlexConnect mode enables you to configure and control APs in a branch or remote office from the corporate office through a WAN link without having to deploy a controller in each office. This mode can switch client data traffic locally and perform client authentication locally when the connection to the controller is lost. When connected to the controller, an AP in FlexConnect mode can also tunnel traffic back to the controller. • **Flex+Mesh mode:** In this mode, both the FlexConnect and Bridge mode configuration options are available on the access point. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Connected mode

A Wave 2 Flex+Mesh AP (RAP or child MAP) is in connected mode when it can access and join the Catalyst 9800 and can exchange periodic keep-alive messages with the Catalyst 9800. In this mode, a Flex+Mesh AP will be able to support locally and centrally switched WLANs. It will allow regular client and child mesh APs to join.

Standalone mode

A Wave 2 Flex+Mesh AP is in standalone mode if it loses connection to the controller but can access the local gateway. In this mode, the Wave 2 Flex+Mesh AP will disable all the centrally switched WLANs and keep the locally switched WLANs up and running. It will also allow new clients to join on locally switched WLANs using local authentication if the authentication server is reachable in the local network. Child mesh APs will NOT be allowed to join in this mode.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

uplink to the core network in a 'Root' AP or RAP mode. The AP is still managed by a centralized controller over CAPWAP. The AP however is capable of moving to a standalone mode where the AP will be able to continue to serve the 802.11 clients depending on the data switching method of the WLAN configured on the AP. The data can be centrally or locally switched. When data is centrally switched, all data is sent to the WLC that does the switching further. In a local switch network, the data is sent to RAP where RAP switches locally on the wired uplink. There is no difference to the central and local switched WLAN configurations and

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

A Control and Provisioning of Wireless Access Points protocol (CAPWAP) Access Point (AP) can be configured to operate in two different modes:

- FlexConnect mode
- Bridge/Mesh mode

The following are the bridging features for Flex+Bridge mode:

- The Flex+Bridge mode supports the centrally switched 802.11 WLAN. Traffic for this tunneled WLAN is forwarded to and from a CAPWAP controller over an IP tunnel.
- The Flex+Bridge mode supports the Root Ethernet VLAN Bridging. A root AP bridges the traffic for bridged 802.11 WLANs and secondary Ethernet LANs to a local Ethernet LAN over its root Ethernet port.
- The Flex+Bridge mode bridging is supported on Secondary Ethernet Access Ports and Secondary Ethernet VLAN Trunk Ports.
- Fault Tolerant Resilient Mode enables an AP to continue bridging traffic when the connection to the CAPWAP controller is lost. Both mesh and non-mesh root APs continue to bridge traffic. A child mesh AP (MAP) maintains its link to a parent AP and continues to bridge traffic till the parent link is lost. A child mesh AP cannot establish a new parent or child link till it reconnects to the CAPWAP controller. Existing wireless clients on the locally switching WLAN can stay connected with their AP in this mode. Their traffic will continue to flow through the Mesh and wired network. No new or disconnected wireless client can associate to the Mesh AP in this mode.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/flexconnect.html.

FlexConnect APs continue to operate locally "even in the event of … WAN failure." A loss of WAN connectivity isolates the AP from the external network while preserving local functionality. Accordingly, Standalone mode represents an isolated configuration in which external connectivity is lost, not maintained.

FlexConnect modes of operation

Table 3. Flexconnect modes of operation

FlexConnect mode	Description
Connected	FlexConnect is said to be in connected mode when its CAPWAP control plane back to the controller is up and operational, meaning the WAN link is not down.
Standalone	Standalone mode is specified as the operational state the FlexConnect enters when it no longer has the connectivity back to the controller. FlexConnect APs in standalone mode will continue to function with last known configuration, even in the event of power failure and WLC or WAN failure.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/technical-reference/cat-9800-flexconnect-branch-deployment-guide-og.html.

FlexConnect Feature Matrix - Legacy and New Features in Release 7.0.116 and Later

Security - Client

Security support on FlexConnect varies with different modes and states. This table summarizes the security features that are supported:

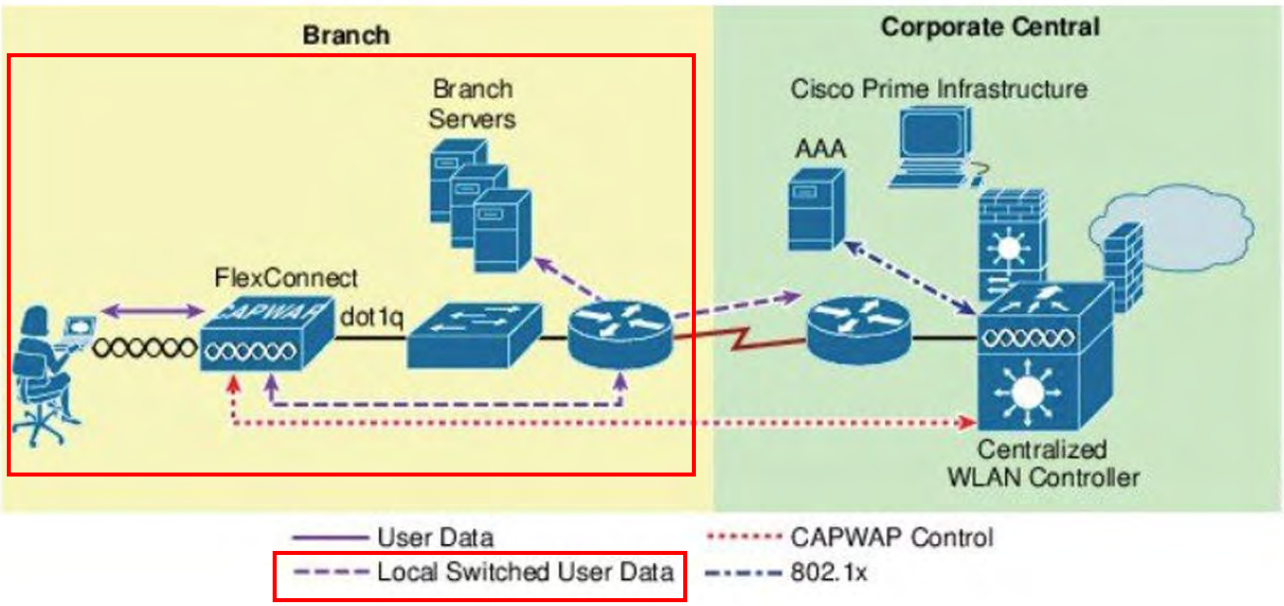
	WAN Up (Central Switching)	WAN Up (Local Switching)	WAN Up (Local Switching, Local Authentication)	WAN Down (Standalone)
Open/Static WEP	Yes	Yes	Yes	Yes
WPA-PSK	Yes	Yes	Yes	Yes
802.1x (WPA/WPA2)	Yes	Yes	Yes	Yes
MAC filter Authentication	Yes	Yes	No	No
CCKM Fast Roaming	Yes	Yes	No	Yes, for connected clients. No, for new clients.

Security - Infrastructure

	WAN Up (Central Switching)	WAN Up (Local Switching)	WAN Down (Standalone)
Data DTLS Encryption	Yes	N/A	N/A
Local EAP (7.0 to 7.4)	Yes (LEAP/EAP-FAST)	Yes (LEAP/EAP-FAST)	Yes (LEAP/EAP-FAST)
LocaL EAP (7.5 and later)	Yes (LEAP/EAP-FAST/PEAP/EAP-TLS)	Yes (LEAP/EAP-FAST/PEAP/EAP-TLS)	Yes (LEAP/EAP-FAST/PEAP/EAP-TLS)
Backup Radius	Yes (7.0.116)	Yes (7.0.116)	Yes
MIC	Yes	Yes	Not applicable

	Source: https://www.cisco.com/c/en/us/support/docs/wireless/5500-series-wireless-controllers/112042-technote-wlc-00.html. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's system achieves the same function (maintaining mesh connectivity when the connection to the external network is lost), in substantially the same way (via a mesh node acting as parent or root node), to yield the same result (continuing to bridge traffic with wireless clients staying connected). Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (radio-based upstream communication and/or tunneling) are insubstantially different, and the manner of establishing that connectivity (through AWPP-determined routing and/or controller tunneling) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1(f):wherein clients of nodes of the structured network in a second isolated configuration retain full connectivity within the structured network during the isolated configuration and wherein said client connectivity is uninterrupted during any transition from the isolated configuration to the first connected configuration wherein distinct downlink frequencies are used by nodes for communication with one or more nodes lower on the tree configuration.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function in the second isolated configuration wherein clients of nodes of the structured network in a second isolated configuration retain full connectivity (*e.g.*, when the locally switched WLANs are up and running despite the loss of a connection to the controller) within the structured network during the isolated configuration (*e.g.*, fault tolerant resilient mode) and wherein said client connectivity is uninterrupted during any transition from the isolated configuration to the first connected configuration; wherein distinct downlink frequencies are used by nodes for communication with one or more nodes lower on the tree configuration. As shown below, and in the Cisco documentation, client data traffic is switched locally when a connection to the controller is lost. Clients stay connected with their AP and traffic will continue to flow. The AP moves to standalone mode and will continue to serve clients: Below is a general FlexMesh architecture. The CAPWAP AP is in Flex connect + Bridge mode with a wired uplink to the core network in a 'Root' AP or RAP mode. The AP is still managed by a centralized controller over CAPWAP. The AP however is capable of moving to a standalone mode where the AP will be able to continue to serve the 802.11 clients depending on the data switching method of the WLAN configured on the AP. The data can be centrally or locally switched. When data is centrally switched, all data is sent to the WLC that does the switching further. In a local switch network, the data is sent to RAP where RAP switches locally on the wired uplink. There is no difference to the central and local switched WLAN configurations and functionality between a Flex-Connect and Flex+Mesh mode AP. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.



Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

As shown above, a child access point (MAP) remains connected to a parent access point (another MAP or a RAP).

As shown below, the distinct downlink frequencies used by nodes for communication with one or more nodes lower on the tree configuration are enabled through Cisco features such as Mesh Serial Backhaul, and Universal Backhaul Access, which assign different frequencies on a per-hop basis,

Information About Mesh Serial Backhaul

The Mesh Serial Backhaul feature in a mesh access point (MAP), allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. This allows the in-bound and out-bound traffic to flow through exclusive communication channels, thereby improving performance and avoiding problems associated with a shared access medium.

The Mesh Serial Backhaul feature is supported in the controller from Cisco IOS XE Cupertino 17.7.1 onwards, for Cisco Catalyst 9124AXE outdoor APs. A new knob is introduced under the radio profile, and that radio profile is associated with a radio frequency (RF) tag to enable the Mesh Serial Backhaul feature. When you enable this feature, the mesh configuration is shared by all the APs that share the same mesh profile. Radio configuration is shared by all the APs that are configured with the same radio profile.

Basic client access functionality is offered on the 2.4-GHz radio and the 5-GHz radio, which are not used in serial backhaul. Universal access is made available on the downlink radio.

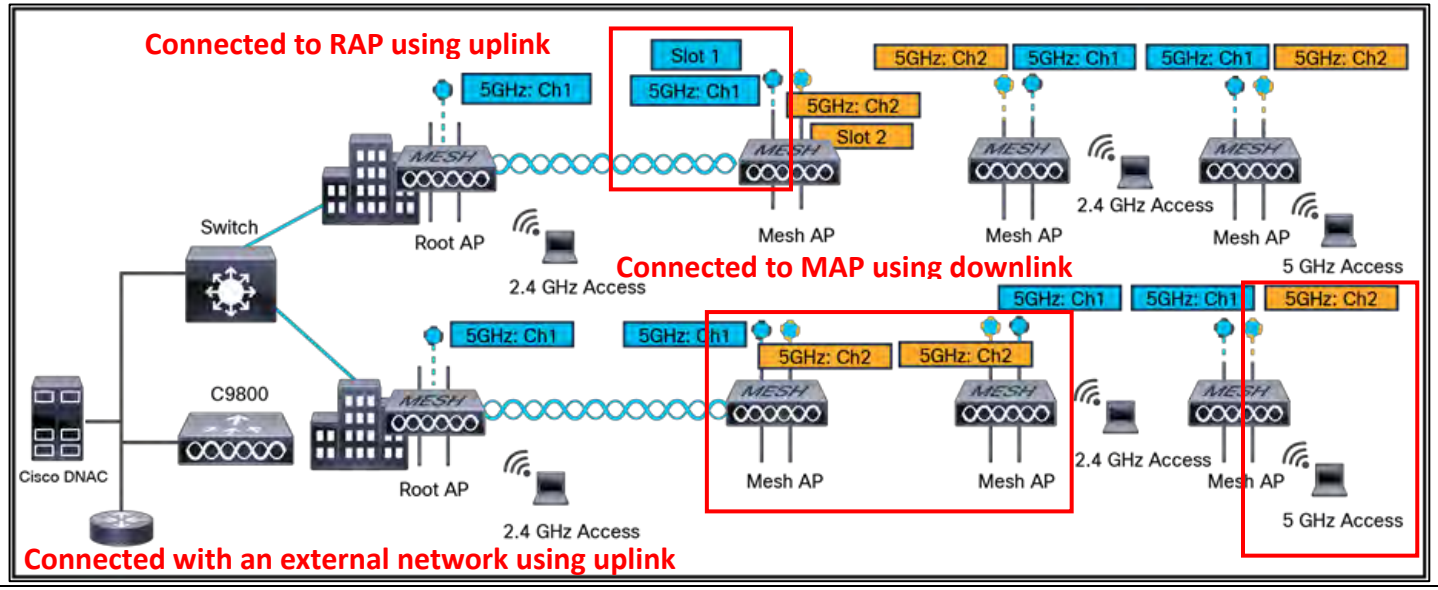
Note Slot 1 and slot 0 are supported as mesh backhaul. Slot 2 is not supported as a mesh backhaul. You can utilize slot 2 for client serving or for serial backhaul downlink.

Channel Assignment

For the Mesh Serial Backhaul feature, channels are assigned according to the following rules:

- Uplink and downlink channels are different.
- All the 5-GHz radios maintain a frequency guard between their operating channels. For example, 100-MHz channel spacing between radios in Cisco Catalyst 9124AXE outdoor APs.
- Dynamic Frequency Selection (DFS) channels are supported.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html.



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Furthermore, Mesh Serial Backhaul and Universal Backhaul Access allow for distinct frequencies or channels to be assigned to different slots designated for Uplink, Downlink, or Client Access.

Information About Mesh Serial Backhaul

The Mesh Serial Backhaul feature in a mesh access point (MAP), allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. This allows the in-bound and out-bound traffic to flow through exclusive communication channels, thereby improving performance and avoiding problems associated with a shared access medium.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html.

Mesh – Designated Downlink



- Mesh Serial Backhaul(SBH) allows us to maximize throughput over multiple mesh hops
- Allows network segregation.
- Enables us to use different channels on different backhauls, which helps in avoiding localized link interferences
- In above deployment, Slot 1 is used for Uplink, Slot 0 is for Client Access and Slot 2 is the designated Downlink

Source: https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/IBOEWN-2000.pdf.

On information and belief, the client connectivity is uninterrupted during any transition from an isolated configuration back to the first connected configuration. Resiliency is a design feature of the Cisco Wireless Mesh Network. Clients do not disconnect from an access point if the connection to the controller and the external network fails.

High availability

High availability has been a requirement on WLCs to minimize downtime in live networks. This section describes the theory of operation and configuration for Catalyst 9800 Series WLCs as it pertains to supporting stateful switchover of Access Points and clients (AP and client SSO). The Catalyst 9800 Series is the next generation in wireless controllers and can run on multiple platforms with different scalability goals, from low to high scale. AP and client SSO is supported on the physical appliances and the virtual cloud platforms of the Catalyst 9800 Series, namely the 9800-L, 9800-40, 9800-80, and 9800-CL. The underlying SSO functionality is the same on all platforms, with some differences in the setup process.

The high availability SSO capability on WLCs allows the AP to establish a CAPWAP tunnel with the active WLC and the active WLC to share a mirror copy of the AP and client database with the standby WLC. The APs do not go into the discovery state, and clients do not disconnect when the active WLC fails and the standby WLC takes over the network as the active WLC. Only one CAPWAP tunnel is maintained at a time between the AP and the wireless controller that is in an active state. Release 16.10 supports full AP and client SSO. Client SSO is supported for clients that have already completed the authentication and DHCP phase and have started passing traffic. With client SSO, a client's information is synced to the standby WLC when the client associates to the WLC or the client's parameters change. Fully authenticated clients, that is, the ones in the run state, are synced to the standby WLC, and thus client reassociation is avoided on switchover, making the failover seamless for the APs as well as for the clients, resulting in zero client service downtime and zero SSID outage. The overall goal in adding AP and client SSO support to the Catalyst 9800 Series WLCs is to reduce major downtime in wireless networks due to failure conditions that may occur due to box failover, network failover, or a power outage at the primary site.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's system achieves the same function (maintaining mesh connectivity when the connection to the external network is lost), in substantially the same way (via a mesh node acting as parent or root node), to yield the same result (continuing to bridge traffic with wireless clients staying connected). Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (radio-based upstream communication and/or tunneling) are insubstantially different, and the manner of establishing that connectivity (through AWPP-determined routing and/or controller tunneling) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.

	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 9	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
The structured mesh network of claim 1 wherein the connectivity logic contained by the nodes realigns connections between the nodes upon detection of an external network connection to form the connected configuration.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function in the structured mesh network of claim 1 wherein the connectivity logic contained by the nodes realigns connections between the nodes upon detection of an external network connection to form the connected configuration. In addition to the elements in claim 1 above, AWPP continually monitors and scans conditions (*e.g.*, addition or removal of RAPs, ease values, etc.) and changes routes to reflect changes in conditions such as detection of an external network connection (*e.g.*, establishing a path to the CAPWAP controller).

Claim 9	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

AWPP uses ease values to realign connections between the nodes to offer the best route back to the RAP.

Claim 9	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
  - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.
- Mesh networks are half duplex, meaning after the first hop (RAP to MAP) overall throughput is decreased by 50% for each additional hop (MAP to MAP). Where Ethernet bridging clients are used in MAPs and heavy traffic is passed, high throughput consumption nay result, which may cause the downlink MAPs to disassociate from the network due to throughput starvation.

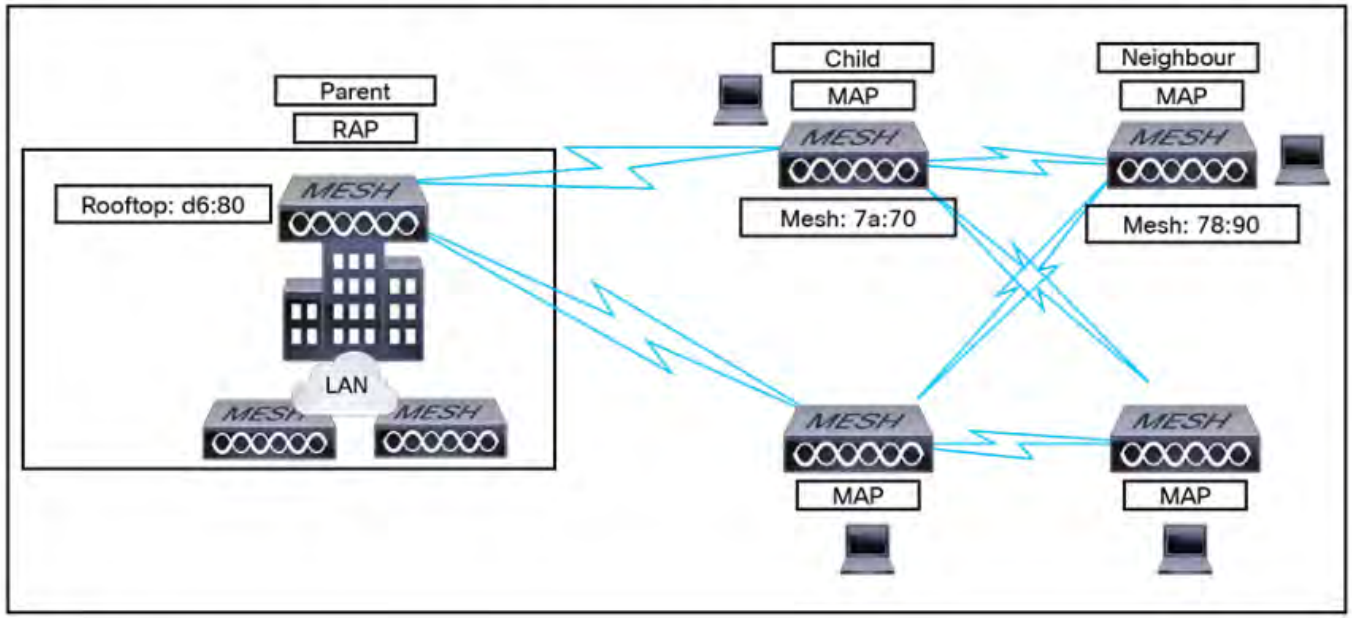
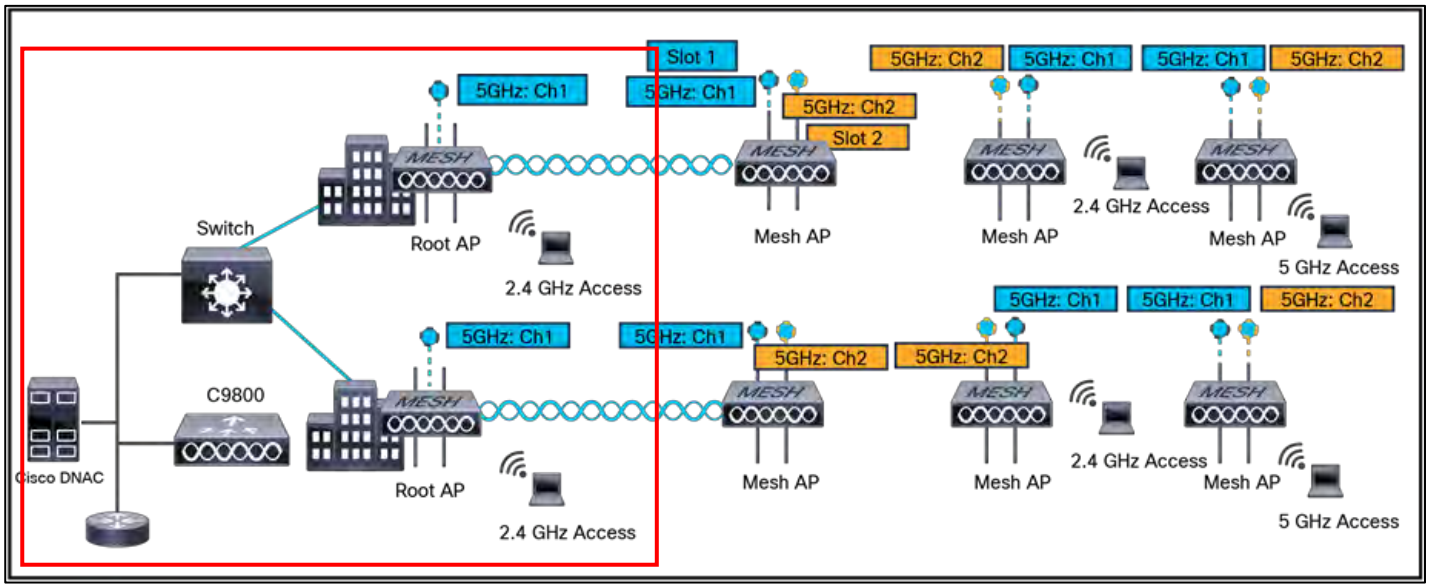


Figure 73.

