

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

CLOUD BYTE LLC,

Plaintiff,

v.

DELL INC., and DELL
TECHNOLOGIES, INC.,

Defendants.

Civil Action No. 2:24-cv-00637

JURY TRIAL DEMANDED

CLOUD BYTE’S PRELIMINARY INFRINGEMENT CONTENTIONS

Plaintiff Cloud Byte LLC (“Cloud Byte” or “Plaintiff”) submits its Patent Rule (“P.R.”) 3-1 and 3-2 Disclosures to Defendants Dell Inc. and Dell Technologies, Inc. (collectively “Dell”) pursuant to the Court’s Order granting the Joint Motion for Extension of Time to Serve Infringement and Invalidity Contentions (Dkt. 27).

Cloud Byte submits these Disclosures based upon information it has acquired to date, as it presently understands this information and the significance thereof, without yet having had the benefit of formal discovery, claims, or defenses from Dell. Accordingly, Cloud Byte reserves the right to modify, amend, retract, and/or further supplement the disclosures made herein as additional evidence and information becomes available, after the Court has construed the claim terms of the Asserted Patents and as otherwise allowed by the Local Rules and Federal Rules of Civil Procedure. Cloud Byte also reserves the right to accuse different instrumentalities or

alternative literal and/or equivalent infringing elements in Dell's Accused Instrumentalities as disclosed below.

Cloud Byte submits its P.R. 3-1 and 3-2 Disclosures to Dell with respect to U.S. Patent No. 7,739,544 ("the '544 Patent"), U.S. Patent No. 9,482,632 ("the '632 Patent"), U.S. Patent No. 9,560,177 ("the '177 Patent"), U.S. Patent No. 9,629,265 ("the '265 Patent"), U.S. Patent No. 9,651,320 ("the '320 Patent"), U.S. Patent No. 9,900,249 ("the '249 Patent"), and U.S. Patent No. 10,628,273 ("the '273 Patent") (collectively, the "Asserted Patents").

Although Cloud Byte is concurrently producing documents to Dell with its Infringement Contentions herewith, to date, Dell has not produced documents to Cloud Byte. Accordingly Cloud Byte's investigation regarding infringement is necessarily preliminary and ongoing, and these disclosures are therefore based on public information that Cloud Byte has been able to obtain to date regarding the function and operation of the Accused Instrumentalities (as defined in the attached P.R. 3-1(c) claim charts), together with Cloud Byte's present understanding of the meaning and scope of the asserted claims.¹ Cloud Byte reserves the right to supplement and/or amend its disclosures as additional information is ascertained through discovery and after the Court has construed the Asserted Patents.

A. CLOUD BYTE'S RULE 3-1 DISCLOSURES

(a) Each claim of each patent in suit that is allegedly infringed by each opposing party

¹ In many portions of Cloud Byte's 3-1 and 3-2 disclosures, Cloud Byte has gone beyond the disclosure requirements required by P.R. 3-1 and 3-2. These additional disclosures and citations are made for the benefit of Dell. Cloud Byte in no way confines its contentions to the specific citations of evidence and explicitly reserves the right to rely on additional or different evidence as the case moves forward through fact discovery, expert discovery, and trial. Furthermore, any omission of any other specific citation to the Accused Instrumentalities or documents related thereto does not constitute a waiver of Cloud Byte's right to raise any issues related to the Accused Instrumentalities or other documents related thereto at a later date.

Cloud Byte asserts that Dell has infringed and continues to infringe, has contributed to and continues to contribute to the infringement of, and/or has actively induced and continues to actively induce others to infringe at least the following claims of the Asserted Patents:

Claims 1-9, 11-20, 22 of the '544 Patent (*see* Exs. 1–4);

Claims 1, 3, 5-9 of the '632 Patent (*see* Ex. 5);

Claims 1-20 of the '177 Patent (*see* Ex. 6);

Claims 1-8 of the '265 Patent (*see* Exs. 7–8);

Claims 1-8 of the '320 Patent (*see* Ex. 9);

Claims 1-27 of the '249 Patent (*see* Ex. 10); and

Claims 1-21 of the '273 Patent (*see* Ex. 11).

Cloud Byte reserves the right to supplement and/or amend its disclosures to identify additional claims infringed by Dell that are revealed through discovery, or in response to Dell's P.R. 3-4 production, or as permitted under the Patent Rules.

(b) Separately for each asserted claim, each accused apparatus, product, device, process, method, act, or other instrumentality ("Accused Instrumentality") of each opposing party of which the party is aware. This identification shall be as specific as possible. Each product, device, and apparatus must be identified by name or model number, if known. Each method or process must be identified by name, if known, or by any product, device, or apparatus which, when used, allegedly results in the practice of the claimed method or process

Dell infringes Cloud Byte's patented technologies across many different parts of its business, including in: all ethernet switches, SmartFabric services, and edge networking products and servers incorporating same ("Dell Networking Products"); all products and services relating to Redundant Array of Independent Disks (RAID) technology and servers incorporating same ("Dell RAID Products"); and all server systems supporting high availability or fault tolerance features, abnormality detection, airflow cooling, and/or temperature control ("Dell Server

Products”), that Dell makes, uses, sells, offers to sell, and/or imports into the United States and in the associated consulting and support services that it provides to its customers (as disclosed in the P.R. 3-1(c) claim charts) that are capable of and/or configured to perform the identified infringing functionalities.