Parent, child, and neighbor relationships

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Claim 9	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's mesh network achieves the same function (*e.g.*, restructuring network topology in response to external connectivity), by substantially the same way (*e.g.*, using AWPP and ease-based path selection to reassign parent-child relationships and rebuild the mesh tree toward a RAP), and yields the same result (*e.g.*, creating a converged mesh topology linked to an external network). The use of dynamic ease-based reassignment is an insubstantial difference. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 10	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
The structured mesh network of claim 9 wherein solely the isolated network root node logic searches for the external network connection and establishes the external network connection.	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network functions in the structured mesh network of claim 9 wherein solely the isolated network root node logic searches for the external network connection and establishes the external network connection. In the first configuration, the root node (*e.g.*, the Root AP (RAP)) is directly connected to the external network. 

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

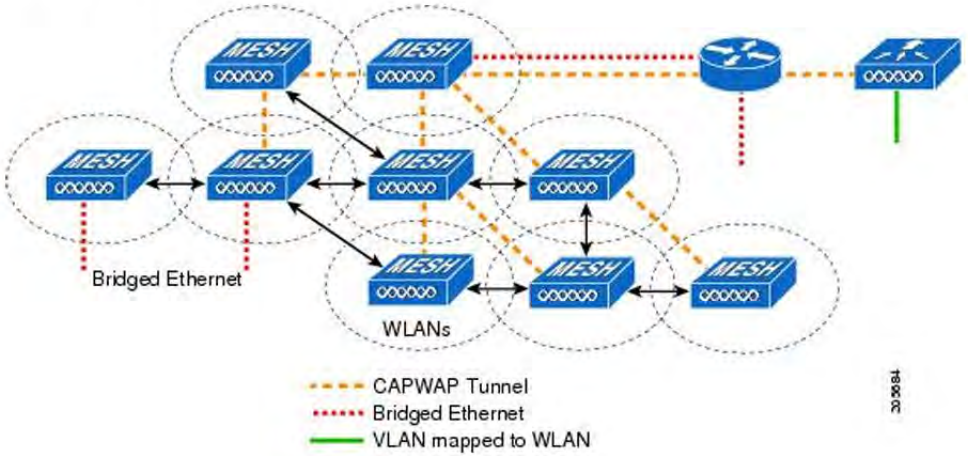
Only nodes connected to the network via Ethernet uplink are designated as the Root AP (RAP) and would search to establish the external network connection.

Mesh background scanning

Cisco MAPs are interconnected over wireless links in a spanning tree–like topology. A MAP connected to the network via Ethernet uplink is designated as the root AP. AWPP is used to maintain and form the tree. When a MAP comes up, it tries to look for another MAP (parent) to join to reach the gateway eventually via a RAP. The same happens when a MAP lose connectivity with its existing parent. This procedure is known as mesh tree convergence. This feature aims to make the convergence procedure faster and more robust.

This procedure is time consuming. To simplify the process, we have introduced mesh background scanning and automatic parent selection. This feature is available in Cisco IOS XE starting with Release 17.11.1 in the Aironet 1562 and Catalyst 9124AX APs. This mechanism allows the MAP to find and connect faster to a better potential parent across channels and always maintain its uplink with the best parent.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

	Each mesh access point within the mesh forms an CAPWAP session with a controller. WLAN traffic is encapsulated inside CAPWAP and is mapped to a VLAN interface on the controller. Bridged Ethernet traffic can be passed from each Ethernet interface on the mesh network and does not have to be mapped to an interface on the controller (see Figure 5: Logical Bridge and WLAN Mapping, on page 12). ***Figure 5: Logical Bridge and WLAN Mapping***  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's RAP (Root Access Point) performs the same function (*e.g.*, initiating discovery of and establishing connectivity with an external controller), by substantially the same means (*e.g.*, using CAPWAP discovery via a static IP or assigned controller), and achieves the same result (*e.g.*, re-establishing external network communication for the mesh). The RAP remains the sole node initiating and negotiating the external connection. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

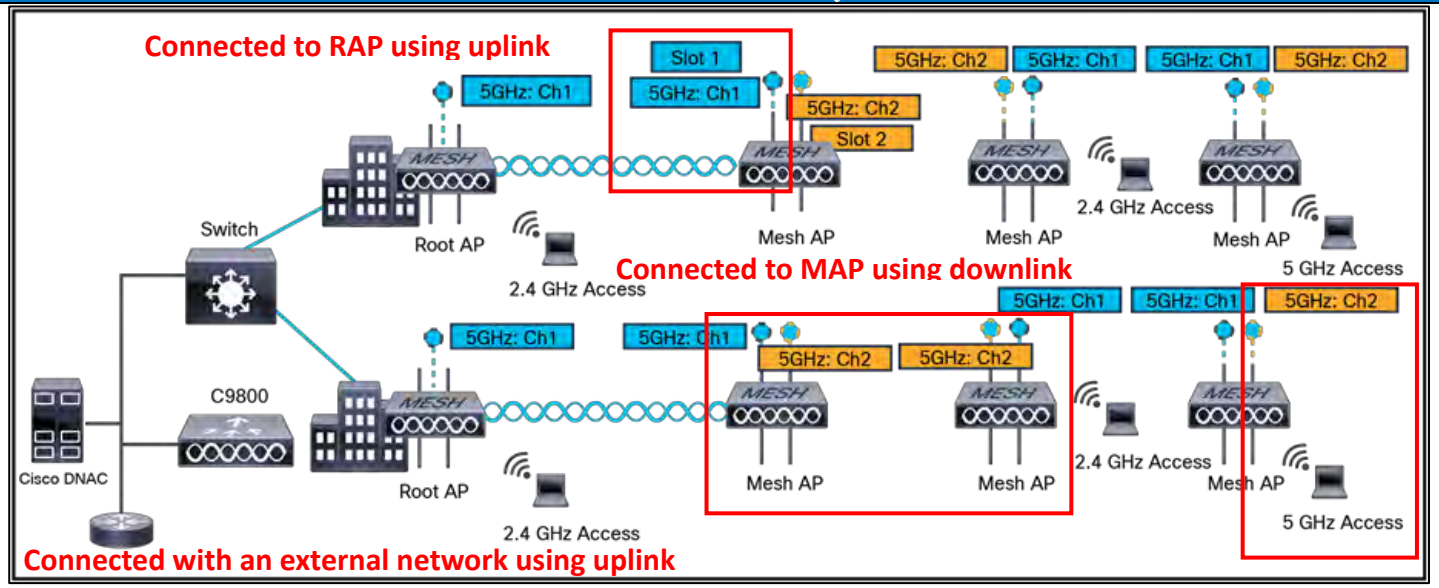
Claim 11	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
The structured mesh network of claim 10 wherein the previously	Cisco's Catalyst 9100 Access Points in the Cisco Wireless Mesh Network functions in the structured mesh network of claim 10 wherein the previously unused uplink radio of the isolated network root node is used to connect to the external network.

Claim 11	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
unused uplink radio of the isolated network root node is used to connect to the external network.	As shown in Cisco's documentation, when an AP operates in standalone mode, it may lose its connection and enter an isolated configuration (*e.g.*, standalone mode). In that state, the RAP's uplink radio becomes unused with respect to the external network. While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh. **Connected mode** A Wave 2 Flex+Mesh AP (RAP or child MAP) is in connected mode when it can access and join the Catalyst 9800 and can exchange periodic keep-alive messages with the Catalyst 9800. In this mode, a Flex+Mesh AP will be able to support locally and centrally switched WLANs. It will allow regular client and child mesh APs to join. **Standalone mode** A Wave 2 Flex+Mesh AP is in standalone mode if it loses connection to the controller but can access the local gateway. In this mode, the Wave 2 Flex+Mesh AP will disable all the centrally switched WLANs and keep the locally switched WLANs up and running. It will also allow new clients to join on locally switched WLANs using local authentication if the authentication server is reachable in the local network. Child mesh APs will NOT be allowed to join in this mode. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh. In standalone mode, the RAP's uplink radio becomes unused with respect to the external network. Upon detecting a viable connection, the RAP re-establishes connectivity by using its previously unused uplink radio to reconnect to the external network and moves from standalone mode to connected mode. • After moving back to connected mode, from standalone mode, best neighbor selection timer and BGN timer will be restarted, so allow the child mesh AP to roam to the best possible neighbor. Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

Claim 11	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's RAP achieves the same function (*e.g.*, restoring connectivity to the external network), by substantially the same means (*e.g.*, activating an unused uplink interface that previously lost connection), and yields the same result (*e.g.*, establishing a connection between the mesh and external networks). The use of a wired connection instead of a radio is an insubstantial difference. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
35pre: A structured mesh network capable of isolated operation, comprising:	*See* Claim 1pre above, the contents of which are incorporated herein by reference.
35(a): at least two structured mesh nodes; wherein each structured mesh node comprises at least a connectivity logic and a radio operating on an uplink frequency and a distinct downlink frequency;	*See* Claim 1(a) and 1(b) above, the contents of which are incorporated herein by reference.
35(b): wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node on the uplink frequency and client devices or other mesh nodes connect to each node on the downlink frequency;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network comprise a connectivity logic, wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node on an uplink frequency and client devices or other mesh nodes connect to each node on an downlink frequency. For example, and as shown below, and in the Cisco documentation, the connectivity logic in Cisco's Catalyst 9100 and Aironet Access Points determine whether each access point connects with the external network (*e.g.*, is a MAP connecting to an external network via a RAP and controller) or another access point (*e.g.*, is a MAP connecting to a RAP or another MAP) on an uplink frequency. Similarly, the logic determines when and whether the downlink frequency is connected to another mesh node (*e.g.*, MAP) or a client device (*e.g.*, a smartphone or laptop).

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points



Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

The connectivity logic (*e.g.*, AWPP) determines the best path back to the controller for the mesh network, and therefore determines whether a mesh node connects to the controller or a RAP or another MAP on an uplink or downlink frequency:

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

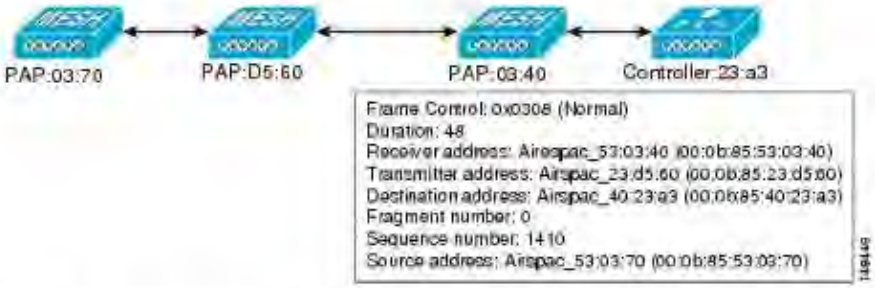
AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Additionally, the connectivity logic (*e.g.*, AWPP) determines whether each structured mesh node connects with an external network or another node on a hop-by-hop basis.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Figure 4: Wireless Mesh Frame



As this frame is sent, the transmitter and receiver addresses change on a hop-by-hop basis. AWPP is used to determine the receiver address at each hop. The transmitter address is known because it is the current mesh access point. The source and destination addresses are the same over the entire path.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

The connectivity logic (*e.g.*, AWPP) may also determine whether each node connects with an external network to exit the isolated configuration (*e.g.*, move from standalone mode to connected mode after a connection is lost and restablished):

- After moving back to connected mode, from standalone mode, best neighbor selection timer and BGN timer will be restarted, so allow the child mesh AP to roam to the best possible neighbor.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf

The mesh system supports various user-defined and preset configurations for backhaul and client communications (*e.g.*, Mesh Serial Backhaul, Universal Backhaul Access) enabling the connectivity logic to determine whether each structured mesh node (*e.g.*, RAP or MAP) connects to an external network or another mesh node on an uplink frequency, and whether client devices or other mesh nodes connect to each node on a downlink frequency.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Information About Mesh Serial Backhaul

The Mesh Serial Backhaul feature in a mesh access point (MAP), allows different channels for uplink and downlink access, thus improving backhaul bandwidth and extending universal access. One radio is used as the uplink radio and a different one is used as the downlink radio. This allows the in-bound and out-bound traffic to flow through exclusive communication channels, thereby improving performance and avoiding problems associated with a shared access medium.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html.

Enterprise 11n/ac mesh is added to the Catalyst 9800 controller features to work with the 802.11n/ac APs. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but add higher backhaul and client access speeds. The 802.11ac indoor APs are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the AP, and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4-GHz radios in Release 17.1 can be used for local (client) access as well as backhaul. Enterprise 11ac mesh supports point-to-point, point-to-multipoint, and mesh types of architectures.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh

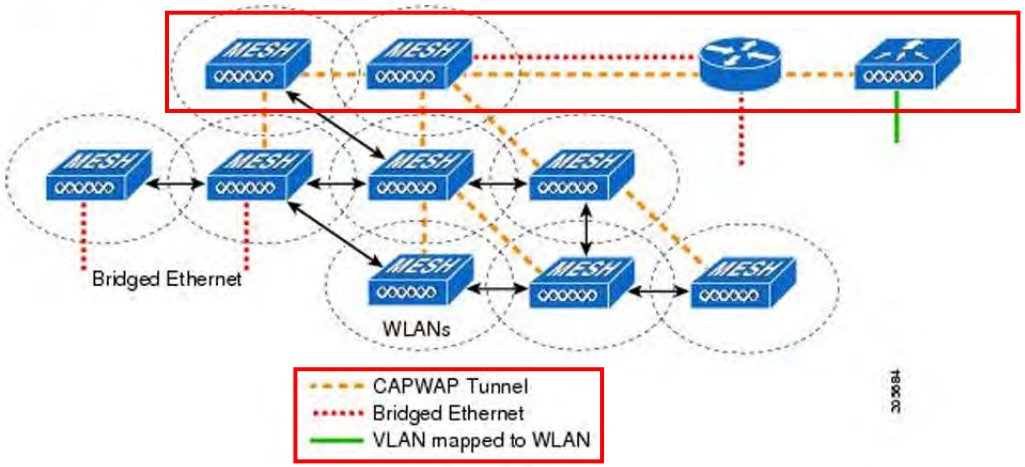
.The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios:

- From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable
- 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled

The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands.

Source: https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/technotes/8-5/cisco-enterprise-mobility-design-guide-8-5.pdf

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's implementation using AWPP and CAPWAP with VLAN tunneling performs the same function (*e.g.*, routing traffic between mesh nodes and the external network), in substantially the same way (*e.g.*, by dynamically selecting the next hop toward a RAP whose traffic is tunneled to a controller mapped to an external VLAN), and achieves the

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
	same result (*e.g.*, establishing and managing a connection between mesh nodes and external networks). Connection to an external network via VLAN tunneling, is an insubstantial difference from determining external and mesh connections via internal logic. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. Each mesh access point within the mesh forms an CAPWAP session with a controller. WLAN traffic is encapsulated inside CAPWAP and is mapped to a VLAN interface on the controller. Bridged Ethernet traffic can be passed from each Ethernet interface on the mesh network and does not have to be mapped to an interface on the controller (see Figure 5: Logical Bridge and WLAN Mapping, on page 12). ***Figure 5: Logical Bridge and WLAN Mapping***  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
35(c): wherein the connectivity logic in each node contains a distributed routing table;	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function wherein the connectivity logic in each node contains a distributed routing table. The documentation below outlines functionality that is indicative of a distributed routing table on each node, as it demonstrates the ability of each node to independently evaluate, select, and adjust paths within the mesh network based on dynamic conditions: The description of background scanning and MAP Fast Ancestor finding shows that each node independently discovers and evaluates connections to other nodes and maintains their topology by using AWPP (*e.g.*, connectivity logic), implying each has awareness of available paths—a foundational component of a routing table. **Background scanning and MAP fast ancestor find mode (Concept)** Cisco mesh access points (MAPs) perform the following functions: • Interconnects over wireless links in a tree topology, • Uses Adaptive Wireless Path Protocol (AWPP) to create and maintain their topology, and • Supports additional features: Background Scanning and MAP Fast Ancestor Finding. When a MAP comes up, it tries to look for another MAP (parent) to join and reach the gateway through a RAP. The same happens when a MAP loses connectivity with its existing parent. This procedure is known as mesh tree convergence. https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/17-12/config-guide/b_wl_17_12_cg/m_mesh_ewlc.html The Adaptive Wireless Path Protocol (AWPP) enables each MAP to dynamically determine the best path back to a RAP evaluation of network path metrics, such as ease values and hop counts, a decision-making process indicative of the use of a distributed routing table, *e.g.*, wherein the connectivity logic in each node contains a distributed routing table.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/dam/en/us/td/docs/wireless/controller/9800/17-1/deployment-guide/c9800-mesh-rel-17-1.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Mesh neighbors, parents, and children

Relationships among mesh APs are as a parent, child, or neighbor (see Figure 73).

- A parent AP offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
  - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an AP with a higher ease value is selected.
- A child AP selects the parent AP as its best route back to the RAP.
- A neighbor AP is within RF range of another AP but is not selected as its parent or a child because its ease value is lower than that of the parent.
- Mesh networks are half duplex, meaning after the first hop (RAP to MAP) overall throughput is decreased by 50% for each additional hop (MAP to MAP). Where Ethernet bridging clients are used in MAPs and heavy traffic is passed, high throughput consumption nay result, which may cause the downlink MAPs to disassociate from the network due to throughput starvation.

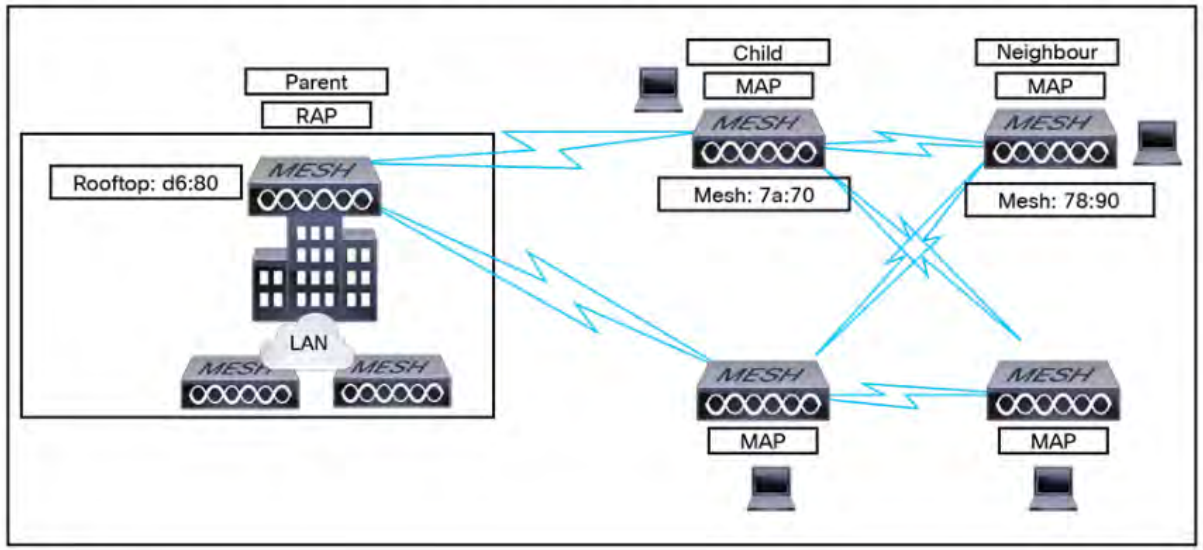


Figure 73.

Parent, child, and neighbor relationships

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

Criteria for choosing the best parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

1. A list of channels with a neighbor is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
2. The channels with a neighbor are sought by actively scanning in the seek state, and the backhaul channel is changed to the channel with the best neighbor.
3. The parent is set to the best neighbor, and the parent-child handshake is completed in the seek state.
4. Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and it is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Parent decision

A parent mesh AP is chosen by using the adjusted ease, which is the ease of each neighbor divided by the number of hops to the RAP: Adjusted ease = Min (ease at each hop) Hop count.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

The node-initiated NEIGHBOR_REQUEST and evaluation of neighbor responses show that routing decisions are made locally at each node, based on real-time link assessments using AWPP. This further supports the existence of node-specific routing data contained in the connectivity logic.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh AP provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Each mesh access point maintains routing data—specifically, it checks for routing loops by comparing route information (including MAC addresses of each hop to the root access point) against its own MAC address. This

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

further suggests that each node stores and processes routing information locally, consistent with maintaining a routing table to support AWPP (Adaptive Wireless Path Protocol) operations.

> **Loop prevention**
>
> To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh AP can easily detect and discard routes that loop.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

The connectivity logic in each node containing a distributed routing table is further shown by the show command line interface manual and the "show ap name mesh path" command, which retrieves the mesh AP's path.

show ap name mesh path

To see information about the mesh AP's path, use the **show ap name** *ap-name* **mesh path** command.

show ap name *ap-name* **mesh path** [**chassis** {*chassis-number* | **active** | **standby**} **R0**]

Syntax Description

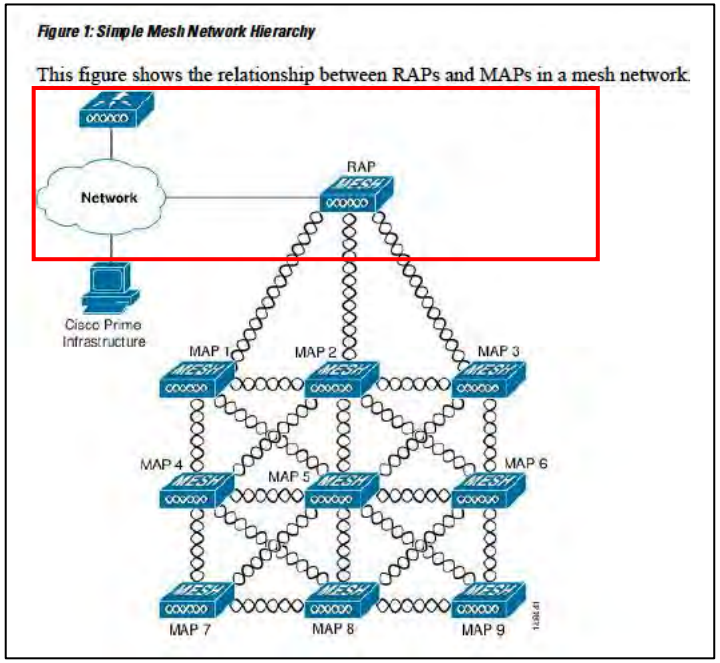
chassis-number	Enter the chassis number as either 1 or 2.
active R0	Active instance of the AP filters in Route-processor slot 0.
standby R0	Standby instance of the AP filters in Route-processor slot 0.

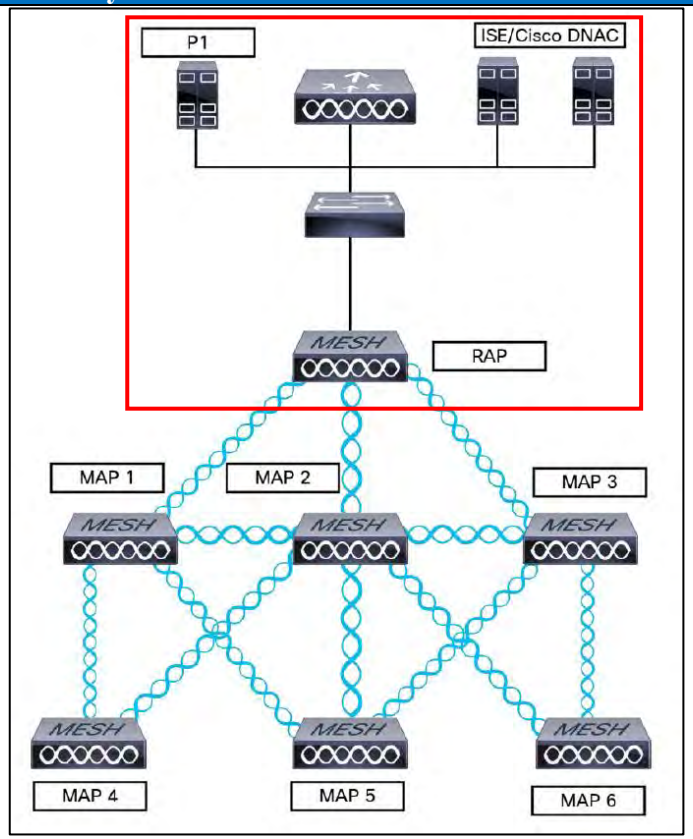
Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/cmd-ref/b_wl_16_10_cr/show-commands.html#wp5155031900

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's MAPs and RAPs achieve the same function (*e.g.*, enabling dynamic path selection and routing decisions across the mesh), by substantially the same means (*e.g.*, storing and updating neighbor information, signal quality, and hop metrics locally at each node, or referencing routing decisions derived from a centralized routing table maintained on the controller and communicated to the APs), and yield the same result (*e.g.*, establishing and maintaining optimized

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
	routes through the mesh network). The absence of a formal data structure labeled as a "routing table" is an insubstantial difference. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(d): wherein the structured mesh network functions in two configurations selected depending on whether a connection to an external network is present; in the first connected configuration the structured mesh network includes at least one structured mesh node's uplink connection comprising a connection to an external network; and	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function in two configurations (*e.g.*, a "first connected configuration" wherein the connection to the external network via the controller is not lost, and an "isolated configuration" wherein the connection to the external network via the controller is lost) selected depending on whether a connection to an external network is present. In normal operation, the Cisco Wireless Mesh Network is connected to an external network (to make/facilitate additional connections to outside networks and destinations, for example to connect clients to an external network). When so connected, the system is configured to pass information from inside the network via the uplink connection of at least one structured mesh node (*i.e.*, RAP), as shown below and in Cisco's documentation:

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points





Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf; https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

In some cases, the typical operating mode is explicitly referred to as "connected mode."

Connected mode

A Wave 2 Flex+Mesh AP (RAP or child MAP) is in connected mode when it can access and join the Catalyst 9800 and can exchange periodic keep-alive messages with the Catalyst 9800. In this mode, a Flex+Mesh AP will be able to support locally and centrally switched WLANs. It will allow regular client and child mesh APs to join.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

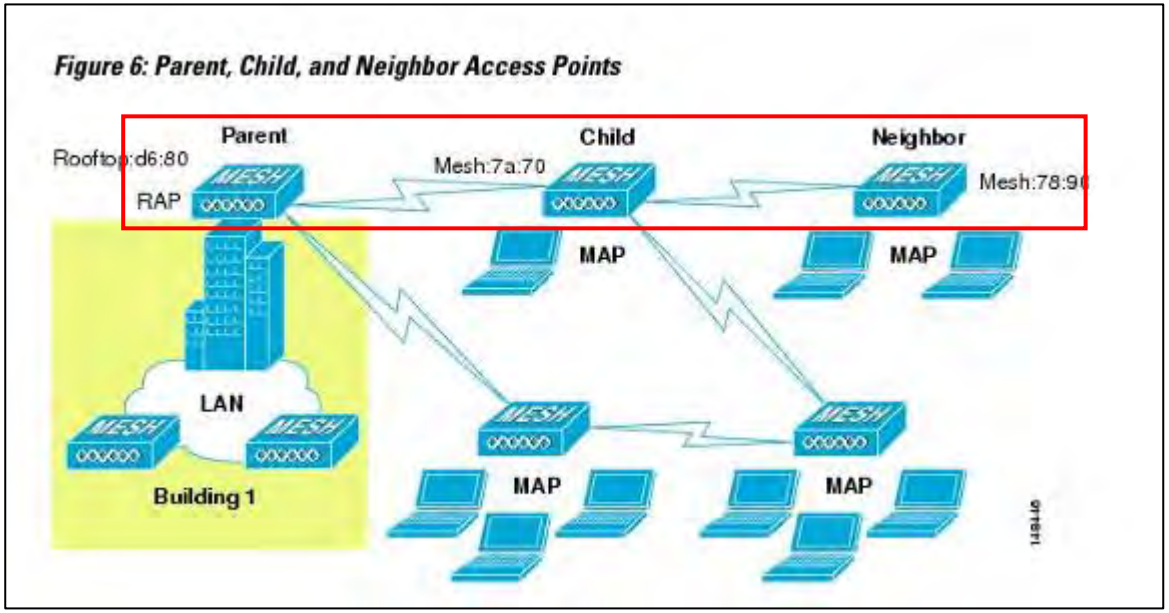
Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

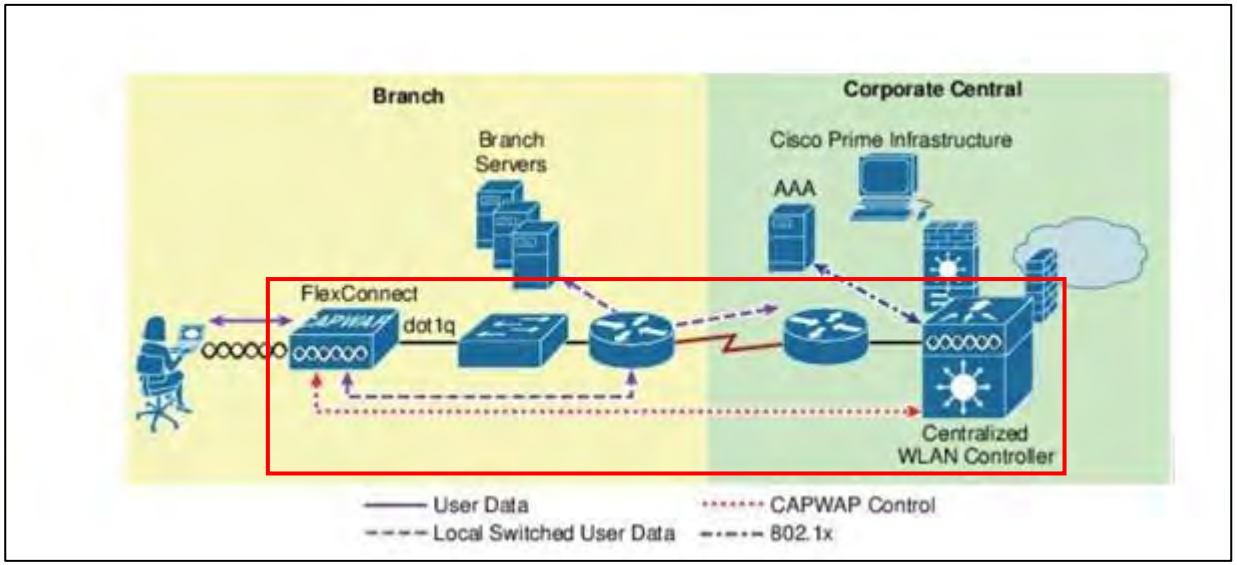
The mesh nodes have a wireless connection to the controller (and the external network) via the RAP:

> While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's system achieves the same function (external network connectivity), in substantially the same way (via the mesh node's uplink connection establishing a path to a controller), to yield the same result (connectivity beyond the local mesh network). Whether or not the controller itself is considered part of the "external network," the resulting connectivity functionally extends beyond the mesh. Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (upstream communication) are insubstantially different, and the manner of establishing that connectivity (through AWPP-determined routing and/or controller tunneling) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial.



Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
	Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.  Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(e): in the second isolated configuration none of the structured mesh nodes' uplink connections comprises a connection to an external network, and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root	Cisco's Catalyst 9100 and Aironet Access Points in the Cisco Wireless Mesh Network function in the second isolated configuration (*e.g.*, fault tolerant resilient mode) where none of the structured mesh nodes' uplink connection comprises a connection to an external network, (*e.g.*, when Cisco's Catalyst 9100 and Aironet Access Points fails to connect to an external network via the controller) and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node (*e.g.*, a Cisco's Catalyst 9100 and Aironet Access Points operating as a RAP prior to a loss of connection to the controller) as isolated root children nodes (*e.g.*, a Cisco's Catalyst 9100 and Aironet Access Points operating as MAPs) forming a tree configuration. When the Cisco Wireless Mesh Network is not connected to an external network, the network is capable of functioning in a variety of localized modes. Cisco's Catalyst 9100 and Aironet Access Points are access points that can operate in several different modes, including, for example, Flex + Mesh mode, fault tolerate resilient mode, and

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
node as isolated root children and descendant nodes forming a tree configuration;	standalone mode. Each of these modes maintain the connectivity of mesh access points when the connection to the controller (and thus to the external network) is lost. These modes can switch client data traffic locally when a connection to the controller is lost. As shown below and in Cisco's documentation, the tree structure is maintained, as a child access point (MAP) remains connected to a parent access point (another MAP or a RAP) in the second isolated configuration. **Cisco outdoor mesh access points** In addition to the mesh mode, the mesh APs can operate in the following modes: • **Local mode:** In this mode, the AP can handle clients on its assigned channel or while monitoring all channels on the band over a 180-second period. During this time, the AP listens on each channel for 50 ms for rogue client beacons, noise floor measurements, interference, and IDS events. The AP also scans for CleanAir interference on the channel. • **FlexConnect mode:** FlexConnect mode enables you to configure and control APs in a branch or remote office from the corporate office through a WAN link without having to deploy a controller in each office. This mode can switch client data traffic locally and perform client authentication locally when the connection to the controller is lost. When connected to the controller, an AP in FlexConnect mode can also tunnel traffic back to the controller. • **Flex+Mesh mode:** In this mode, both the FlexConnect and Bridge mode configuration options are available on the access point. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh. **Connected mode** A Wave 2 Flex+Mesh AP (RAP or child MAP) is in connected mode when it can access and join the Catalyst 9800 and can exchange periodic keep-alive messages with the Catalyst 9800. In this mode, a Flex+Mesh AP will be able to support locally and centrally switched WLANs. It will allow regular client and child mesh APs to join. **Standalone mode** A Wave 2 Flex+Mesh AP is in standalone mode if it loses connection to the controller but can access the local gateway. In this mode, the Wave 2 Flex+Mesh AP will disable all the centrally switched WLANs and keep the locally switched WLANs up and running. It will also allow new clients to join on locally switched WLANs using local authentication if the authentication server is reachable in the local network. Child mesh APs will NOT be allowed to join in this mode. Source: https://www.cisco.com/c/en/us/products/collateral/wireless/unified-wireless-network-sg.html#12Mesh.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points

uplink to the core network in a 'Root' AP or RAP mode. The AP is still managed by a centralized controller over CAPWAP. The AP however is capable of moving to a standalone mode where the AP will be able to continue to serve the 802.11 clients depending on the data switching method of the WLAN configured on the AP. The data can be centrally or locally switched. When data is centrally switched, all data is sent to the WLC that does the switching further. In a local switch network, the data is sent to RAP where RAP switches locally on the wired uplink. There is no difference to the central and local switched WLAN configurations and

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-8/b_mesh_88.pdf.

A Control and Provisioning of Wireless Access Points protocol (CAPWAP) Access Point (AP) can be configured to operate in two different modes:

- FlexConnect mode
- Bridge/Mesh mode

The following are the bridging features for Flex+Bridge mode:

- The Flex+Bridge mode supports the centrally switched 802.11 WLAN. Traffic for this tunneled WLAN is forwarded to and from a CAPWAP controller over an IP tunnel.
- The Flex+Bridge mode supports the Root Ethernet VLAN Bridging. A root AP bridges the traffic for bridged 802.11 WLANs and secondary Ethernet LANs to a local Ethernet LAN over its root Ethernet port.
- The Flex+Bridge mode bridging is supported on Secondary Ethernet Access Ports and Secondary Ethernet VLAN Trunk Ports.
- Fault Tolerant Resilient Mode enables an AP to continue bridging traffic when the connection to the CAPWAP controller is lost. Both mesh and non-mesh root APs continue to bridge traffic. A child mesh AP (MAP) maintains its link to a parent AP and continues to bridge traffic till the parent link is lost. A child mesh AP cannot establish a new parent or child link till it reconnects to the CAPWAP controller. Existing wireless clients on the locally switching WLAN can stay connected with their AP in this mode. Their traffic will continue to flow through the Mesh and wired network. No new or disconnected wireless client can associate to the Mesh AP in this mode.

Source: https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-5/config-guide/b_cg85/flexconnect.html.

Claim 35	Cisco Wireless Mesh Network with Catalyst and Aironet Access Points
	If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's system achieves the same function (maintaining mesh connectivity when the connection to the external network is lost), in substantially the same way (via a mesh node acting as parent or root node), to yield the same result (continuing to bridge traffic with wireless clients staying connected). Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (upstream communication and/or tunneling) are insubstantially different, and the manner of establishing that connectivity (through AWPP-determined routing and/or controller tunneling) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(f): wherein clients of the at least two structured mesh nodes in a second isolated configuration retain full connectivity within the structured network during the isolated configuration and wherein said client connectivity is uninterrupted during any transition from the isolated configuration to the first connected configuration; wherein distinct downlink frequencies are used by nodes for communication with one or more nodes lower on the tree configuration.	*See* Claim 1(f), the contents of which are incorporated herein by reference.

Exhibit 10 – U.S. Patent No. 8,520,691 (Cisco Wireless Mesh Networks with Meraki and Catalyst CW Access Points)

Cisco Systems, Inc. ("Cisco") has made, used, sold, offered for sale, and/or imported the Cisco Meraki Wireless Mesh Network including but not limited to Cisco's Meraki and Catalyst CW Access Points, with related cloud architecture and software (the "Cisco Accused Products," *see* Exhibit 11), and infringed at least Claims 1, 9-11, 15-16, 18-19, and 35 (the "Asserted Claims") of U.S. Patent No. 8,502,691 (the "'691 Patent").[1]

Cisco has infringed (literally and/or under the doctrine of equivalents), directly and/or through subsidiaries, agents, representatives, or intermediaries, one or more claims of the '691 Patent by making, using, testing, supplying, causing to be supplied, importing, selling, and/or offering for sale in the United States the Cisco Accused Products.

Cisco directly infringed under 35 U.S.C. § 271(a) by making, using, selling, offering for sale and importing the Cisco Accused Products. Cisco further directly infringed by testing, setting up, and using the Cisco Meraki Wireless Mesh Networks on Cisco's campuses and locations throughout the United States,[2] and/or assisting customers in setting up and servicing their Cisco Wireless Mesh Networks. On information and belief, Cisco has partnered with several customers to set up an infringing Cisco Meraki Wireless Mesh Network in the United States. As indicated in the Meraki Mesh Networking Guide, each Meraki access point has unique mesh characteristics independent of general technology specifications listed on the datasheet. Cisco recommends a Meraki Systems Engineer as a resource to ensure that an access point's capabilities match the desired design targets.[3] The Meraki Trial Program allows customers to run a real world proof of concept to test the design, which is arranged by a Cisco Meraki account representative. Meraki is also cloud-based architecture, and the infringing mesh networking methods and systems were controlled and put into service by Cisco.[4] Cisco knew how its customers are configuring the specific access points sold, as customers must register to the Meraki cloud and use the invoice or serial number of each purchased access point.[5]

---

[1] The identified Accused Products also include Cisco's products, services, and systems that are considered sold or related to the infringement in the United States and subject to a reasonable royalty according to the Federal Circuit's decisions in *Carnegie Mellon Univ. v. Marvell Tech. Group. Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015), *Cal. Inst. of Tech. v. Broadcom Ltd*., 25 F.4th 976 (Fed. Cir. 2022), and *Brumfield v. IBG LLC*, 97 F.4th 854 (Fed. Cir. 2024).

[2] *See, e.g*., https://www.cisco.com/c/dam/en_us/about/ciscoitatwork/downloads/ciscoitatwork/pdf/Richardson_Topology.pdf.

[3] *See, e.g*., https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

[4] *See Centillion Data Systems, LLC v. Qwest Communications International, Inc.*, 631 F.3d 1279 (Fed. Cir. 2011); *NTP, Inc. v. Research in Motion, Ltd*., 418 F. 3d 1282, 1317-1318 (Fed. Cir. 2005); *see also SiRF Tech. Inc. v. ITC*, No. 2009-1262 (Fed. Cir. Apr. 12, 2010) (direct infringement where the claims do not specifically require action by its customers or end users).