The Accused Products for the ’544 Patent include at least the Dell RAID Products and the following Dell products: Dell PowerEdge servers that use PowerEdge RAID Controllers, which include at least the 13th generation (C4130, C6300, C6320, C6320p, FC430, FC630, FC830, FD332, FM120x4 (for PE FX2/FX2s), FX2/FS2s, M630, M630 (for PE VRTX), M830, M830 (for PE VRTX), R230, R330, R430, R530, R530xd, R630, R730, R730xd, R830, R930, T30, T130, T330, T430, T630); the 14th generation (C4140, C6400, C6420, FC640, M640, M640 (for PE VRTX), MX5016s, MX7000, MX740c, MX840c, R240, R340, R440, R540, R640, R6415, R740, R740xd, R740xd2, R7415, R7425, R840, R940, R940xa, T40, T140, T340, T440, T640, XE2420, XE7100, XE7420, XE7440, XR2); the 15th generation (C6520, C6525, MX750c, R250, R350, R450, R550, R650, R650xs, R6515, R6525, R750, R750xa, R750xs, R7515, R7525, T150, T350, T550, XE8545, XR4000r, XR4000w, XR4000z, XR4510c, XR4520c, XR11, XR12); the 16th generation (C6600, C6615, C6620, HS5610, HS5620, MX760c, R260, R360, R660, R660xs, R6615, R6625, R760, R760xa, R760xd2, R760xs, R7615, R7625, R860, R960, T160, T360, T560, XE8640, XE9640, XE9680, XR5610, XR7620, XR8000r, XR8610t, XR8620t); and the 17th generation (R670 CSP Edition, R770 CSP Edition) of the PowerEdge servers, and all other Dell servers that use PowerEdge RAID Controllers with substantially similar functionality. The PowerEdge RAID Controllers can be either hardware or software based and include at least the following Dell products: Series 9 Family (H830 Adapter, H730P Adapter, H730P MX, H730P Mini Mono, H730P Mini Blade, H730 Adapter, H730 Mini Mono, H730 Mini Blade, H330

Adapter, H330 Mini Mono, H330 Mini Blade, S130 Software RAID); Series 10 Family (H840 Adapter, H745P MX Adapter, H740P Adapter, H740P Mini Mono, S150 Software RAID, S140 Software RAID, H745 Adapter, H745 Front, H345 Adapter, H345 Front); Series 11 Family (H755 Adapter, H750 Adapter, H755N Front (NVMe Only), H755 Front (SAS/SATA Only), HBA355 Adapter (internal), HBA355 Front (internal), HBA350 Adapter (internal), HBA355e Adapter (external), HBA350 MX (internal), H350 Adapter, H355 Adapter, H355 Front, S150 Software RAID); and Series 12 Family (PERC H965I Adapter, PERC H965I Front, PERC H965I MX, PERC H965E Adapter, HBA465I Adapter, HBA465I Front, S160 Software RAID). Other Dell PowerEdge RAID Controllers may also be used in conjunction with the Dell PowerEdge Rack Servers to infringe the '544 Patent. Charts establishing Dell's infringement, each of which show an exemplary PowerEdge RAID Controller compatible with each of the 13th, 14th, 15th, and 16th Generations of the Dell PowerEdge Rack Servers, have been provided herewith. See Ex. 1 (S130 Controller, which is compatible with 13th Generation PowerEdge servers); Ex. 2 (S140 Controller, which is compatible with 14th Generation PowerEdge servers); Ex. 3 (S150 Controller, which is compatible with 15th Generation PowerEdge servers); Ex. 4 (S160 Controller, which is compatible with 16th Generation PowerEdge servers). The S130, S140, S150, and S160 controllers are merely exemplary, as there are numerous PowerEdge RAID Controllers that are used by the Dell PowerEdge servers. See <https://www.dell.com/support/kbdoc/en-us/000131648/list-ofpowerededge-raid-controller-perc-types-for-dell-emc-systemsand>.

The Accused Products for the '632 Patent include at least the Dell Server Products and the following Dell products: Dell PowerEdge servers that use Integrated Dell Remote Access Controller ("iDRAC") 9 with firmware version 4.x through 7.x, which include at least the 14th generation (C4140, C6400, C6420, FC640, M640, M640 (for PE VRTX), MX5016s, MX7000,

MX740c, MX840c, R240, R340, R440, R540, R640, R6415, R740, R740xd, R740xd2, R7415, R7425, R840, R940, R940xa, T40, T140, T340, T440, T640, XE2420, XE7100, XE7420, XE7440, XR2), the 15th generation (C6520, C6525, MX750c, R250, R350, R450, R550, R650, R650xs, R6515, R6525, R750, R750xa, R750xs, R7515, R7525, T150, T350, T550, XE8545, XR4000r, XR4000w, XR4000z, XR4510c, XR4520c, XR11, XR12), the 16th generation (C6600, C6615, C6620, HS5610, HS5620, MX760c, R260, R360, R660, R660xs, R6615, R6625, R760, R760xa, R760xd2, R760xs, R7615, R7625, R860, R960, T160, T360, T560, XE8640, XE9640, XE9680, XR5610, XR7620, XR8000r, XR8610t, XR8620t), and the 17th generation (R670 CSP Edition, R770 CSP Edition) of the PowerEdge servers, and all other Dell servers that incorporate iDRAC 9 with firmware version 4.x through 7.x with substantially similar functionality.

The Accused Products for the '177 Patent include at least the Dell Networking Products and the following Dell products: Dell switches that support OpenFlow (and server systems including same), which at least include Dell PowerSwitch S-series (S3048-ON, S3100, S3124F, S3124P, S3148P, S4048T-ON, S4048-ON, S4100-ON, S4112T-ON, S4128F-ON, S4128T-ON, S4147FE-ON, S4148U-ON, S4148T-ON, S4248FB-ON, F4248FBL-ON, S4180, S4820T, S5000, S5048F-ON, S5148F-ON, S5200-ON, S5212F-ON, S5232F-ON, S5224F-ON, S5296F-ON, S5448F-ON, S6000, S6000-ON, S6010-ON, S6100-ON); N-series (N1100-ON, N1108EP-ON, N1108T-ON, N1124P-ON, N1124T-ON, N1148P-ON, N1148T-ON, N1500, N1524, N1524P, N1548, N1548P, N2000, N2024, N2024P, N2048, N2048P, N2100-ON, N2128PX-ON, N2224PX-ON, N2224X-ON, N2248PX-ON, N2248X-ON, N3000, N3024, N3024EF-ON, N3024EP-ON, N3024ET-ON, N3024F, N3024P, N3048, N3048EP-ON, N3048ET-ON, N3048P, N3100-ON, N3132PX-ON, N3208PX-ON, N3224T-ON, N3224P-ON, N3224F-ON, N3224PX-ON, N3248P-ON, N3248X-ON, N3248PXE-ON, N3248TE-ON, N3248X-ON, N4000, N4032,

N4032F, N4064, N4064F); Z-series (Z9000, Z9100, Z9100-ON, Z9264F-ON, Z9332F-ON, Z9432F-ON, Z9500, Z9664F-ON, Z9864F-ON); E-series (E3224F-ON, E3248P-ON, and E3248PXE-ON); FN IOM; and M-series (MXL), and all other Dell switches with substantially similar functionality.

The Accused Products for the '265 Patent include at least the Dell Server Products and the following Dell products: Dell PowerEdge servers, which include at least the 14th Generation PowerEdge servers (R940); 15th Generation PowerEdge servers (R650, R6525, R750, R750xa, R7525); 16th Generation PowerEdge servers (R660, R960, R6615, R6625, R760, R7625, R860); and 17th Generation PowerEdge servers (R670 CSP Edition, R770 CSP Edition), and all other Dell servers with substantially the same internal layout as these systems.

The Accused Products for the '320 Patent include at least the Dell Server Products and the following Dell products: Dell PowerEdge servers, which include at least the 13th generation (C4130, C6300, C6320, C6320p, FC430, FC630, FC830, FD332, FM120x4 (for PE FX2/FS2s), FX2/FS2s, M630, M630 (for PE VRTX), M830, M830 (for PE VRTX), R230, R330, R430, R530, R530xd, R630, R730, R730xd, R830, R930, T30, T130, T330, T430, T630), the 14th generation (C4140, C6400, C6420, FC640, M640, M640 (for PE VRTX), MX5016s, MX7000, MX740c, MX840c, R240, R340, R440, R540, R640, R6415, R740, R740xd, R740xd2, R7415, R7425, R840, R940, R940xa, T40, T140, T340, T440, T640, XE2420, XE7100, XE7420, XE7440, XR2), the 15th generation (C6520, C6525, MX750c, R250, R350, R450, R550, R650, R650xs, R6515, R6525, R750, R750xa, R750xs, R7515, R7525, T150, T350, T550, XE8545, XR4000r, XR4000w, XR4000z, XR4510c, XR4520c, XR11, XR12), the 16th generation (C6600, C6615, C6620, HS5610, HS5620, MX760c, R260, R360, R660, R660xs, R6615, R6625, R760, R760xa, R760xd2, R760xs, R7615, R7625, R860, R960, T160, T360, T560, XE8640, XE9640, XE9680,

XR5610, XR7620, XR8000r, XR8610t, XR8620t), and the 17th generation (R670 CSP Edition, R770 CSP Edition) of the PowerEdge servers, and all other Dell servers with substantially similar functionality.