[5] *See, e.g*., https://documentation.meraki.com/MR/MR_Installation_Guides/CW9163E-MR_Installation_Guide.

Cisco indirectly infringed under 35 U.S.C. § 271(b) by providing specific guidance and instructions to its customers and end users on how to infringe the ’691 Patent and how to make and use the infringing Cisco Meraki Wireless Mesh Network. Specifically, Cisco’s Meraki Mesh Deployment Guides, Meraki Mesh Networking Guide, Technical Assistance Center, Cisco Tech Talks, technical support, Cisco Community, presentations, and technical documentation on Cisco’s website identified in the chart *infra* provide design and deployment guidelines for the deployment of, *e.g.*, secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Meraki Wireless Mesh Network solutions. These guides and instructions specifically teach and describe each and every claim element of Claims 1, 9-11, 15-16, 18-19, and 35 as demonstrated *infra*, including the necessary components to make and use the infringing Cisco Meraki Wireless Mesh Network, and the infringing configuration and operation of the cloud, access points, and software that comprise the infringing Cisco Meraki Wireless Mesh Network. These guides and instructions were published and distributed by Cisco for its customers for the sole purpose to instruct and encourage its customers to purchase the Cisco components to make and use the infringing Cisco Meraki Wireless Mesh Network (and directly infringe literally and/or under the doctrine of equivalents). Cisco has had knowledge of the ’691 Patent and/or its predecessors since 2009 when Cisco executives met with MeshDynamics to discuss partnering opportunities and discussed the application that would result in the ’691 Patent. MeshDynamics further put Cisco on actual notice of its infringement by its making, using, selling, offering for sale and importing the Cisco Accused Products that comprise the Cisco Meraki Wireless Mesh Networks in 2022. Cisco has knowingly and intentionally induced its customers to infringe the ’691 Patent, or in the alternative was willfully blind to the ’691 Patent and the infringement, following the MeshDynamics meeting in 2009.

Cisco further indirectly infringed under 35 U.S.C. § 271(c) by supplying the components used in the infringing Cisco Meraki Wireless Mesh Network, with knowledge of its infringing use (as discussed *supra*) and without substantial non-infringing uses. Further, the Cisco Accused Products were designed specifically to be used in a manner that infringed the ’691 Patent. Cisco Accused Products come enabled as mesh access points and automatically connect to establish an infringing wireless mesh network. Certain Cisco access points and proprietary software are particularly suited for Cisco’s Meraki Wireless Mesh Network, and Cisco provides specific instructions and guides on how to make and use the Cisco Meraki Wireless Mesh Network using the Cisco Accused Products. As a result of Cisco’s selling and/or offering for sale of the Cisco Accused Products to other entities, the other entities use these products for their intended purpose and according to their instructions with the result that such entities directly infringe the ’691 Patent, literally or under the doctrine of equivalents, for the reasons stated *infra*. Cisco’s Accused Products are specially made and adapted for this use in an infringing manner, are not staple articles of commerce, and do not have substantial non-infringing uses.

MeshDynamics asserts that the Cisco Accused Products literally infringed the Asserted Claims of the ’691 Patent. With respect to any claim limitation that may be found not to be literally infringed, MeshDynamics contends, in the alternative, that Cisco Accused Products infringed such claim limitations under the doctrine of equivalents and that any claim element not found to be literally met is equivalently met because any difference between the claim limitation and Cisco Accused Products is not a substantial difference. Accordingly, MeshDynamics contends that if any Cisco Accused Products is found not to have infringed any Asserted Claims literally it is nevertheless embodied by Cisco Accused Products under an operative doctrine of equivalents test (*e.g.*, function-way-result or insubstantial difference).

The below chart presents MeshDynamics’s preliminary infringement analysis of Cisco Accused Products based solely on public information. MeshDynamics has not received any discovery from Cisco, including any confidential information or source code demonstrating the operation and structure of the Cisco Accused Products. All information contained in this disclosure is the result of publicly available information regarding the Cisco Accused Products. Based on these circumstances, MeshDynamics presents these infringement allegations reflecting its analysis of the facts presently known to it. To the best of MeshDynamics’s knowledge, information, and belief, formed after an inquiry that is reasonable under the circumstances, the information contained in this disclosure is correct.

Certain information about Cisco Accused Products is not available without discovery, such as proprietary source code, documents, and deposition testimony, from Cisco. Accordingly, MeshDynamics reserves the right to supplement and/or amend this infringement contention claim chart as new information becomes available and as MeshDynamics has the opportunity to review Cisco discovery, including information related to the implementation of Cisco Accused Products, relationships with other parties who participated or contributed to the operation of Cisco Accused Products, and/or elements of Cisco Accused Products that cannot be fully examined based solely on publicly available information. Infringement investigations are ongoing, and MeshDynamics anticipates that additional facts and relevant documents will be uncovered and disclosed that create cause for supplementation and/or amendment.

Claim 1	Cisco's Meraki Wireless Mesh Network with Meraki and Catalyst CW Access Points
1pre: A structured mesh network capable of isolated operation, comprising:	Cisco's Meraki Wireless Mesh Network is a structured mesh network capable of isolated operation, which comprises a plurality of Wi-Fi nodes (*e.g.*, Cisco Meraki and Catalyst CW Access Points). The Meraki cloud-managed wireless series is a line of enterprise-grade WLAN access points ("APs") that leverage the Cisco Meraki cloud-managed architecture. Both Cisco Meraki and Catalyst CW Access Points support Wi-Fi capabilities and can be paired with the Meraki cloud platform. Meraki AP model numbers start with MR, *e.g.* MR33, MR36, MR46E, and Catalyst CW APs with CW, *e.g.* CW9162, CW9164, CW9166, are dual-mode APs which can be used with the Catalyst 9800 controller or the Meraki cloud architecture. Source: Meraki Family Datasheet; Meraki Wi-Fi Models **Cisco Meraki Cloud Management** Pairing the Cisco Catalyst 9166 Series Access Points with the Meraki cloud platform gives organizations a unified IT experience for network monitoring and management. The Meraki dashboard provides an intuitive and interactive web interface connecting your network to the industry's leading cloud IT platform. Through the dashboard, Meraki provides sophisticated and scalable tools to automate network optimization, deploy policy and segmentation configurations across thousands of sites and devices, and manage a full-stack network from SD-WAN to Access to IoT technologies. The platform supports over 3.5 million active networks around the world. Working together, the Catalyst 9166 Series and Cisco Meraki offer such features as: • Cisco Spaces • Cisco Identity Services Engine • Meraki Health intelligent optimization and assurance • Meraki Vision, smart cameras, and sensors for network closet monitoring Source: CW9166 Datasheet; CW9166 The Meraki and Catalyst CW Access Points contain the same mesh networking capabilities, as shown in Cisco's documentation and the Meraki Family Datasheet, and the Meraki Wireless Mesh Networking and Mesh Deployment Guide apply to both product lines for setting up a Cisco Meraki Wireless Mesh Network. Multiple access points communicate over wireless interfaces to form a single network. Meraki's mesh networking functionality is automatic, self-healing and available on all Meraki (including Catalyst CW) Access Points.



Meraki

Wireless Mesh Networking

In a wireless mesh deployment, multiple Access points (with or without Ethernet connections) communicate over wireless interfaces to form a single network. This wireless communication between Access points is called Mesh Networking. Meraki's mesh networking functionality is automatic, self-healing and available on all access points.

A detailed article on designing wireless Mesh networks can be found on Mesh Deployment Guide.

Source: Wireless Mesh Networking; *see also* Meraki Best Practices - Wireless Mesh Networking

Mesh functionality is enabled by default in every Meraki AP. All Meraki APs support wireless meshing without any additional configuration ("zero-configuration mesh") or licensing.

Access Point Capabilities

All Meraki access points support wireless meshing without any additional configuration or licensing. The most successful mesh deployments undergo both pre-deployment design and post-design testing phases.

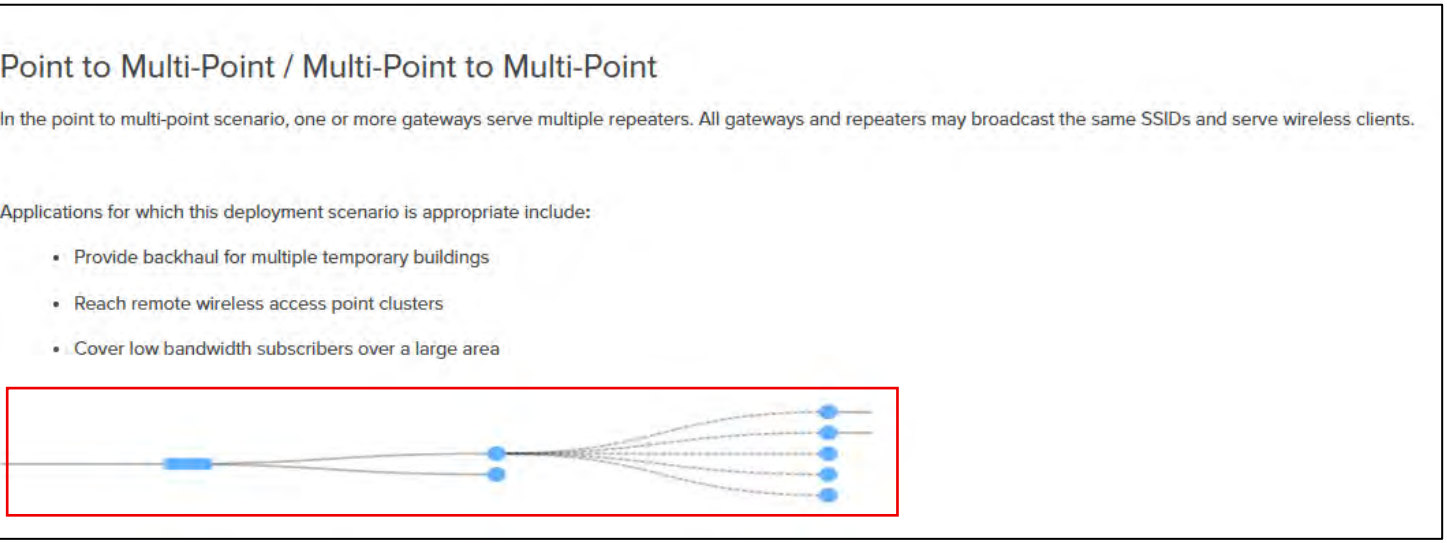
Source: Mesh Deployment Guide

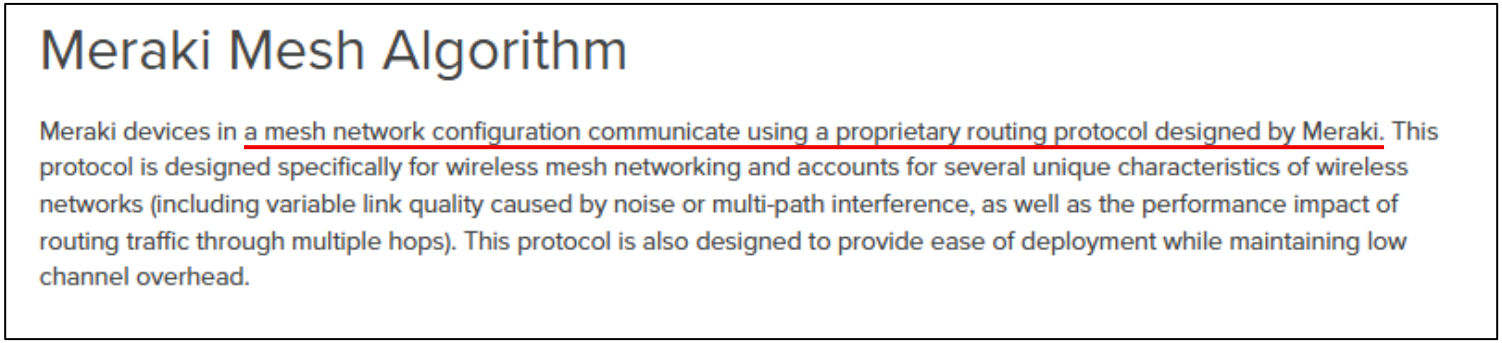
Class-leading enterprise features

Meraki MR wireless access points come equipped with industry-leading features that make them ideal for demanding enterprise deployments:

- Self-configuring plug-and-play deployment
- 802.11be MU-MIMO with up to sixteen spatial streams built for voice and video
- Integrated Bluetooth® IoT radio
- IoT ready (ESL integration)
- Self-healing, zero-configuration mesh

Source: Meraki Family Datasheet

	Mesh networking, included in every Meraki AP, extends coverage to hard-to-wire areas and creates a self-healing network that is resilient to cable and switch failures, continuing to operate despite failures or configuration changes in the rest of the network without the need for manual configuration or optimization. Source: Meraki Family Datasheet Cisco Meraki Wireless Mesh Networks organize the wireless access points in a structured mesh network and tree-like configuration with a gateway connected to the external network and mesh access points (*i.e.*, repeaters) connected to the gateway, as shown in red below:  Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.
1(a): at least two structured mesh nodes;	Cisco's Meraki Wireless Mesh Network comprises at least two structured mesh nodes (*e.g.*, Meraki access points operating as gateways and repeaters), wherein each structured mesh node comprises at least a connectivity logic.

wherein each structured mesh node comprises at least a connectivity logic;	 Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. Meraki access points in a Meraki Wireless Mesh Network comprise at least a connectivity logic (*e.g.*, the Meraki Mesh Routing Algorithm, Gateway Selection Logic).  Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

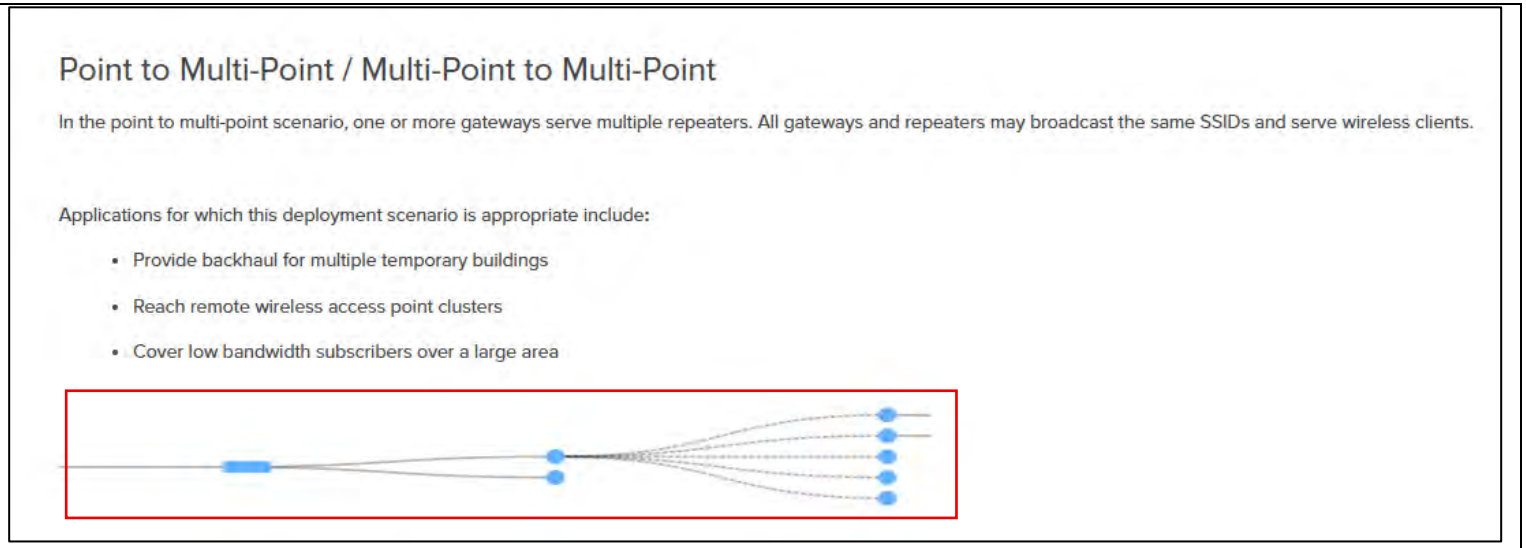
1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

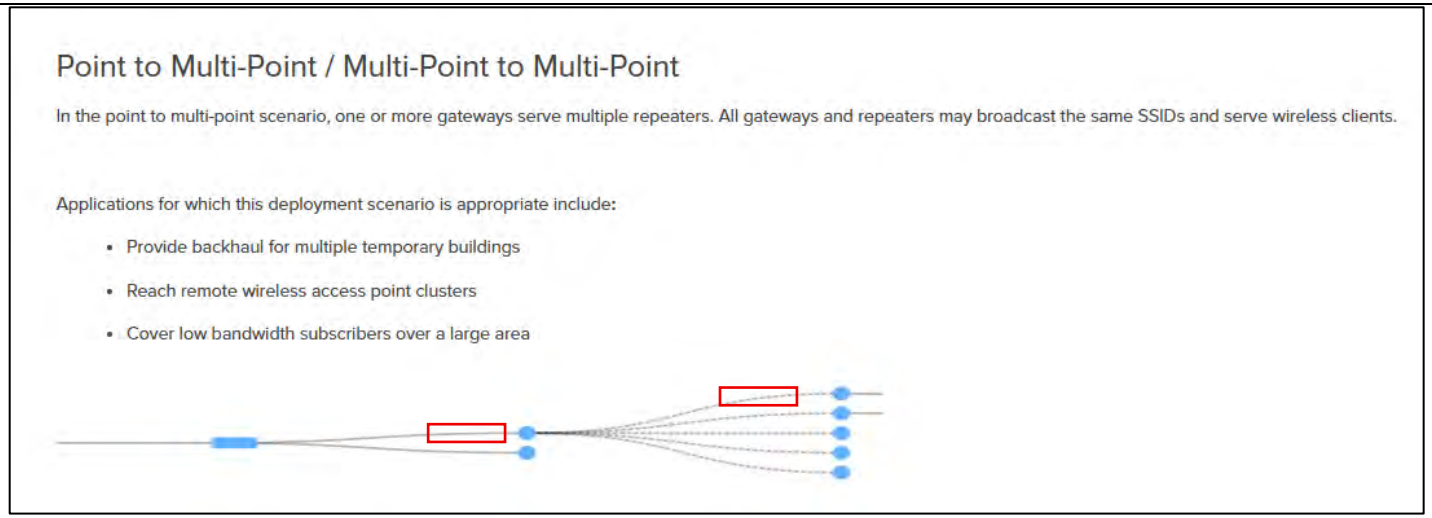
   ⓘ If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Meraki access points are capable of meshing in a tree topology configuration ("structured mesh nodes"), as evidenced by the Meraki documentation shown above below (*e.g.*, point-to-multipoint deployment). One or more gateway access points can serve at least two downstream repeaters arranged in a tree topology configuration.

	 Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Meraki's wireless mesh system performs the same function (enabling mesh nodes to determine and maintain appropriate connections within the network), in substantially the same way (through the Meraki's connectivity logic (*e.g.*, Mesh Routing Algorithm, Gateway Selection Logic) , to achieve the same result (maintaining structured mesh network connections. A person of ordinary skill in the art would view any differences in implementation as insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1(b): an uplink radio operating on an uplink frequency and a downlink radio operating on a distinct downlink frequency;	Meraki access points in a Meraki Wireless Mesh Network comprise an uplink radio operating on an uplink frequency and a downlink radio operating on a distinct downlink frequency. Meraki MR Series Access Points include multiple radios capable of operating on both 2.4 GHz and 5 GHz bands. These radios operate on distinct frequencies (*i.e.*, distinct channels), allowing one radio to act as an uplink and another as a downlink. Annotations in the image below, are representative of an uplink or downlink with respect to the associated nodes.



Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

As shown in Cisco Meraki's documentation, access points to mesh may operate on a 5 GHz uplink channel, while using the distinct 2.4 GHz channels for downlink client access.

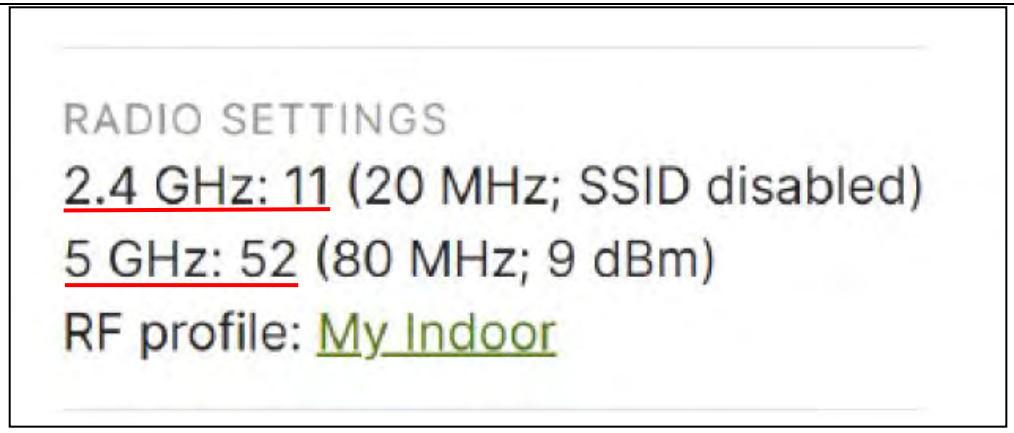
As shown *supra*, APs in Cisco's Meraki Wireless Mesh Network include radios capable of operating on distinct frequencies for uplink and downlink communication. For example, the MR86 datasheet identifies two 2.4/5 GHz radios that can be flexibly used—one for mesh backhaul (uplink) and one for client access (downlink)—with each radio operating on a distinct frequency. The 2.4 GHz and 5 GHz bands support multiple non-overlapping channels (*e.g.*, channels 1, 6, and 11 on 2.4 GHz; and 24 non-overlapping channels across the 5 GHz band), allowing the radios to operate concurrently without interference. Thus, an uplink radio can operate on one frequency (*e.g.*, 5.25 GHz) while the downlink radio operates on a distinct frequency (*e.g.*, 2.4 GHz or a separate 5 GHz channel).

High Performance 802.11ax Wireless

The Cisco Meraki MR86 is a cloud-managed 4x4:4 802.11ax access point that raises the bar for wireless performance and efficiency. Designed for next-generation deployments in offices, schools, hospitals, shops, and hotels, the MR86 offers high throughput, enterprise-grade security, and simple management.

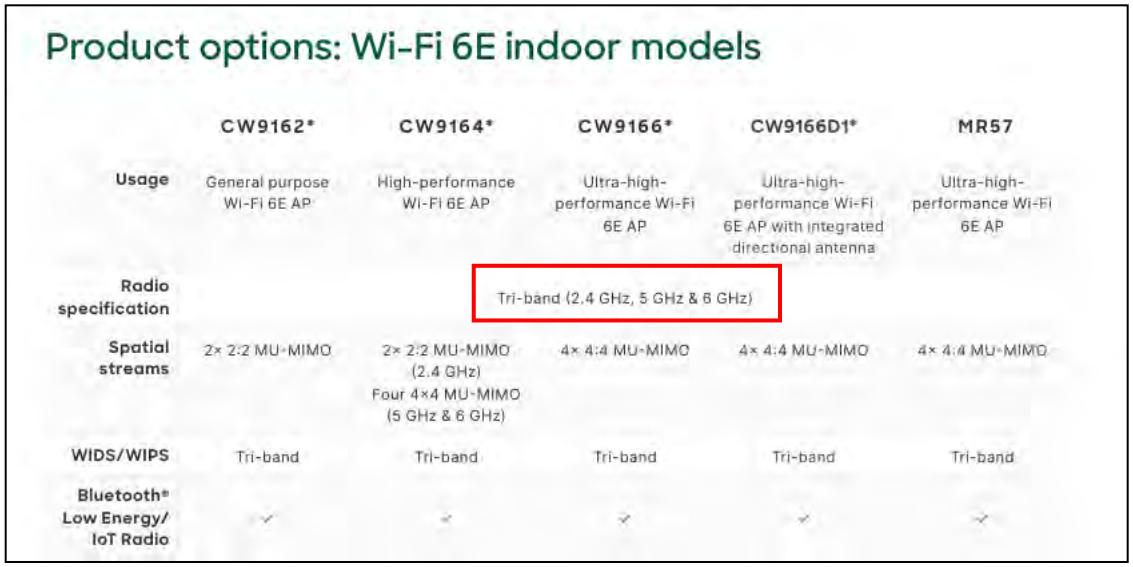
The MR86 provides a maximum of 3.0 Gbps* aggregate frame rate with concurrent 2.4 GHz and 5 GHz radios. A dedicated third radio provides real-time WIDS/WIPS with automated RF optimization, and a fourth integrated radio delivers Bluetooth scanning and beaconing.

Source: https://documentation.meraki.com/MR/MR_Overview_and_Specifications/MR86_Datasheet.

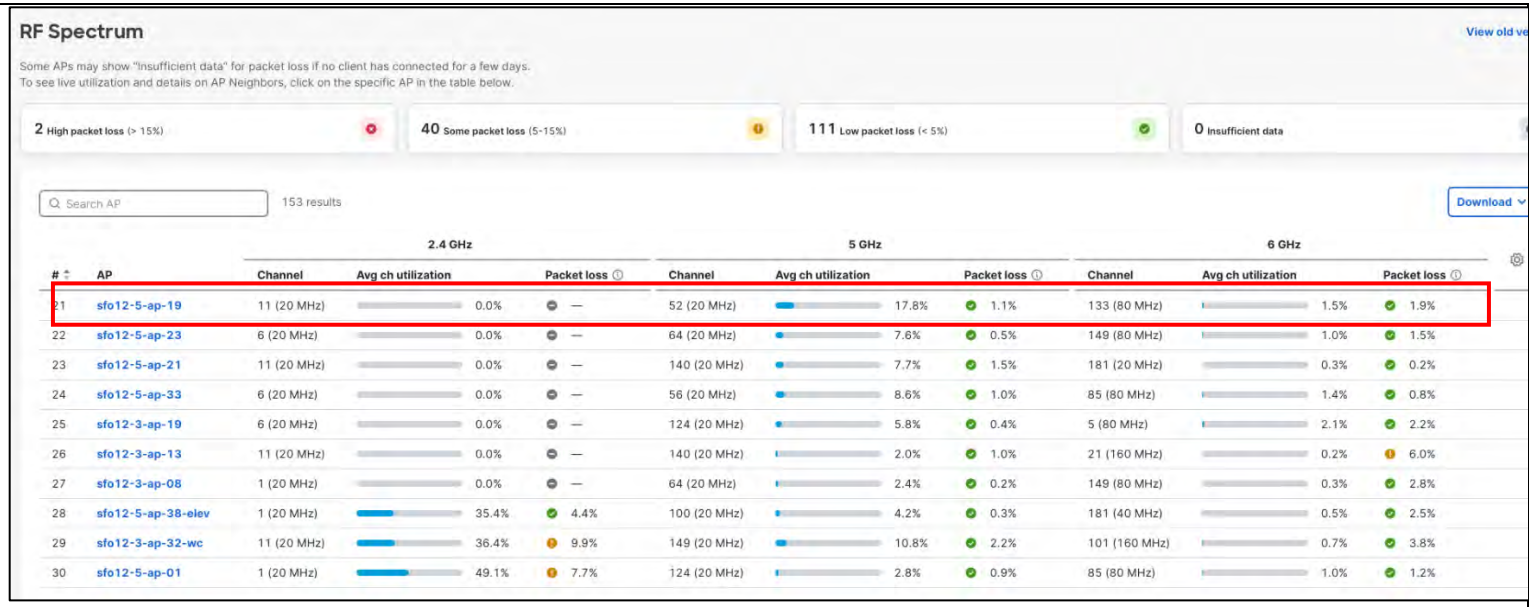


Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking#Mesh_Network_Components.

Additional radios are tri-band, and each radio is operating on a different frequency and channel.



Source: https://meraki.cisco.com/product-collateral/meraki-wireless-cloud-family-datasheet-managed-wireless-access-points/?file.



Source: https://documentation.meraki.com/MR/Monitoring_and_Reporting/AP_Neighbors.

The radios may be used concurrently for meshing and client access, with one radio operating as an uplink radio on a distinct uplink frequency (e.g., for backhaul communication with another mesh node) and the other operating as a downlink radio on a distinct downlink frequency.

Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

Extending the LAN with a Wireless Mesh Link

Cisco Meraki MR access points (APs) can operate as mesh repeaters, which allows them to extend the wireless network range off of a limited number of gateway APs. Since repeaters also support wired clients plugged into their wired interface, a repeater can be used to bridge a remote LAN segment back to the main network.

Wireless Mesh repeaters can back-haul to access points acting as wireless gateways and:

- Utilize the same radio as used for mesh to serve wireless clients
- Utilize non-meshing radios to serve wireless clients
- Utilize wired uplink for use cases described below

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Extending_the_LAN_with_a_Wireless_Mesh_Link.

APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference as possible using channel selection algorithms.

Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in.

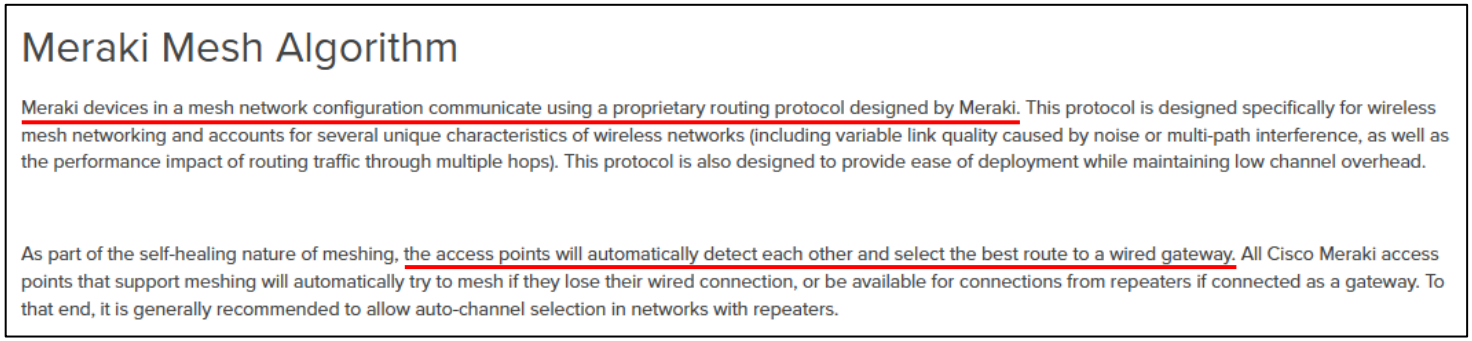
Source: https://documentation.meraki.com/MR/Radio_Settings/Manually_Changing_Channels_in_a_Mesh_Network.

Further, Meraki access points support external antennas and directional configurations for mesh backhaul, enabling the radios to operate on distinct frequencies based on antenna port assignment and channel configuration. In some deployments, Distinct channels and frequencies are on a per hop basis (*e.g.*, between mesh nodes creating a backhaul mesh link).

Can I Mix Antenna Types on Outdoor APs?

It depends. On 2x2:2 outdoor access points, such as MR74/76, there are use cases for mixing antenna types, such as when creating a mesh link. For example, there might be a distant outdoor area that needs client coverage. An MR76 mesh AP can be deployed with ANT-20 omni antennas on the 2.4GHz ports to serve clients in the area around the access point, while an ANT-27 sector antenna can be connected to the 5GHz ports of the MR76 and aimed back at a gateway AP on the building to create a 5GHz backhaul mesh link.

On 4x4:4 outdoor access points, such as MR84 and MR86, all four antenna ports are dual-band and diplexed to both radios. On these APs, all antennas should be identical and have the same orientation. Do not mix omni and patch antennas or patch and sector antennas, for example, on 4x4:4 access points.

	Source: https://documentation.meraki.com/MR/Other_Topics/Frequently_Asked_Questions_regarding_Cisco_Meraki_Antennas. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Meraki's mesh networking solution achieves the same function (separating uplink and downlink communication paths), by substantially the same means (employing multiple radios that operate on non-overlapping channels—often within the 5 GHz band—where one radio is used for mesh backhaul and another for client-facing traffic), and yields the same result (enabling concurrent bidirectional communication and minimizing interference). The Meraki system can be configured to assigns radios and channels in such a manner, which is functionally equivalent to the claim requirement. A person skilled in the art would consider any difference between the claimed invention and the Meraki Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1(c): wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node using its uplink radio and client devices or other mesh nodes connect to each node using each node's downlink radio;	Cisco Meraki Access Points, including those operating in mesh configurations, comprise a connectivity logic, wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node using its uplink radio, and whether client devices or other mesh nodes connect to each node using the node's downlink radio. Meraki's connectivity logic (*e.g.*, the Meraki Mesh Routing Algorithm, Gateway Selection Logic) determines whether a node connects to an external network (*e.g.*, through gateways) or another node (other repeaters).  Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

The connectivity logic (*e.g.*, Meraki Mesh Algorithm, Gateway Selection Logic) determines whether to connect to an external network through a gateway or to another repeater based on gateway availability and link quality. Additionally, the connectivity logic scans to determine if potentially viable gateways connect to an external network.

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

   ⓘ If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

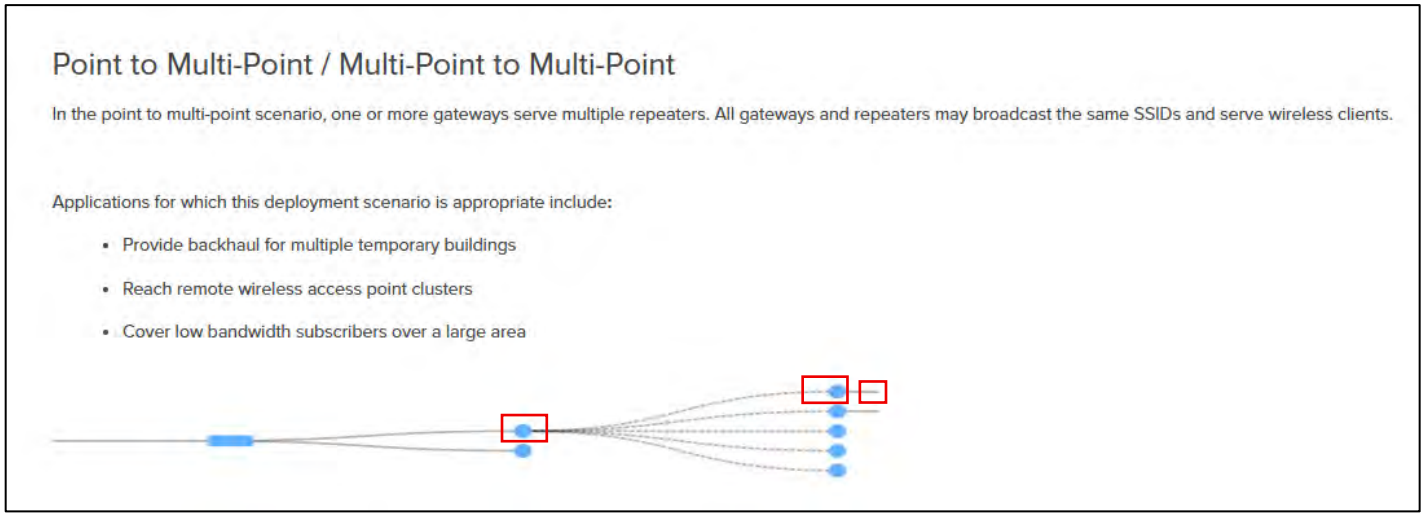
Meraki repeaters determine when to connect, via its downlink radio to another mesh node (*e.g.*, repeater) or to a client device (*e.g.,*a laptop of other user device).

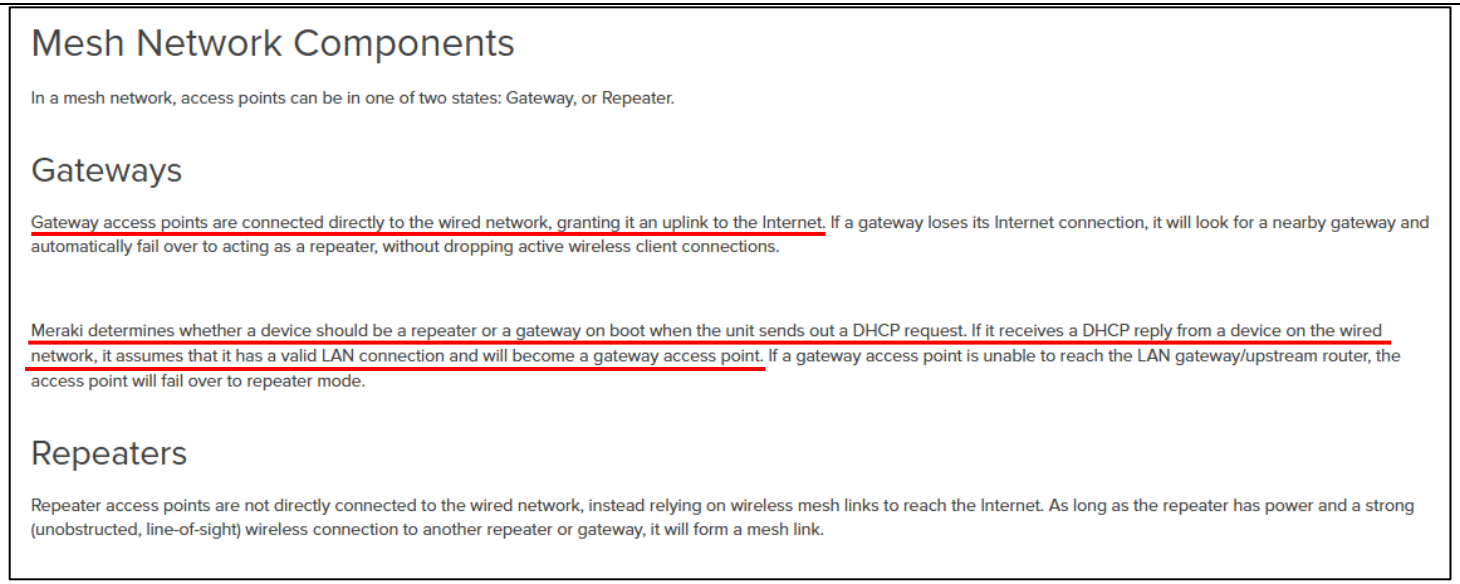
Maximum Mesh Hops

There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Annotations in the image below, are representative of a nodes connected via a backhaul with a downlink client connection.

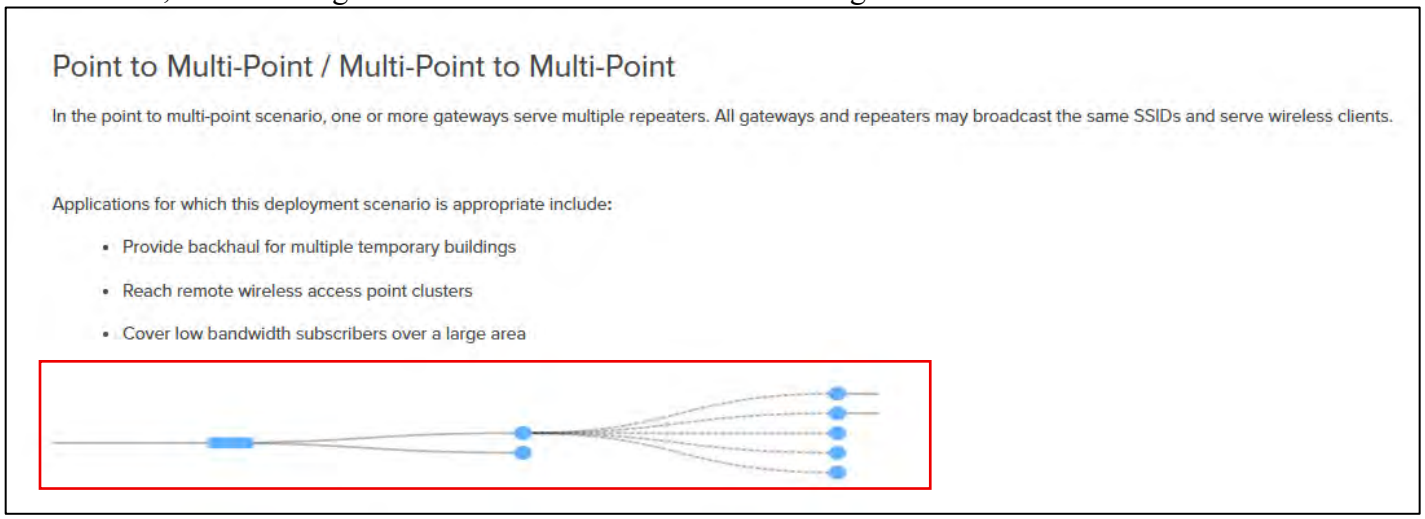
	 Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide. If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Cisco's implementation using the Meraki Mesh Algorithm and/or Gateway Selection Logic performs the same function (routing traffic between mesh nodes and external networks or among mesh nodes themselves), in substantially the same way (by evaluating link quality, hop count, and gateway availability to determine optimal upstream and downstream paths), to yield the same result (establishing a connection to the external network or client). A person skilled in the relevant art would consider this difference in how the uplink is selected to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
1(d): wherein the structured mesh network functions in two configurations selected depending on whether a connection to an external network is present; in the first connected configuration the structured mesh network includes at least one structured mesh	Meraki access points in a Meraki Wireless Mesh Network function in two configurations selected depending on whether a connection to an external network is present. In the first configuration (*e.g.*, the gateway is connected to the internet), at least one access point functions as a gateway node with an active uplink connection to the external network (*e.g.*, the internet).

node's uplink radio comprising a connection to an external network;	 Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (external network connectivity), in substantially the same way (via the mesh node's uplink radio establishing a path to a router/controller/switch), to yield the same result (connectivity beyond the local mesh network). Whether or not a router, controller, or switch itself is considered part of the "external network," the resulting connectivity functionally extends beyond the mesh. The means of enabling external network communication (upstream communication) are insubstantially different, and the manner of establishing that connectivity (through the use of Meraki's Mesh Network algorithm, Gateway Selection Logic) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to. *See* below description in 1(e) of the second configuration, incorporated herein by reference.
1(e): and in the second isolated configuration none of the structured mesh nodes' uplink radio comprises a connection to	Meraki access points in a Meraki Wireless Mesh Network function in a second isolated configuration wherein none of the structured mesh nodes' uplink radio comprises a connection to an external network, and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node as isolated root children nodes forming a tree configuration;

an external network, and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node as isolated root children nodes forming a tree configuration;	As shown below, the Gateway AP does not drop active wireless client connections despite losing its connection to the external network (*e.g.*, internet). As such the Gateway AP acts as the isolated network root node of the isolated configuration: **Mesh Network Components** In a mesh network, access points can be in one of two states: Gateway, or Repeater. **Gateways** Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. In one method outlined in Cisco documentation, an isolated configuration may be achieved to allow LAN access such that the gateway acts as an isolated network root that retains its function as a gateway. **Gateway AP Switches to Repeater Mode** **Last updated:** Jan 6, 2025 TABLE OF CONTENTS 1. Troubleshooting A gateway AP is an access point that has a wired interface configured with an IP address, that is connected to the LAN and has a route to the Internet. The AP only converts to a repeater when one of the following is true: • The AP cannot receive an ARP reply packet from the default gateway on the LAN which is usually a local firewall or router • The AP is unable to obtain a valid IP address via DHCP If the AP is configured with a static IP, it will advertise "*ssid name*-bad gateway" when the default gateway is unavailable. If it has no other APs to mesh with that provide a route to the Internet, the AP will remain offline. If it does find another AP with a route to the Internet, it will act as a repeater but the Dashboard will report that it has an invalid IP configuration. If the Internet is unavailable and the SSID is configured to allow LAN access, the AP will continue to act as a gateway because it still has a valid IP and can reach the local router or firewall.