The Accused Products for the '249 Patent include at least the Dell Networking Products and the following Dell products: Dell switches that support OpenFlow (and server systems including same), which at least include Dell PowerSwitch E-series (E3224F-ON, E3248P-ON, E3248PXE-ON), FN-series (FN IOM), MX-series (MXL), N-series (N1108EP-ON, N1108T-ON, N1124P-ON, N1124T-ON, N1148P-ON, N1148T-ON, N1524, N1524P, N1548, N1548P, N2024, N2024P, N2048, N2048P, N2128PX-ON, N2224PX-ON, N2224X-ON, N2248PX-ON, N2248X-ON, N3024, N3024EF-ON, N3024EP-ON, N3024ET-ON, N3024F, N3024P, N3048, N3048EP-ON, N3048ET-ON, N3048P, N3132PX-ON, N3208PX-ON, N3224T-ON, N3224P-ON, N3224F-ON, N3224PX-ON, N3248P-ON, N3248X-ON, N3248PXE-ON, N3248TE-ON, N3248X-ON, N4032, N4032F, N4064, N4064F), S-series (S3048-ON, S3124, S3124F, S3124P, S3148, S3148P, S4048-ON, S4048T-ON, S4112F-ON, S4112T-ON, S4128F-ON, S4128T-ON, S4148F-ON, S4148FE-ON, S4148U-ON, S4148T-ON, S4248FB-ON, S4248FBL-ON, S4810, S4820T, S5000, S5048F-ON, S5148F-ON, S5212F-ON, S5232F-ON, S5224F-ON, S5248F-ON, S5296F-ON, S5448F-ON, S6000, S6000-ON, S6010-ON, S6100-ON), and Z-series (Z9000, Z9100-ON, Z9264F-ON, Z9332F-ON, Z9432F-ON, Z9500, Z9664F-ON, Z9864F-ON) switch models, and all other Dell switches with substantially similar functionality.

The Accused Products for the '273 Patent include at least the Dell Server Products and the following Dell products: Dell PowerEdge servers, which include at least the 14th generation (C4140, C6400, C6420, FC640, M640, M640 (for PE VRTX), MX5016s, MX7000, MX740c, MX840c, R240, R340, R440, R540, R640, R740, R740xd, R740xd2, R840, R940, R940xa,

R6415, R7415, R7425, T40, T140, T340, T440, T640, XE2420, XE7100, XE7420, XE7440, XR2), the 15th generation (C6520, C6525, MX750c, R250, R350, R450, R550, R650, R650xs, R750, R750xa, R750xs, R6515, R6525, R7515, R7525, T150, T350, T550, XE8545, XR4000r, XR4000w, XR4000z, XR11, XR12, XR4510c, XR4520c), the 16th generation (C6615, C6620, HS5610, HS5620, MX760c, R260, R360, R660, R660xs, R760, R760xa, R760xd2, R760xs, R860, R960, R7615, R7625, T160, T360, T560, XE8640, XE9640, XE9680, XR5610, XR7620, XR8000, XR8610t, XR8620t), and the 17th generation (R670 CSP Edition, R770 CSP Edition) of PowerEdge models, and all other Dell servers with substantially similar functionality.

Based on the information that it has been able to obtain to date, Cloud Byte separately identifies each Dell Accused Instrumentality for each Asserted Claim applicable to Dell in Cloud Byte's P.R. 3-1(c) claim charts attached as Exhibits 1 to 11.

Cloud Byte believes that discovery from Dell and others will provide additional evidence of the nature and/or scope of Dell's infringement. Accordingly, Cloud Byte reserves the right to supplement and/or amend this disclosure to identify and accuse additional Dell Accused Instrumentalities released, developed, or made available by Dell after the date on which these disclosures are served, or of which Cloud Byte was not aware at the time of these disclosures.

(c) A chart identifying specifically where each element of each asserted claim is found within each Accused Instrumentality, including for each 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

Pursuant to P.R. 3-1(c), claim charts for Dell's infringement of the Asserted Claims of the Patents-in-Suit are set forth as Exhibits 1 to 11. The charts are based on information currently known to Cloud Byte and contain illustrative (not exhaustive) examples of presently known infringement of the Asserted Claims.

As explained above, Cloud Byte believes that discovery from Dell and others will provide additional evidence of the nature and/or scope of Dell's infringement. Accordingly, Cloud Byte reserves the right to amend, supplement, or otherwise revise the Accused Instrumentalities and its infringement contentions and claim charts.