Source: https://documentation.meraki.com/MR/Monitoring_and_Reporting/Gateway_AP_Switches_to_Repeater_Mode.

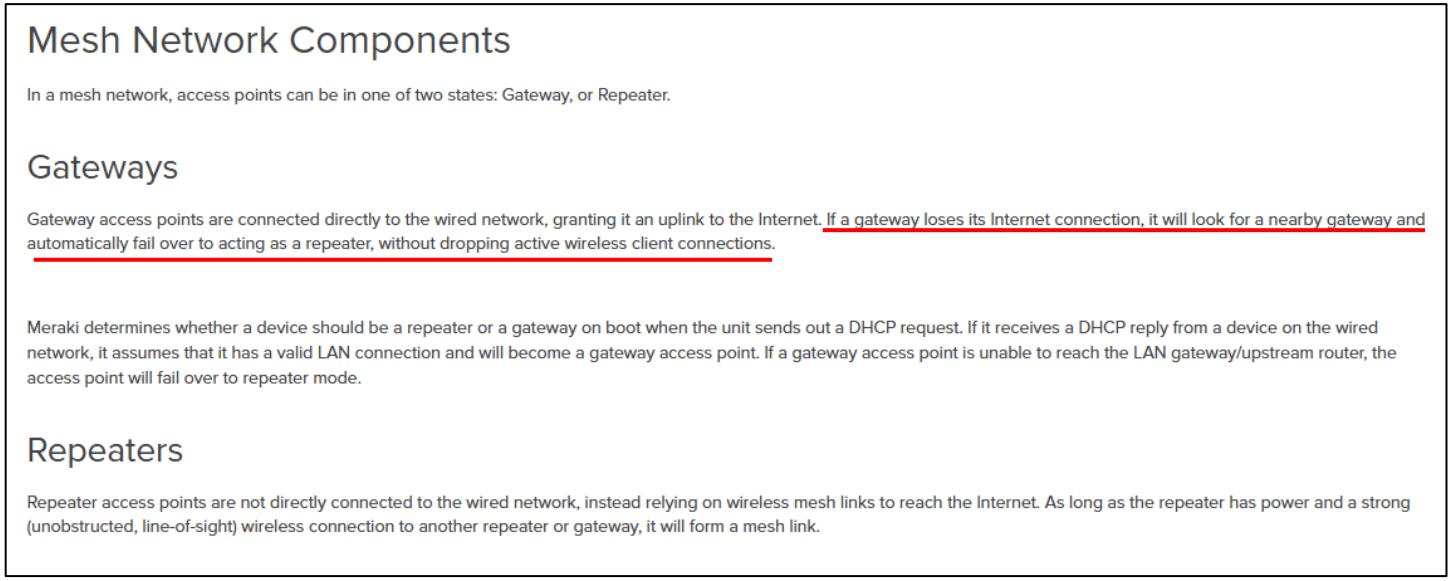
Furthermore, all remaining nodes remain connected in a tree configuration.



Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (maintaining mesh connectivity when the connection to the external network is lost), in substantially the same way (via a mesh node acting as gateway or repeater node), to yield the same result (continuing to bridge traffic with wireless clients staying connected). Thus, the functional result of preserving communication is preserved, and the means of accomplishing it (a continuing structured mesh networked) are insubstantially different and fulfill the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Accused Products to be minor and insubstantial.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

1(f): wherein clients of nodes of the structured network in a second isolated configuration retain full connectivity within the structured network during the isolated configuration and wherein said client connectivity is uninterrupted during any transition from the isolated configuration to the first connected configuration wherein distinct downlink frequencies are used by nodes for communication with one or more nodes lower on the tree configuration.	Meraki access points in a Meraki Wireless Mesh Network function in the second isolated configuration wherein clients of nodes of the structured network in a second isolated configuration retain full connectivity (*e.g.*, when the Gateway AP maintains its role as a gateway despite the loss of a connection to the internet) within the structured network during the isolated configuration (*e.g.*, SSID configured to allow LAN access) and wherein said client connectivity is uninterrupted during any transition from the isolated configuration to the first connected configuration; wherein distinct downlink frequencies are used by nodes for communication with one or more nodes lower on the tree configuration. As shown below, and in the Cisco documentation, the gateway AP does not drop active wireless client connections despite losing its connection to the external network (*e.g.*, internet).  Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. In one method outlined in Cisco documentation, an isolated configuration may allow LAN access such that the network acts as an isolated network that retains connectivity.

Gateway AP Switches to Repeater Mode

Last updated: Jan 6, 2025

A gateway AP is an access point that has a wired interface configured with an IP address, that is connected to the LAN and has a route to the Internet. The AP only converts to a repeater when one of the following is true:

TABLE OF CONTENTS
1. Troubleshooting

- The AP cannot receive an ARP reply packet from the default gateway on the LAN which is usually a local firewall or router
- The AP is unable to obtain a valid IP address via DHCP

If the AP is configured with a static IP, it will advertise "*ssid name*-bad gateway" when the default gateway is unavailable. If it has no other APs to mesh with that provide a route to the Internet, the AP will remain offline. If it does find another AP with a route to the Internet, it will act as a repeater but the Dashboard will report that it has an invalid IP configuration.

If the Internet is unavailable and the SSID is configured to allow LAN access, the AP will continue to act as a gateway because it still has a valid IP and can reach the local router or firewall.

Source: https://documentation.meraki.com/MR/Monitoring_and_Reporting/Gateway_AP_Switches_to_Repeater_Mode.

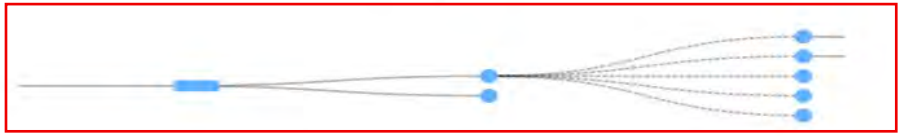
Furthermore, all remaining nodes remain connected in a tree configuration.

Point to Multi-Point / Multi-Point to Multi-Point

In the point to multi-point scenario, one or more gateways serve multiple repeaters. All gateways and repeaters may broadcast the same SSIDs and serve wireless clients.

Applications for which this deployment scenario is appropriate include:

- Provide backhaul for multiple temporary buildings
- Reach remote wireless access point clusters
- Cover low bandwidth subscribers over a large area



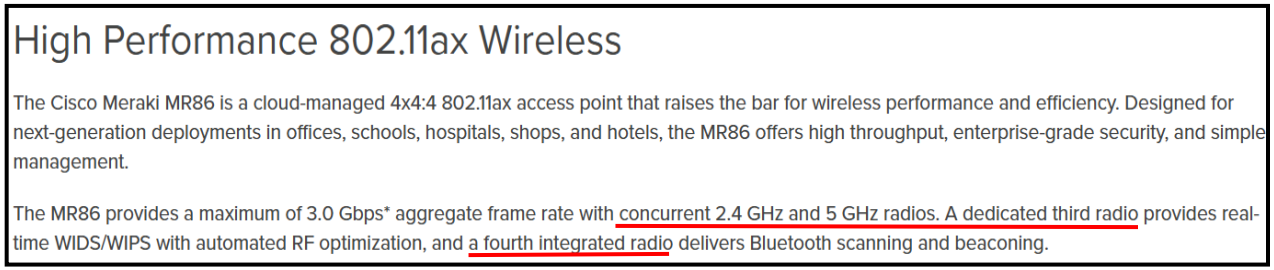
Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.

As shown in Cisco Meraki's documentation access points can mesh on a 5 GHz uplink channel, while using the distinct 2.4 GHz channels for downlink client access.

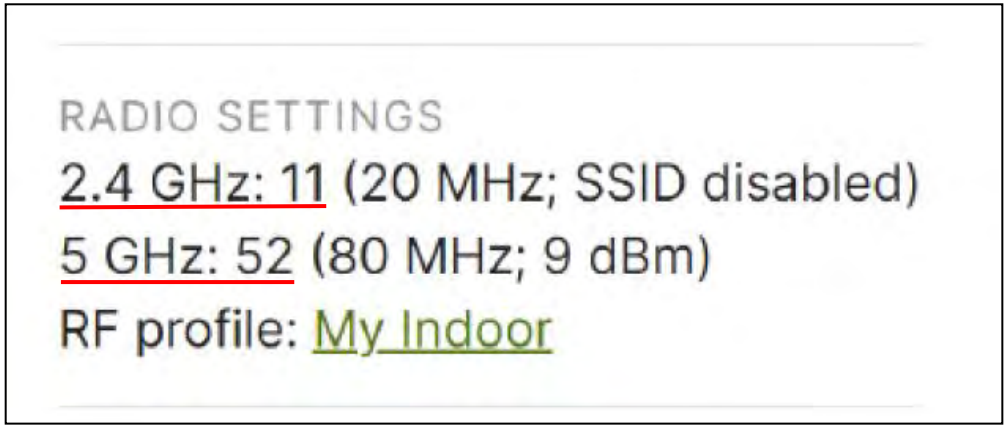
See claim 1(b) *supra*, which is incorporated herein by reference.

In Cisco's Meraki Wireless Mesh Network, distinct downlink frequencies are used by nodes to communicate with one or more downstream nodes in the tree topology.

As shown *supra*, APs in Cisco's Meraki Wireless Mesh sNetwork include radios capable of operating on distinct frequencies for uplink and downlink communication. For example, the MR86 datasheet identifies two 2.4/5 GHz radios that can be flexibly used—one for mesh backhaul (uplink) and one for client access (downlink)—with each radio operating on a distinct frequency. The 2.4 GHz and 5 GHz bands support multiple non-overlapping channels (*e.g.*, channels 1, 6, and 11 on 2.4 GHz; and 24 non-overlapping channels across the 5 GHz band), allowing the radios to operate concurrently without interference. Thus, an uplink radio can operate on one frequency (*e.g.*, 5.25 GHz) while the downlink radio operates on a distinct frequency (*e.g.*, 2.4 GHz or a separate 5 GHz channel).



Source: https://documentation.meraki.com/MR/MR_Overview_and_Specifications/MR86_Datasheet.



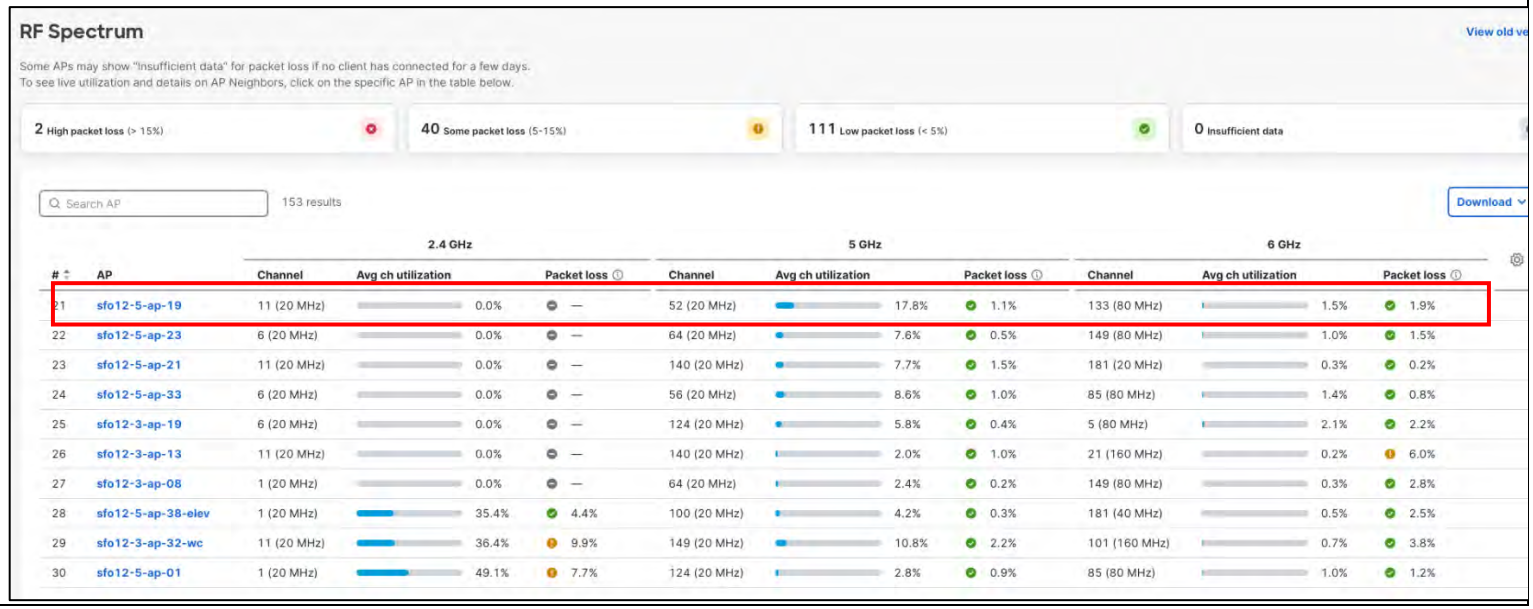
Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking#Mesh_Network_Components.

Additional radios are tri-band, and each radio is operating on a different frequency and channel.

Product options: Wi-Fi 6E indoor models

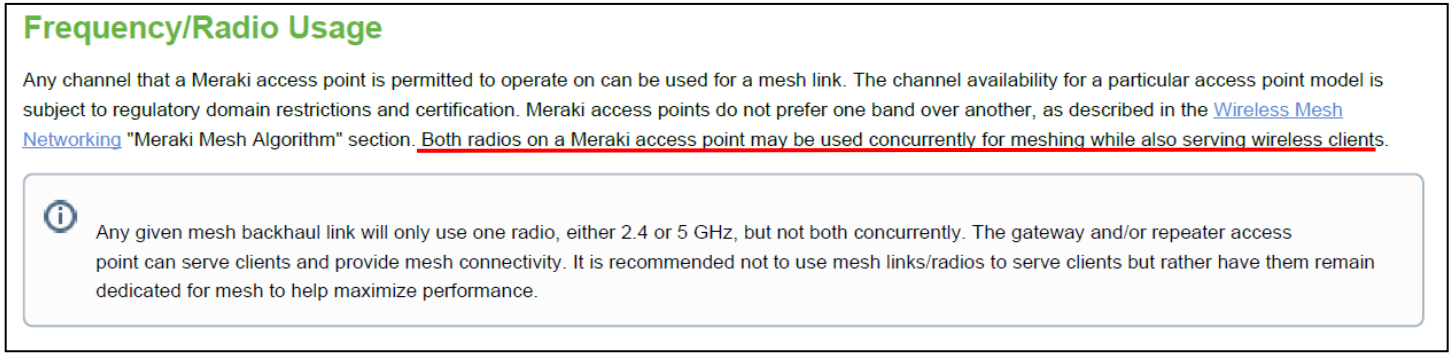
	CW9162*	CW9164*	CW9166*	CW9166D1*	MR57
Usage	General purpose Wi-Fi 6E AP	High-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP	Ultra-high-performance Wi-Fi 6E AP with integrated directional antenna	Ultra-high-performance Wi-Fi 6E AP
Radio specification	Tri-band (2.4 GHz, 5 GHz & 6 GHz)				
Spatial streams	2× 2:2 MU-MIMO	2× 2:2 MU-MIMO (2.4 GHz) Four 4×4 MU-MIMO (5 GHz & 6 GHz)	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO	4× 4:4 MU-MIMO
WIDS/WIPS	Tri-band	Tri-band	Tri-band	Tri-band	Tri-band
Bluetooth® Low Energy/ IoT Radio	✓	✓	✓	✓	✓

Source: https://meraki.cisco.com/product-collateral/meraki-wireless-cloud-family-datasheet-managed-wireless-access-points/?file.



Source: https://documentation.meraki.com/MR/Monitoring_and_Reporting/AP_Neighbors.

The radios may be used concurrently for meshing and client access, with one radio operating as an uplink radio on a distinct uplink frequency (e.g., for backhaul communication with another mesh node) and the other operating as a downlink radio on a distinct downlink frequency.

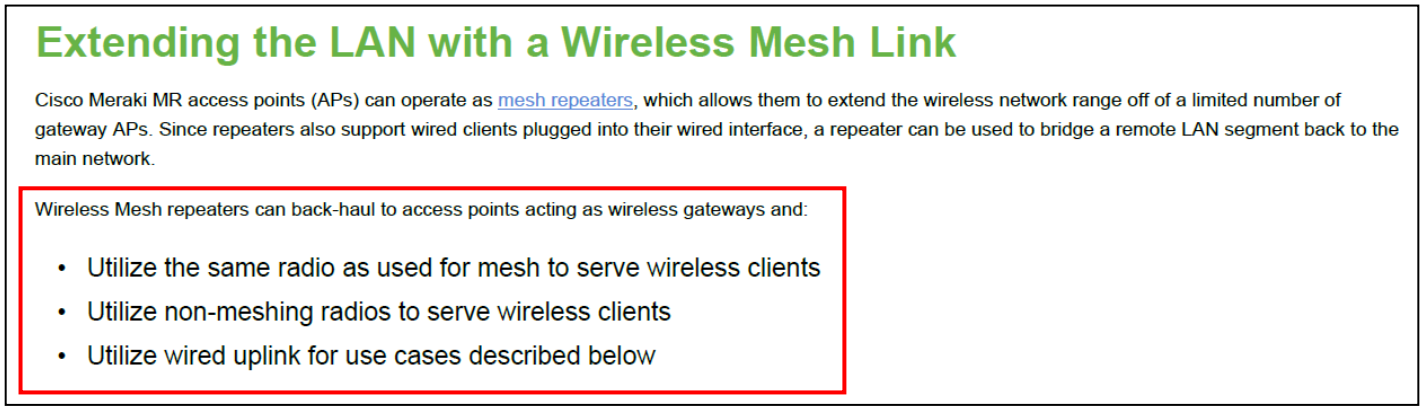


Frequency/Radio Usage

Any channel that a Meraki access point is permitted to operate on can be used for a mesh link. The channel availability for a particular access point model is subject to regulatory domain restrictions and certification. Meraki access points do not prefer one band over another, as described in the Wireless Mesh Networking "Meraki Mesh Algorithm" section. Both radios on a Meraki access point may be used concurrently for meshing while also serving wireless clients.

Any given mesh backhaul link will only use one radio, either 2.4 or 5 GHz, but not both concurrently. The gateway and/or repeater access point can serve clients and provide mesh connectivity. It is recommended not to use mesh links/radios to serve clients but rather have them remain dedicated for mesh to help maximize performance.

Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide.



Extending the LAN with a Wireless Mesh Link

Cisco Meraki MR access points (APs) can operate as mesh repeaters, which allows them to extend the wireless network range off of a limited number of gateway APs. Since repeaters also support wired clients plugged into their wired interface, a repeater can be used to bridge a remote LAN segment back to the main network.

Wireless Mesh repeaters can back-haul to access points acting as wireless gateways and:

- Utilize the same radio as used for mesh to serve wireless clients
- Utilize non-meshing radios to serve wireless clients
- Utilize wired uplink for use cases described below

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Extending_the_LAN_with_a_Wireless_Mesh_Link.

APs in Cisco's Meraki Wireless Mesh Network mesh automatically on channels with the least interference as possible using channel selection algorithms.

Access points mesh automatically on the channels with the least interference possible using channel selection algorithms. It is not recommended to manually change channels in a mesh network unless absolutely necessary. However, there are special cases which merit the reason to manually select the specific channels that a collection of access points mesh in.

Source: https://documentation.meraki.com/MR/Radio_Settings/Manually_Changing_Channels_in_a_Mesh_Network.

Further, Meraki access points support external antennas and directional configurations for mesh backhaul, enabling the radios to operate on distinct frequencies based on antenna port assignment and channel configuration. In some deployments, Distinct channels and frequencies are on a per hop basis (*e.g.*, between mesh nodes creating a backhaul mesh link).

Can I Mix Antenna Types on Outdoor APs?

It depends. On 2x2:2 outdoor access points, such as MR74/76, there are use cases for mixing antenna types, such as when creating a mesh link. For example, there might be a distant outdoor area that needs client coverage. An MR76 mesh AP can be deployed with ANT-20 omni antennas on the 2.4GHz ports to serve clients in the area around the access point, while an ANT-27 sector antenna can be connected to the 5GHz ports of the MR76 and aimed back at a gateway AP on the building to create a 5GHz backhaul mesh link.

On 4x4:4 outdoor access points, such as MR84 and MR86, all four antenna ports are dual-band and diplexed to both radios. On these APs, all antennas should be identical and have the same orientation. Do not mix omni and patch antennas or patch and sector antennas, for example, on 4x4:4 access points.

Source: https://documentation.meraki.com/MR/Other_Topics/Frequently_Asked_Questions_regarding_Cisco_Meraki_Antennas.

On information and belief, client connectivity remains uninterrupted during any transition from an isolated configuration back to a connected configuration. Resiliency and seamless failover are design features of the Meraki Wireless Mesh Network. As described in Meraki documentation, if a gateway access point loses its connection to the Internet, it will automatically fail over to repeater mode without dropping active wireless client connections. Similarly, if the SSID is configured to allow LAN access, the access point may continue to function locally, maintaining client connectivity even without external network access.

If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (maintaining mesh connectivity when the connection to the external network is lost), in substantially the same way (via a mesh node acting as gateway or repeater node), to yield the same result (continuing to bridge traffic with wireless clients staying connected). Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it are insubstantially different, and the manner of establishing that connectivity (through the use of Meraki's Mesh Network algorithm), fulfills the same claimed

	operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 9	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 1 wherein the connectivity logic contained by the nodes realigns connections between the nodes upon detection of an external network connection to form the connected configuration.	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 1 wherein the connectivity logic contained by the nodes realigns connections between the nodes upon detection of an external network connection to form the connected configuration. Meraki's connectivity logic (*e.g.*, the Meraki Mesh Routing Algorithm, Gateway Selection Logic) continually monitors and scans conditions (*e.g.*, availability of gateways with an internet connection, quality of links, number of hops, etc.) and changes routes to reflect changes in conditions such as detection of an external network connection (*e.g.*, establishing a path to the gateway.) Meraki's connectivity logic (*e.g.*, the Meraki Mesh Routing Algorithm, Gateway Selection Logic) determines the mesh node alignment when connected to an external network. Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Claim 9	Cisco Wireless Mesh Network with Meraki

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Mesh Network Components

In a mesh network, access points can be in one of two states: Gateway, or Repeater.

Gateways

Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections.

Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode.

Repeaters

Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki mesh network achieves the same function (*e.g.*, restructuring network topology in response to external connectivity), in substantially the same way (*e.g.*, using any overlayed SDN or otherwise proprietary mesh routing protocol and channel/link quality evaluation to dynamically reassign roles and reestablish optimal paths), and yields the same

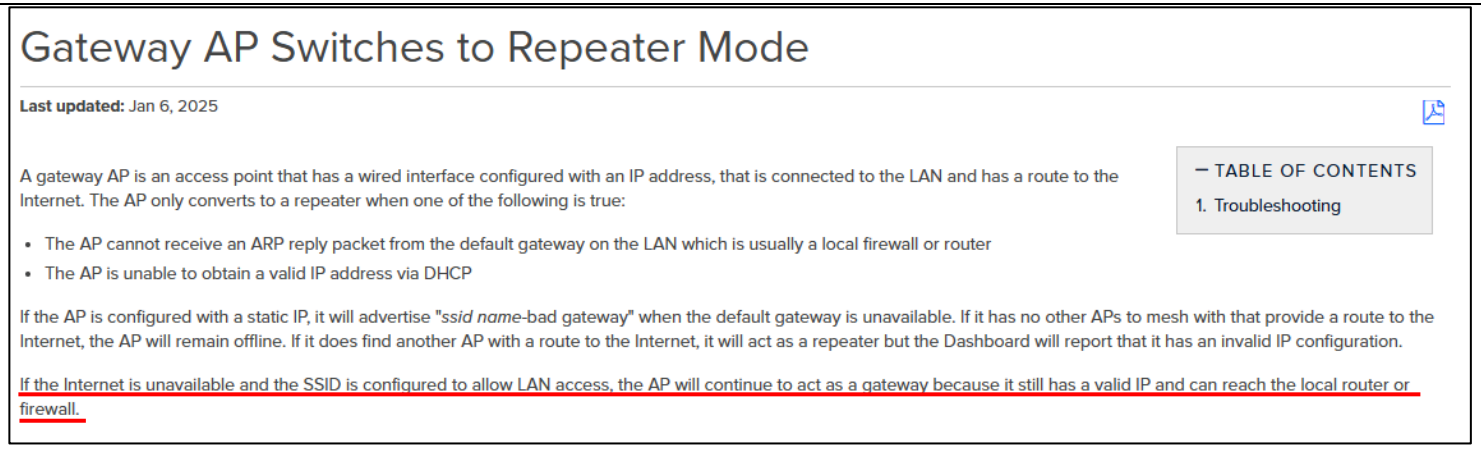
Claim 9	Cisco Wireless Mesh Network with Meraki
	result (*e.g.*, a self-healing mesh topology connected to the external network). Any difference between the claimed invention and Cisco's implementation would be considered insubstantial by a person of ordinary skill in the art. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 10	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 9 wherein solely the isolated network root node logic searches for the external network connection and establishes the external network connection.	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 9 wherein solely the isolated network root node logic searches for the external network connection and establishes the external network connection. Nodes connected to the external network via Ethernet uplink can be designated as the Gateway access points and would search to establish the external network connection. While any node with a wired connection can operate as a repeater, if the node has already been designated as a Gateway and has active downstream connections, it will continue attempting to establish or reestablish the external network connection without dropping active wireless client connections. Gateways Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's RAP (Root Access Point) performs the same function (initiating discovery of and establishing connectivity with an external controller), by substantially the same means (using connectivity logic (*e.g.*, the Meraki Mesh Routing Algorithm, Gateway Selection Logic)) and achieves the same result (re-establishing external network communication for the mesh). A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 11	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 10 wherein the previously unused uplink radio of the isolated network root node is used to connect to the external network.	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 10 wherein the previously unused uplink radio of the isolated network root node is used to connect to the external network. As shown in Cisco's Meraki documentation, when a gateway access point loses its Internet connection, it enters a failover state and begins operating as a repeater. In that state, its wired uplink is no longer used for external network connectivity. Upon detecting a nearby gateway with Internet access, the access point dynamically re-establishes external connectivity—either by reassuming the gateway role or by forming a wireless mesh link to another gateway. In either case, the previously unused uplink radio—wired or wireless—is then used to re-establish the connection to the external network. Gateways Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. If this claim element is not literally infringed, it is infringed under the doctrine of equivalents. Cisco's Meraki Wireless Mesh Network performs the same function (restoring connectivity to the external network), by substantially the same means (*e.g.*, activating an uplink interface—not previously in use), and achieves the same result (re-establishing a connection between the mesh network and the external network). The use of a wired uplink instead of a wireless radio, or establishing a new connection instead of reusing a prior one, represents an insubstantial difference. A person of ordinary skill in the art would regard any such differences between the claimed invention and the Cisco Accused Products as minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 15	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 1 wherein a first mesh network in the isolated configuration	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 1 wherein a first mesh network in the isolated configuration comprises an isolated root node and one or more isolated children nodes and a second mesh network in the isolated configuration comprises an isolated root node and at least

comprises an isolated root node and one or more isolated children nodes and a second mesh network in the isolated configuration comprises an isolated root node and at least one or more isolated children nodes and the root node of the first network establishes communication with a child node of the second network, thereby the root node of the first network becomes a child node of the second network.	one or more isolated children nodes and the root node of the first network establishes communication with a child node of the second network, thereby the root node of the first network becomes a child node of the second network. The isolated root node of a first mesh network in an isolated configuration (*e.g.*, a gateway access point that has failed over to repeater mode) can establish communication with a child node of a second mesh network in an isolated configuration (*e.g.*, a repeater node connected to a gateway with an SSID configured for LAN access), such that the root node of the first network transitions to become a child node of the second network. As shown below, the Gateway AP does not drop active wireless client connections despite losing its connection to the external network (*e.g.*, internet). As such the Gateway AP acts as the isolated network root node of the isolated configuration: **Mesh Network Components** In a mesh network, access points can be in one of two states: Gateway, or Repeater. **Gateways** Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. In one method outlined in Cisco documentation, an isolated configuration may allow LAN access such that it acts as an isolated network root that retains its function as a gateway.



Gateway AP Switches to Repeater Mode

Last updated: Jan 6, 2025

A gateway AP is an access point that has a wired interface configured with an IP address, that is connected to the LAN and has a route to the Internet. The AP only converts to a repeater when one of the following is true:

- The AP cannot receive an ARP reply packet from the default gateway on the LAN which is usually a local firewall or router
- The AP is unable to obtain a valid IP address via DHCP

If the AP is configured with a static IP, it will advertise "*ssid name*-bad gateway" when the default gateway is unavailable. If it has no other APs to mesh with that provide a route to the Internet, the AP will remain offline. If it does find another AP with a route to the Internet, it will act as a repeater but the Dashboard will report that it has an invalid IP configuration.

If the Internet is unavailable and the SSID is configured to allow LAN access, the AP will continue to act as a gateway because it still has a valid IP and can reach the local router or firewall.

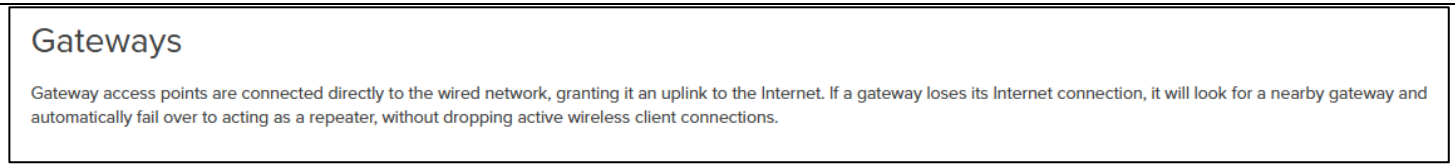
TABLE OF CONTENTS
1. Troubleshooting

Source: https://documentation.meraki.com/MR/Monitoring_and_Reporting/Gateway_AP_Switches_to_Repeater_Mode.

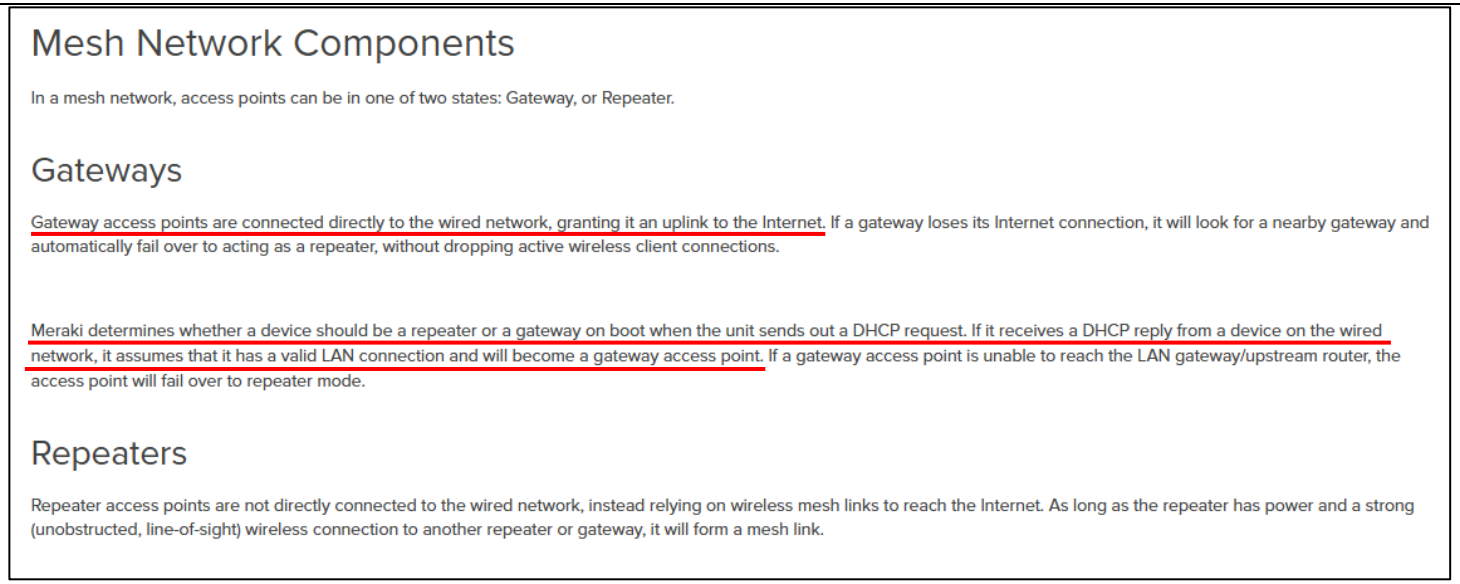
The network will continue to automatically update the mesh configuration to find the optimal path, including by making the root node of the first network the child node of the second network. See 1(a) *supra.*

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 16	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 15 wherein the unused uplink radio on the root node of the first network is used to connect to the child node of the second network.	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 15 wherein the unused uplink radio on the root node of the first network is used to connect to the child node of the second network. *See supra* discussion for claim element 15, the contents of which are incorporated herein by reference. When a gateway loses its internet connection over the wired uplink, it may fail over to repeater mode and attempt to re-establish external network access by connecting to a nearby viable gateway via its previously unused wireless uplink radio.

	 Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 18	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 1 wherein each network node further comprises a DHCP server used to assign addresses to client devices communicating with the node using the node's downlink radio.	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 1 wherein each network node further comprises a DHCP server used to assign addresses to client devices communicating with the node using the node's downlink radio. As shown in Cisco's documentation, each Meraki MR access point can operate with a built-in DHCP server that assigns IP addresses to client devices.  Source: https://documentation.meraki.com/MR/Other_Topics/Frequently_Asked_Questions_regarding_Cisco_Meraki_AP_Devices The DHCP server, and its commands, requests and responses, are responsible for routing traffic, including, on information and belief traffic to client devices via the downlink to assigned client device addresses.



Mesh Network Components

In a mesh network, access points can be in one of two states: Gateway, or Repeater.

Gateways

Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections.

Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode.

Repeaters

Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 19	Cisco Wireless Mesh Network with Meraki
The structured mesh network of claim 18 wherein the DHCP server within each node assigns IP addresses to client devices containing one or more random numbers within the IP addresses.	Meraki access points in a Meraki Wireless Mesh Network function in the structured mesh network of claim 18 wherein the DHCP server within each node assigns IP addresses to client devices containing one or more random numbers within the IP addresses. *See supra* discussion for claim element 18, the contents of which are incorporated herein by reference. On information and belief each node assigns IP addresses to client devices containing one or more random numbers within the IP addresses. If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (assigning unique IP addresses to client devices), in substantially the same way (via a DHCP server embedded within each node), to yield the same result (ensuring addressability and connectivity

	for client devices on the network). Whether the IP addresses contain random numbers or are drawn from a sequential pool or predefined subnet range is an insubstantial difference. A person skilled in the relevant art would consider any difference between the claimed invention and the Cisco Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.

Claim 35	Cisco Wireless Mesh Network with Meraki
35pre: A structured mesh network capable of isolated operation, comprising:	*See* Claim 1pre above, the contents of which are incorporated herein by reference.
35(a): at least two structured mesh nodes; wherein each structured mesh node comprises at least a connectivity logic and a radio operating on an uplink frequency and a distinct downlink frequency;	*See* Claim 1(a) and 1(b) above, the contents of which are incorporated herein by reference.
35(b): wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node on the uplink frequency and client devices or other mesh nodes connect to each node on the downlink frequency;	Cisco Meraki Access Points, including those operating in mesh configurations, comprise a connectivity logic, wherein the connectivity logic determines whether each structured mesh node connects with an external network or another node on an uplink frequency, and whether client devices or other mesh nodes connect to each node using the node's downlink frequency. Meraki's connectivity logic (*e.g.*, the Meraki Mesh Routing Algorithm, Gateway Selection Logic) determines whether a node connects to an external network (*e.g.*, through gateways) or to another node (*e.g.,* other repeaters).

Claim 35	Cisco Wireless Mesh Network with Meraki
	Meraki Mesh Algorithm Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead. As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. **Gateway Selection** Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel. Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops. While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. The connectivity logic (*e.g.*, Meraki Mesh Algorithm, Gateway Selection Logic, etc.) determines whether to connect to an external network through a gateway or to another repeater based on gateway availability and link quality. Additionally, the connectivity logic scans to determine if potentially viable gateways connect to an external network and the corresponding link quality.

Claim 35	Cisco Wireless Mesh Network with Meraki

New Gateway Selection Logic

A repeater starts looking for an access point in two cases:

1. When a gateway is down: If a gateway is not reachable for 3 minutes gateway is marked as down. Immediately after a gateway is marked as down the access point starts scanning for new gateways. Will scan the entire spectrum (including 2.4 GHz and 5 GHz) and then select the best one that is available based on metrics. Higher preference is given to the configured static channel.

> ⓘ If an access point received ARPs from the upstream gateway, access point will not go into mesh mode.

2. Repeater finds a better gateway: A repeater constantly evaluates the current channel it is operating on for better gateways but each access point will send mesh info every 15 seconds. If the repeater finds an access point with a better link quality metric of even 1 the access point will move to the new gateway.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

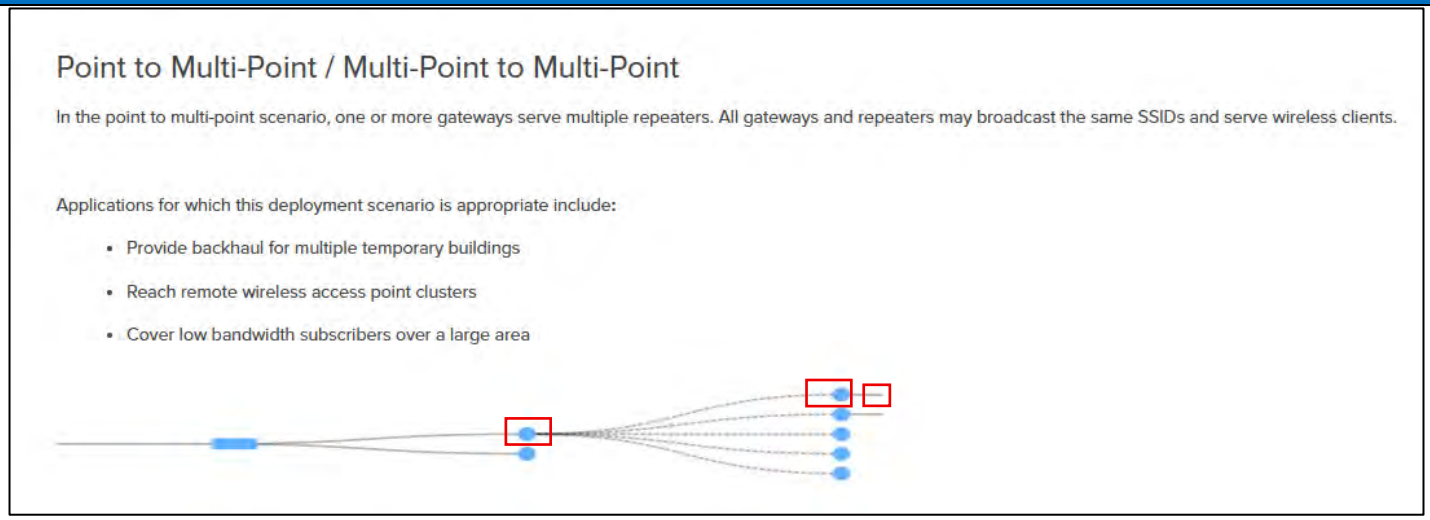
Meraki repeaters determine when to connect, via its downlink frequency, to another mesh node (*e.g.*, repeater) or to a client device (*e.g.,* a laptop or other user device).

Maximum Mesh Hops

There will be a throughput reduction (about 50%) with each "hop" in a mesh. It is recommended that a mesh network be designed for no more than one mesh hop from the gateway access point to client device.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Annotations in the image below, are representative of a nodes connected via a backhaul with a downlink client connection.

Claim 35	Cisco Wireless Mesh Network with Meraki
	 Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide. If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Cisco's implementation using the Meraki Mesh Algorithm and/or Gateway Selection Logic performs the same function (routing traffic between mesh nodes and external networks or among mesh nodes themselves), in substantially the same way (by evaluating link quality, hop count, and gateway availability to determine optimal upstream and downstream paths), to yield the same result (*i.e.*, establishing a connection to the external network or client). A person skilled in the relevant art would consider this difference in how the uplink is selected to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(c): wherein the connectivity logic in each node contains a distributed routing table;	Meraki access points in a Meraki Wireless Mesh Network function in two configurations selected depending on whether a connection to an external network is present. Meraki's self-healing mesh algorithm enables each access point to automatically detect other nodes and select the best route to a gateway. This is indicative of a distributed routing table maintained on each node to locally evaluate and update routing decisions.