(d) Whether each element of each asserted claim is claimed to be literally present or present under the doctrine of equivalents in the Accused Instrumentality

Pursuant to P.R. 3-1(d), Cloud Byte asserts that, under the proper construction of the Asserted Claims and their claim terms, every limitation of the Asserted Claims of the Patents-in-Suit is literally present in Dell's instrumentalities accused of infringing the claims. However, to the extent any claim element or step is found to be not literally present in or performed by the Accused Instrumentalities, based on the Court's claim construction or Dell's arguments, such claim element or step is satisfied under the doctrine of equivalents because any difference between such claim element or step and the element or step of the Accused Instrumentalities is insubstantial. In other words, the element or step of the Accused Instrumentalities performs substantially the same function, in substantially the same way, to achieve substantially the same result. Further detail regarding the claim limitations Cloud Byte asserts that Dell infringes under the doctrine of equivalents can be found in Cloud Byte's P.R. 3-1(c) claim charts attached as Exhibits 1 to 11.

(e) For any patent that claims priority to an earlier application, the priority date to which each asserted claim allegedly is entitled

The Asserted Claims of the '544 Patent are entitled to a priority date no later than March 10, 2005. The Asserted Claims of the '632 Patent are entitled to a priority date no later than September 5, 2012. The Asserted Claims of the '177 Patent are entitled to a priority date no later than February 17, 2011. The Asserted Claims of the '265 Patent are entitled to a priority date no later than March 5, 2012. The Asserted Claims of the '320 Patent are entitled to a priority date no

later than March 2, 2012. The Asserted Claims of the '249 Patent are entitled to a priority date no later than September 14, 2009. The Asserted Claims of the '273 Patent are entitled to a priority date no later than January 30, 2015. For each Asserted Claim, Cloud Byte reserves the right to assert priority to earlier conception and reduction to practice dates earlier than the earliest priority date.

Cloud Byte may rely on additional evidence to demonstrate earlier dates of conception, reduction to practice, and/or diligence in reducing to practice. *See EMG Tech., LLC v. Chrysler Grp., LLC*, No. 6:12-CV-259, 2013 WL 12147662, at *2 (E.D. Tex. July 3, 2013).

(f) If a party claiming patent infringement wishes to preserve the right to rely, for any purpose, on the assertion that its own apparatus, product, device, process, method, act, or other instrumentality practices the claimed invention, the party must identify, separately for each asserted claim, each such apparatus, product, device, process, method, act, or other instrumentality that incorporates or reflects that particular claim

Cloud Byte is still investigating this matter, however, at this time Cloud Byte is not relying on the assertion that its own apparatus, product, device, process, method, act, or other instrumentality of its own practices the claimed inventions. Cloud Byte reserves the right to supplement and/or amend this disclosure to identify any apparatus, product, device, process, method, act, or other instrumentality of its own that practices the Asserted Claims of which Cloud Byte was not aware at the time of these disclosures.

B. CLOUD BYTE'S RULE 3-2 PRODUCTION OF DOCUMENTS

(a) Documents (e.g., contracts, purchase orders, invoices, advertisements, marketing materials, offer letters, beta site testing agreements, and third party or joint development agreements) 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

Cloud Byte is not presently aware of documents responsive to P.R. 3-2(a). Cloud Byte's investigation is ongoing and Cloud Byte will supplement its document production in the event responsive documents are located.

(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 P. R. 3-1(e), whichever is earlier

In addition to information included in the file histories noted below, documents responsive to P.R. 3-2(b) have been produced to Dell as follows: CLOUDBYTE0002647–CLOUDBYTE0002781. As discovery is ongoing and Cloud Byte's search for documents is ongoing, Cloud Byte reserves the right to supplement its production and disclosures as information is ascertained during discovery.

Cloud Byte's investigation is ongoing and Cloud Byte will supplement its document production in the event responsive documents are located.

(c) A copy of the file history for each patent in suit

Documents responsive to P.R. 3-2(c) have been produced to Dell as follows:

The '544 Patent file history: CLOUDBYTE0000691–CLOUDBYTE0000914;

The '632 Patent file history: CLOUDBYTE0000915–CLOUDBYTE0001190;

The '177 Patent file history: CLOUDBYTE0001191–CLOUDBYTE0001353;

The '265 Patent file history: CLOUDBYTE0001354–CLOUDBYTE0001672;

The '320 Patent file history: CLOUDBYTE0001673–CLOUDBYTE0001929;

The '249 Patent file history: CLOUDBYTE0001930–CLOUDBYTE0002646; and

The '273 Patent file history: CLOUDBYTE0000001–CLOUDBYTE0000690.

Cloud Byte's investigation is ongoing and Cloud Byte will supplement its document production in the event responsive documents are located.

Dated: November 21, 2024

/s/ Yury Kapgan

Yury Kapgan (admitted in EDTX)
yurykapgan@quinnemanuel.com
Ryan Goldstein (admitted in EDTX)
ryangoldstein@quinnemanuel.com
**QUINN EMANUEL URQUHART &
SULLIVAN, LLP**
865 S. Figueroa St., 10th Floor
Los Angeles, California 90017
Telephone: (213) 443-3000
Facsimile: (213) 443-3100

Brian Mack (admitted in EDTX)
brianmack@quinnemanuel.com
**QUINN EMANUEL URQUHART &
SULLIVAN, LLP**
50 California Street, 22nd Floor
San Francisco, CA 94111
Telephone: (415) 875-6400
Facsimile: (415) 875-6700

Of Counsel:

Claire Abernathy Henry
Texas State Bar No. 24053063
Andrea Fair
Texas State Bar No. 24078488
MILLER FAIR HENRY PLLC
1507 Bill Owens Parkway
Longview, TX 75604
Telephone: (903) 757-6400
Fax: (903) 757-2323
claire@millerfairhenry.com
andrea@millerfairhenry.com

Attorneys for Plaintiff Cloud Byte LLC.