Claim 35	Cisco Wireless Mesh Network with Meraki

Meraki Mesh Algorithm

Meraki devices in a mesh network configuration communicate using a proprietary routing protocol designed by Meraki. This protocol is designed specifically for wireless mesh networking and accounts for several unique characteristics of wireless networks (including variable link quality caused by noise or multi-path interference, as well as the performance impact of routing traffic through multiple hops). This protocol is also designed to provide ease of deployment while maintaining low channel overhead.

As part of the self-healing nature of meshing, the access points will automatically detect each other and select the best route to a wired gateway. All Cisco Meraki access points that support meshing will automatically try to mesh if they lose their wired connection, or be available for connections from repeaters if connected as a gateway. To that end, it is generally recommended to allow auto-channel selection in networks with repeaters.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Additionally, the documentation explicitly states that each repeater access point "develops a table of all the detected gateways and their corresponding link quality metrics," which describes a routing table used to dynamically determine optimal paths within the mesh network.

Gateway Selection

Once the access point goes in the mesh mode, The access point scans all channels to collect info from all neighbors. If a valid neighbor (in-network access point or Meraki access point) is found, it goes to that channel. The configured channel has higher precedence if a valid neighbor is found on it. If no valid neighbor is found at all from all channels, it stays on the configured channel.

Based on the scan results, the repeater access point develops a table of all the detected gateways and their corresponding link quality metrics. Additionally, the number of hops are also considered and preference is given to gateways with lower hops. Once the access point hears all the neighboring access points, it finalizes a route based on the link quality and the number of hops.

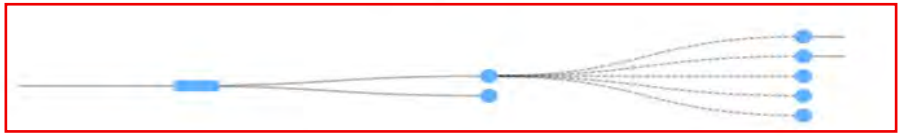
While it is not possible to select which frequency band should be used for meshing, it is possible to manually adjust channel selections to direct the access point toward a desired behavior. To do this, refer to the article on manually changing channels in a mesh network. If it is desired for two access points to mesh on 5Ghz as opposed to 2.4Ghz, then the access points should both be set to the same 5Ghz channel, but different 2.4Ghz channels. Keep in mind though that a frequency band cannot be allocated specifically for meshing, and both bands will still be available for servicing clients unless the SSID is configured to use the 5Ghz band only.

Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (routing mesh traffic based on network topology and link metrics), in substantially the same way (via each mesh node maintaining a routing table comprising available gateways and associated link quality metrics), to yield the same result (adaptive and efficient routing decisions across the mesh network). Whether or not the routing tables are distributed from a centralized source or independently generated at each node is an insubstantial difference. A person skilled in the relevant art would consider any difference between the claimed invention and the Accused Products to be minor and insubstantial.

Claim 35	Cisco Wireless Mesh Network with Meraki
	MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(d): wherein the structured mesh network functions in two configurations selected depending on whether a connection to an external network is present; in the first connected configuration the structured mesh network includes at least one structured mesh node's uplink connection comprising a connection to an external network; and	Meraki access points in a Meraki Wireless Mesh Network function in two configurations selected depending on whether a connection to an external network is present. In the first configuration (*e.g.*, the gateway is connected to the internet), at least one access point functions as a gateway node with an active uplink connection to the external network (*e.g.*, the internet). **Mesh Network Components** In a mesh network, access points can be in one of two states: Gateway, or Repeater. **Gateways** Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking. If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (external network connectivity), in substantially the same way (via the mesh node's uplink connection establishing a path to a router/controller/switch), to yield the same result (connectivity beyond the local mesh network). Whether or not a router, controller, or switch itself is considered part of the "external network," the resulting connectivity functionally extends beyond the mesh. Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (upstream communication) are insubstantially different, and the manner of establishing that connectivity (through the use of Meraki's Mesh Network algorithm, Gateway Selection Logic, Meraki Dashboard, Meraki Cloud management, or Custom User Configurations) fulfills the same claimed operation A person skilled in the relevant

Claim 35	Cisco Wireless Mesh Network with Meraki
	art would consider any difference between the claimed invention and the Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(e): in the second isolated configuration none of the structured mesh nodes' uplink connections comprises a connection to an external network, and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node as isolated root children and descendant nodes forming a tree configuration;	Meraki access points in a Meraki Wireless Mesh Network function in a second isolated configuration wherein none of the structured mesh nodes' uplink connection comprises a connection to an external network, and one of the structured mesh nodes acts as an isolated network root of the isolated configuration and all remaining nodes connect to the isolated network root node as isolated root children and descendent nodes forming a tree configuration. As shown below, the Gateway AP does not drop active wireless client connections despite losing its connection to the external network (*e.g.*, internet). As such the Gateway AP acts as the isolated network root node of the isolated configuration: **Mesh Network Components** In a mesh network, access points can be in one of two states: Gateway, or Repeater. **Gateways** Gateway access points are connected directly to the wired network, granting it an uplink to the Internet. If a gateway loses its Internet connection, it will look for a nearby gateway and automatically fail over to acting as a repeater, without dropping active wireless client connections. Meraki determines whether a device should be a repeater or a gateway on boot when the unit sends out a DHCP request. If it receives a DHCP reply from a device on the wired network, it assumes that it has a valid LAN connection and will become a gateway access point. If a gateway access point is unable to reach the LAN gateway/upstream router, the access point will fail over to repeater mode. **Repeaters** Repeater access points are not directly connected to the wired network, instead relying on wireless mesh links to reach the Internet. As long as the repeater has power and a strong (unobstructed, line-of-sight) wireless connection to another repeater or gateway, it will form a mesh link. Source: https://documentation.meraki.com/MR/Wi-Fi_Basics_and_Best_Practices/Wireless_Mesh_Networking.

Claim 35	Cisco Wireless Mesh Network with Meraki
	In one method outlined in Cisco documentation, an isolated configuration may be achieved when the SSID on the gateway AP is configured to allow LAN access such that it acts as an isolated network root that retains its function as a gateway. Gateway AP Switches to Repeater Mode **Last updated:** Jan 6, 2025 A gateway AP is an access point that has a wired interface configured with an IP address, that is connected to the LAN and has a route to the Internet. The AP only converts to a repeater when one of the following is true: – TABLE OF CONTENTS 1. Troubleshooting • The AP cannot receive an ARP reply packet from the default gateway on the LAN which is usually a local firewall or router • The AP is unable to obtain a valid IP address via DHCP If the AP is configured with a static IP, it will advertise "*ssid name*-bad gateway" when the default gateway is unavailable. If it has no other APs to mesh with that provide a route to the Internet, the AP will remain offline. If it does find another AP with a route to the Internet, it will act as a repeater but the Dashboard will report that it has an invalid IP configuration. If the Internet is unavailable and the SSID is configured to allow LAN access, the AP will continue to act as a gateway because it still has a valid IP and can reach the local router or firewall. Source: https://documentation.meraki.com/MR/Monitoring_and_Reporting/Gateway_AP_Switches_to_Repeater_Mode. Furthermore, all remaining nodes remain connected in a tree configuration. Point to Multi-Point / Multi-Point to Multi-Point In the point to multi-point scenario, one or more gateways serve multiple repeaters. All gateways and repeaters may broadcast the same SSIDs and serve wireless clients. Applications for which this deployment scenario is appropriate include: • Provide backhaul for multiple temporary buildings • Reach remote wireless access point clusters • Cover low bandwidth subscribers over a large area 

Claim 35	Cisco Wireless Mesh Network with Meraki
	Source: https://documentation.meraki.com/MR/Deployment_Guides/Mesh_Deployment_Guide. If this claim element is not literally infringed, it is nonetheless infringed under the doctrine of equivalents. Meraki's system achieves the same function (maintaining mesh connectivity when the connection to the external network is lost), in substantially the same way (via a mesh node acting as gateway or repeater node), to yield the same result (continuing to bridge traffic with wireless clients staying connected). Thus, the functional result of enabling external network communication is preserved, the means of accomplishing it (upstream communication) are insubstantially different, and the manner of establishing that connectivity (through the use of Meraki's Mesh Network algorithm, Meraki Dashboard, Meraki Cloud management, or Custom User Configurations) fulfills the same claimed operation. A person skilled in the relevant art would consider any difference between the claimed invention and the Accused Products to be minor and insubstantial. MeshDynamics asserts that this claim element is a software limitation. Pursuant to Judge Gilstrap's standard discovery order, MeshDynamics will identify what Cisco source code satisfies this limitation 30 days after Cisco's complete production of source code or at such other time as the parties may agree to.
35(f): wherein clients of the at least two structured mesh nodes in a second isolated configuration retain full connectivity within the structured network during the isolated configuration and wherein said client connectivity is uninterrupted during any transition from the isolated configuration to the first connected configuration; wherein distinct downlink frequencies are used by nodes for communication with one or more nodes lower on the tree configuration.	*See* Claim 1(f), the contents of which are incorporated herein by reference.

Exhibit 11:Accused Instrumentalities ("Cisco Accused Products")

The Cisco Accused Products include any and all of the Cisco controller, access point, cloud, and software products that are capable of mesh or bridge mode (including mesh access point ("MAP" or repeater) or root access point ("RAP" or gateway) roles, with or without a Cisco wireless LAN controller) made, used, imported, offered for sale or sold in the United States since May 5, 2019. This includes all products that can be used in a Cisco Unified Wireless Mesh Network, Meraki Wireless Mesh Network, Ultra Reliable Wireless Backhaul, Point-to-Point/Point-to-Multipoint Bridge and any other Cisco mesh architecture(s).

This includes, but is not limited to, any and all of the Cisco hardware and software products that can be used in mesh or bridge mode identified in the list below. These products were identified after a reasonable search of publicly-available information and on information and belief comprise a complete list of Cisco's mesh-capable products and licenses since May 5, 2019. MeshDynamics reserves the right to amend or supplement this list following discovery.

Catalyst 9100 Series Wireless Access Points

Catalyst 9105 Series
- C9105
- C9105i
- C9105w
- C9105AX
- C9105AXI
- C9105AXI-x
- C9105AXI-EWC
- C9105AXI-EWC-x
- C9105AXW
- C9105AXW-x
- C9105AXW-EWC
- C9105AXW-EWC-x
- C9105AXIT-x
- C9105AXWT-x

Catalyst 9115 Series
- C9115
- C9115I
- C9115E
- C9115AXI
- C9115AXI-EWC
- C9115AXI-x
- C9115AXI-EWC-x
- C9115AXE
- C9115AXE-EWC
- C9115AXE-x
- C9115AXE-EWC-x

Catalyst 9117 Series
- C9117

- C9117AXI
- C9117AXI-EWC

Catalyst 9120 Series
- C9120AX
- C9120AXI
- C9120AXI-x
- C9120AXI-EWC-x
- C9120AXE
- C9120AXE-x
- C9120AXE-EWC-x
- C9120AXP
- C9120AXP-x
- C9120AXP-EWC-x

Catalyst 9124 Series
- C9124AX
- C9124AXI
- C9124AXE
- C9124AXD
- C9124AXI-x
- C9124AXI-EWC-x
- C9124AXD-x
- C9124AXD-EWC-x
- C9124AXE-x
- C9124AXE-EWC-x

Catalyst 9130 Series
- C9130AX
- C9130AXI
- C9130AXI-EWC
- C9130AXI-x
- C9130AXI-EWC-x
- C9130AXE
- C9130AXE-x
- C9130AXE-EWC
- C9130AXE-EWC-x
- C9130AXE-STA

Catalyst 9136 Series
- C9136I
- C9136I-x

Catalyst 9160 Series
- C9162
- C9162I
- C9163E
- C9164
- C9164I
- C9166
- C9166I

- C9166D1

Catalyst 9170 Series
- C9172I
- C9172H
- C9176E
- C9176I/D1
- C9178E
- C9178I
- C9179F

Catalyst Wireless Series (including for use with the Meraki cloud platform)
- CW9162
- CW9162I-x
- CW9162I-MR
- CW9163E
- CW9164
- CW9164I-x
- CW9164I-MR
- CW9166
- CW9166I-x
- CW9166D1-x
- CW9166D1
- CW9166I-MR
- CW9166D1-MR
- CW9172H
- CW9172I
- CW9176I/D1
- CW9178I

Catalyst Industrial Wireless Access Points
- IW9167E
- IW9167E-HZ
- IW9167I
- IW9165
- IW9165D
- IW9165E

Cisco Meraki and Meraki Go Wireless Access Points

MR20
MR20-HW
MR28
MR28-HW
MR30H
MR30H-HW
MR32-HW
MR33
MR33-HW
MR36
MR36H

MR36-HW
MR42
MR42-HW
MR42E
MR42E-HW
MR44
MR44-HW
MR45
MR45-HW
MR46
MR46-HW
MR46E-HW
MR46E
MR52
MR52-HW
MR53
MR53-HW
MR53E
MR53E-HW
MR55
MR56
MR56-HW
MR57
MR57-HW
MR70
MR70-HW
MR72
MR72-HW
MR74
MR74-HW
MR76
MR76-HW
MR78
MR78-HW
MR84
MR84-HW
MR86
MR86-HW
GR10
GR12
GR60
GR62

Cisco Aironet Series Wireless Access Points

Aironet 4800

AIR-AP4800

Aironet 3800 Series (3802I/E/P)

- AIR-AP3802I
- AIR-AP3802E
- AIR-AP3802P

Aironet 2800 Series (2802I/E)

- AIR-AP2802I
- AIR-AP2802E

Aironet 1850 Series

- AIR-AP1852I
- AIR-AP1852E

Aironet 1830 Series

- AIR-AP1832I
- AIR-AP1832E

Aironet 1815 Series

- AIR-AP1815I
- AIR-AP1815W
- AIR-AP1815T
- AIR-AP1815M

Aironet 1540 Series

- AIR-AP1542I
- AIR-AP1542E

Aironet 1560 Series

- AIR-AP1562I
- AIR-AP1562E
- AIR-AP1562D

Aironet 1570 Series

- AIR-CAP1572E
- AIR-CAP1572IC

Aironet 1530 Series

- AIR-CAP1532I
- AIR-CAP1532E

Aironet 3700 Series

- AIR-CAP3702I
- AIR-CAP3702E

Aironet 3600 Series

- AIR-CAP3602I
- AIR-CAP3602E

Aironet 3500 Series

- AIR-CAP3502I
- AIR-CAP3502E

Aironet 2700 Series

- AIR-CAP2702I
- AIR-CAP2702E

Aironet 2600 Series

- AIR-CAP2602I
- AIR-CAP2602E

Aironet 1700 Series

AIR-CAP1702I
AIR-CAP1702E
Aironet 1600 Series
AIR-CAP1602I
AIR-CAP1602E

Cisco Business Wireless (CBW) Access Points

Cisco Business 100 Series
CBW140AC
CBW141ACM
CBW145AC
CBW150AX
CBW151AXM
CBW153AXM
Cisco Business 200 Series
CBW240AC

Wireless Controllers

Cisco controllers for the wireless access points listed above
Embedded controllers in any Catalyst access point
Cisco AireOS-Based Controllers, including:
AIR-AP1832I (AIR-AP1832I-A-K9 (or -A-K9 / -B-K9))
AIR-AP1852I (AIR-AP1852I-A-K9)
AIR-AP2802I (AIR-AP2802I-A-K9)
AIR-AP3802I (AIR-AP3802I-A-K9)
AIR-AP1562I (AIR-AP1562I-A-K9)
AIR-AP1562D (AIR-AP1562D-A-K9)
AIR-AP1542I (AIR-AP1542I-A-K9)
AIR-AP1815W (AIR-AP1815W-A-K9)
AIR-AP1815I (AIR-AP1815I-A-K9)
Cisco Wireless LAN Controllers
Cisco 2500 Series Controllers
AIR-CT2504
AIR-CT2504-5
AIR-CT2504-15
AIR-CT2504-25
AIR-CT2504-50
AIR-CT2504-HA
Cisco 3500 Series Controllers
Cisco 3504 Wireless Controller
AIR-CT3504
AIR-CT3504-HA
Cisco 5500 Series Controllers
5508
AIR-CT5508-12
AIR-CT5508-25
AIR-CT5508-50

- AIR-CT5508-100
- AIR-CT5508-250
- AIR-CT5508-500
- AIR-CT5508-750
- 5520
- AIR-CT5520
- AIR-CT5520-HA

Cisco 6500 Series Controllers

Cisco 7500 Series Controllers

- 7510 Wireless Controllers

Cisco 7600 Series Controllers

Cisco 8500 Series Controllers

- 8510 Wireless Controllers
- 8540 Wireless Controllers

Cisco 9800 Series Controllers

- Catalyst 9800 Wireless Controllers
- Catalyst 9800
- Catalyst 9800-40
- Catalyst 9800-L
- Catalyst 9800-CL
- Catalyst 9800-80
- CW9800M (Virtual)
- CW9800H1/H2 (Virtual)

Cisco Integrated Services Routers with Controller Network Module (ISR-WLC), including:

- Cisco 3600 Series Integrated Services Routers with Controller Network Module
- WS-C3650
- WS-C365024TD
- WS-C3650-48F

Cisco Wireless Controller Modules, including:

- NME-WLC
- L-AIR-CTVM-5-K9 (Virtual AireOS WLC)

Cisco Wireless Services Modules (WiSM & WiSM2)

- Cisco 6500 Series Wireless Services Module (WiSM & WiSM2)
- Cisco 7600 Series Router Wireless Services Module (WiSM & WiSM2)

Cisco Integrated Wireless LAN Controller Switches, including:

- Catalyst 9000 Series Switches
- Catalyst 9300 Series Switches
- Catalyst 9300-48U
- Catalyst 9300-24U

Cisco Switches with FlexConnect or Embedded Wireless Controllers, including:

- Cisco 9000 Series Switches
- Catalyst 9300
- Catalyst 9400
- Catalyst 9500

Cisco Unified Wireless Network Controllers, including:

- AireOS WLCs and Catalyst 9800

Cisco Wireless Control System (WCS) Controllers, including:
- 3504, 5508, 8540 WLCs
- Aironet products

Cisco Software

Cisco Prime Infrastructure
Cisco Prime Network Control System
Cisco DNA Center
Cisco Catalyst Center
Cisco Meraki
Cisco Meraki Cloud Platform
Operating Systems
- IOS XE
- IOS XE Wireless
- AirOS

Controller Software
- Virtual Wireless LAN Controller (vWLC)
- 9800-CL

Licenses: Catalyst/Aironet Access Points and WLC

Cisco Prime Infrastructure License
All licenses identified at https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9800-series-wireless-controllers/impl-catalyst-9800-wrls-lan-ctrlr-licens-faq.html, including the relevant AIR Network license and AIR DNA license for each access point that joins to the Catalyst 9800 WLC or other Catalyst Controller:
- AIR-DNA
- AIR-DNA-E
- AIR-DNA-E-1Y
- AIR-DNA-E-3Y
- AIR-DNA-E-5Y
- AIR-DNA-E-7Y
- AIR-DNA-A
- AIR-DNA-A-1Y
- AIR-DNA-A-3Y
- AIR-DNA-A-5Y
- AIR-DNA-A-7Y

SW9124AX-CAPWAP-K9
SW9124AX-MESH-K9
SW9124AXE-CW-K9
SW9124AXE-MESH-K9
SW9124AXE-EWC-M-K9
Cisco Smart Licensing/Smart Account
2500
- LIC-CT2504-1A
- LIC-CT2504-5A
- LIC-CT2504-25A
- L-LIC-CT2504-1A

L-LIC-CT2504-5A
L-LIC-CT2504-25A

3500

CON-SNT-CT3504
CON-3YR-CT3504
CON-ECMU-CT3504

5500

5508
CON-SNT-CT5508
L-CT5508
LIC-CT5508-50A
LIC-CT5508-100A
LIC-CT5508-250A
LIC-CT5508-500A
LIC-CT5508-750A
5520
Cisco Smart Net Total Care
CON-SNT-CT5520
CON-3ECMU-CT552

Licensing: Meraki MR access points

LIC-MR-A
LIC-MR-E
LIC-ENT-1YR
LIC-ENT-3YR
LIC-ENT-5YR
LIC-ENT-7YR
LIC-ENT-10YR
LIC-MR-ADV-1Y
LIC-MR-ADV-3Y
LIC-MR-ADV-5Y
LIC-MR-ADV-7Y
LIC-MR-ADV-10Y
LIC-MR-UPGR-1Y
LIC-MR-UPGR-3Y
LIC-MR-UPGR-5Y
LIC-WIFI-1YR
LIC-WIFI-3YR
LIC-WIFI-5YR
LIC-WIFI-7YR
LIC-WIFI-10YR

Cisco Networking Subscriptions

LIC-CW-A
LIC-CW-E