CERTIFICATE OF SERVICE

I certify that a true and correct copy of the above document was served by email with exhibits served via FTP on counsel of record for Dell on November 21, 2024.

/s/ Yury Kapgan

Cloud Byte's Patent Initial Disclosures for Dell
Exhibit 9
UNITED STATES PATENT NO. 9,651,320
INFRINGEMENT CHART FOR CLAIMS 1-8¹

'320 Patent Claim² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found³ (P.R. 3-1(b)-(c))
1[pre] ICT (information and Communication)	The Accused Dell '320 Products include information and communication technology ("ICT") equipment having an electronic component.

¹ The infringement contentions provided herein are based on information obtained to date and may not be exhaustive. Cloud Byte's investigation of Dell's infringement is ongoing. Cloud Byte reserves the right to supplement and/or amend these disclosures to identify additional Asserted Claims (P.R. 3-1(a)), to identify additional Accused Instrumentalities (P.R. 3-1(b)), and to further identify where each element of each Asserted Claim is found in each Accused Instrumentality (P.R. 3-1(c)), including on the basis of discovery obtained from Dell and from third-parties during the course of this litigation.

² All infringement contentions set forth herein for any independent patent claims are hereby incorporated by reference into the infringement contentions alleged for any dependent patent claims that depend on such independent claims, as if fully set forth therein.

³ The Accused Instrumentalities and associated exhibits discussed and/or cited for any claim herein are representative in all material respects of all other Accused Instrumentalities identified for that claim (*e.g.*, although various versions of the Accused Instrumentalities may have immaterial differences in their hardware, firmware, and/or software configuration, the cited references are believed to be illustrative of all such accused servers). For example, on information and belief, the Accused Instrumentalities, including the 13th to 17th generation of Dell PowerEdge servers that implement the Integrated Dell Remote Access Controller ("iDRAC"), are substantially similar to those shown here. On information and belief, there are no material differences between the Accused Instrumentalities, including the 13th to 17th generation of Dell PowerEdge servers that implement the iDRAC, that would substantially alter the infringement read shown here. The illustrative documentation cited herein is for the Dell R760xd, XR8000, XR8610t, or XR8620t PowerEdge servers and can be found at the following URLs: <https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-r760xd2-technical-guide.pdf> (CLOUDBYTE0020666); <https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-xr8000-technical-guide.pdf> (CLOUDBYTE0020907). Other generations of Dell PowerEdge servers, including the 13th to 17th generation of Dell PowerEdge servers, that incorporate iDRAC also infringe the claims of the '320 patent. *See, e.g.*, CLOUDBYTE0004144, CLOUDBYTE0005625, CLOUDBYTE0005627, CLOUDBYTE0005734, CLOUDBYTE0005737, CLOUDBYTE0005738; CLOUDBYTE0010314 –


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																											
Technology) equipment having an electronic component, the ICT equipment comprising: ⁴	See, e.g., <table><tr><th>Feature</th><th colspan="2">Technical Specifications for XR8000r XR8610t and, XR8620t</th></tr><tr><td>Processor</td><td colspan="2">1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores</td></tr><tr><td>Memory</td><td colspan="2"> 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only </td></tr><tr><td>Storage</td><td colspan="2"> 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t </td></tr><tr><td>Chassis</td><td colspan="2"> - XR8000r </td></tr><tr><td>Sleds</td><td colspan="2"> - XR8610t and XR8620t </td></tr><tr><td>Hot swap Redundant Power Supplies</td><td> 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC </td><td> 1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V </td></tr><tr><td>Heating and Cooling Options</td><td colspan="2"> - Air cooling - Optional Heater Manager </td></tr><tr><td>Fans</td><td colspan="2">Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR8620t sled</td></tr></table> https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/powerededge-xr8000-spec-sheet.pdf (“Dell XR8610t and XR8620t Spec Sheet” or CLOUDBYTE0010598), 2	Feature	Technical Specifications for XR8000r XR8610t and, XR8620t		Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores		Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only		Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t		Chassis	XR8000r		Sleds	XR8610t and XR8620t		Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V	Heating and Cooling Options	- Air cooling - Optional Heater Manager		Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR8620t sled	
Feature	Technical Specifications for XR8000r XR8610t and, XR8620t																											
Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores																											
Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only																											
Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t																											
Chassis	XR8000r																											
Sleds	XR8610t and XR8620t																											
Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V																										
Heating and Cooling Options	- Air cooling - Optional Heater Manager																											
Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR8620t sled																											

CLOUDBYTE0010599, CLOUDBYTE0020425-CLOUDBYTE0021003; CLOUDBYTE0004521; CLOUDBYTE0004885; CLOUDBYTE0005253; CLOUDBYTE0004147, CLOUDBYTE0021009.


⁴ Cloud Byte’s inclusion of any claim preamble in this claim chart should not be interpreted as an admission that the preamble is limiting. Cloud Byte reserves the right to take the position that the claim preambles are limiting or not limiting on a claim-by-claim basis.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																																																																																																																																																																																																									
	<table><tr><th>Rack Server</th><th>R750</th><th>R750xa</th><th>R650</th><th>R752s</th><th>R652s</th><th>R751s</th><th>R651s</th><th>R750xs</th><th>R650xs</th><th>R450</th><th>R550</th><th>XR11</th><th>XR12</th><th>R350</th><th>R250</th></tr><tr><td>Key attributes</td><td>Outstanding performance for the most demanding workloads</td><td>Highly intensive GPU workloads</td><td>High scalability, optimized workload performance</td><td>Powerful performance and flexibility</td><td>Dense virtualization</td><td>Powerful performance and scalability</td><td>High density compute</td><td>Purpose-built 2U server for growing scale-out solutions</td><td>Purpose-built, full performance 1U server for dense, fast growing scale-out solutions</td><td>Value and density-focused, built for general purpose IT</td><td>Versatile, value-optimized, ruggedized, built for general purpose IT</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Powerful performance in 1U server for productivity and data intensive applications</td><td>Powerful compute for common business applications and streamlines productivity</td></tr><tr><td>Target workloads</td><td>Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments</td><td>AI, ML or DL training or inferencing, HPC, and virtualization environments</td><td>Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments</td><td>All flash SDS, VDI, and data analytics</td><td>HPC, Dense VDI, and Virtualization</td><td>SDS, Virtualization, and Data Analytics</td><td>Virtualization, HCI and NFV</td><td>Virtualization, medium VM density or VDI, and scale-out database workloads</td><td>Virtualization, cloud, scale-out database and high-performance compute workloads</td><td>Small IT infrastructure, light VM, small business specific workloads</td><td>Small IT infrastructure, light VM density, small business specific workloads</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads</td></tr><tr><td>Type of processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor</td><td colspan="2">2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor</td><td colspan="2">1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor</td><td colspan="2">1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor</td><td colspan="2">1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores</td></tr><tr><td>Memory (DDR4 DIMM slots & max)</td><td colspan="2">32 (8 TB)</td><td colspan="2">32 (4 TB)</td><td colspan="2">16 (2 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2">8 (1 TB)</td><td colspan="2">4 (128 GB)</td></tr><tr><td>Disk drives up to:</td><td>8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)</td><td>6 x 2.5" 8 x 2.5"</td><td>4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)</td><td>26 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>12 x 2.5" 4 x 3.5"</td><td>24 x 2.5" 12 x 3.5" 8 x 2.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 8 x 3.5"</td><td>4 x 3.5" 8 x 2.5" 16 x 2.5" 8 x 3.5"</td><td>4 x 3.5" 8 x 2.5" 16 x 2.5" 8 x 3.5"</td><td>4 x 2.5" 6 x 2.5"</td><td>4 x 2.5" 6 x 2.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)</td></tr><tr><td>NVMe drives up to:</td><td>24</td><td>8</td><td>12</td><td>24</td><td>12</td><td>24</td><td>10</td><td>8</td><td>10</td><td>N/A</td><td>4</td><td>6</td><td>N/A</td><td>N/A</td></tr><tr><td>Gen4 PCIe slots up to:</td><td>8</td><td>8</td><td>3</td><td>8</td><td>3</td><td>2</td><td>1</td><td>5</td><td>3</td><td>2</td><td>3</td><td>3</td><td>5</td><td>3</td><td>2</td></tr><tr><td>Gen3 PCIe slots up to:</td><td colspan="4">N/A</td><td colspan="2">2</td><td>1</td><td>1</td><td colspan="2">N/A</td><td>1</td><td colspan="4">N/A</td></tr><tr><td>Accelerator support up to:</td><td>2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW</td><td>4 x 150 W SW or 4 x 300 W DW 2 x 75 W SW</td><td>3 x 75 W SW or 3 x 300 W DW or 6 x 75 W SW</td><td>3 x 300 W DW or 6 x 75 W SW</td><td>3 x SW</td><td>4 x SW; 1 x DW; 1 x FPGA</td><td>1 x SW</td><td colspan="2">N/A</td><td colspan="2">N/A</td><td>2 x 75 W SW or 150 W SW 2 x 300 W DW</td><td colspan="3">N/A</td></tr><tr><td>Rack height (U)</td><td>2</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td></tr><tr><td>Integrated security</td><td colspan="15">TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks</td></tr></table>														Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250	Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, ruggedized, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity	Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML or DL training or inferencing, HPC, and virtualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads	Type of processor	2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2 nd or 3 rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2 nd or 3 rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3 rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores		Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)		16 (1 TB)		8 (1 TB)		4 (128 GB)		Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 16 x 2.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 16 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A	4	6	N/A	N/A	Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2	Gen3 PCIe slots up to:	N/A				2		1	1	N/A		1	N/A				Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW 2 x 75 W SW	3 x 75 W SW or 3 x 300 W DW or 6 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		N/A		2 x 75 W SW or 150 W SW 2 x 300 W DW	N/A			Rack height (U)	2	2	1	2	1	2	1	2	1	1	2	1	2	1	1	Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks														
Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250																																																																																																																																																																																											
Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, ruggedized, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity																																																																																																																																																																																											
Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML or DL training or inferencing, HPC, and virtualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads																																																																																																																																																																																											
Type of processor	2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2 nd or 3 rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2 nd or 3 rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3 rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores																																																																																																																																																																																													
Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)		16 (1 TB)		8 (1 TB)		4 (128 GB)																																																																																																																																																																																													
Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 16 x 2.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 16 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)																																																																																																																																																																																												
NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A	4	6	N/A	N/A																																																																																																																																																																																												
Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2																																																																																																																																																																																											
Gen3 PCIe slots up to:	N/A				2		1	1	N/A		1	N/A																																																																																																																																																																																														
Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW 2 x 75 W SW	3 x 75 W SW or 3 x 300 W DW or 6 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		N/A		2 x 75 W SW or 150 W SW 2 x 300 W DW	N/A																																																																																																																																																																																													
Rack height (U)	2	2	1	2	1	2	1	2	1	1	2	1	2	1	1																																																																																																																																																																																											
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks																																																																																																																																																																																																									
	https://i.dell.com/sites/csdocuments/Product_Docs/en/Dell-EMC-PowerEdge-Rack-Servers-Quick-Reference-Guide.pdf (“PowerEdge Reference Guide” or CLOUDBYTE0004144), 1.																																																																																																																																																																																																									
	<table><tr><th>Rack Server</th><th>R940</th><th>R940xa</th><th>R840</th><th>R740xd</th><th>R740</th><th>R740xd2</th><th>R640</th><th>R540</th><th>R440</th><th>R340</th><th>R240</th></tr><tr><td>Key attributes</td><td>Powerful performance</td><td>Extreme acceleration</td><td>Turbocharge data analytics</td><td>Scalable storage performance</td><td>Optimal application performance</td><td>Enterprise content server</td><td>Performance and density</td><td>Balanced and adaptable</td><td>Scale-out computing</td><td>Accelerate business growth</td><td>Compute made simple</td></tr><tr><td>Target workloads</td><td>In-memory databases</td><td>GPU database acceleration and machine learning</td><td>Data-intensive workloads, HFT, and dense virtualization</td><td>SDS, service providers, and big data servers</td><td>VDI and cloud workloads</td><td>Media streaming and SDS</td><td>Dense scale-out data center computing and storage</td><td>Mail messaging and virtualization</td><td>HPC, web tech, and scale-out infrastructure</td><td>ROBO productivity and data-intensive applications</td><td>Small business and service provider workloads</td></tr><tr><td>Type of processor</td><td colspan="3">4 x 2nd Generation Intel® Xeon® Scalable processors</td><td colspan="3">2 x 2nd Generation Intel® Xeon® Scalable processors</td><td colspan="3">1 x 2nd Generation Intel® Xeon® Scalable processors</td><td colspan="2">1 x Intel Xeon E-2200, Intel Core i3®, Intel Pentium®, or Intel Celeron® processor</td></tr><tr><td>Memory (DDR4 DIMM slots & max)</td><td colspan="3">48 (15.36 TB)</td><td colspan="3">24 (7.68 TB)</td><td colspan="3">16 (1 TB)</td><td colspan="2">4 (64 GB)</td></tr><tr><td>Disk drives up to:</td><td colspan="3">24 x 2.5"</td><td colspan="3">32 x 2.5" 16 x 3.5"</td><td colspan="3">26 x 3.5" 16 x 3.5" + 10 x 2.5"</td><td colspan="2">8 x 2.5" 4 x 3.5"</td></tr><tr><td>NVMe drives up to:</td><td colspan="3">12</td><td colspan="3">24</td><td colspan="3">10</td><td colspan="2">4</td></tr><tr><td>Gen4 PCIe slots up to:</td><td colspan="3">13</td><td colspan="3">8</td><td colspan="3">5</td><td colspan="2">2</td></tr><tr><td>Gen3 PCIe slots up to:</td><td colspan="3">12</td><td colspan="3">6</td><td colspan="3">3</td><td colspan="2">2</td></tr><tr><td>Accelerator support up to:</td><td colspan="3">N/A</td><td colspan="3">4 x DW GPUs or 4 x DW or 8 x SW FPGAs</td><td colspan="3">2 x DW GPUs or 2 x SW or 4 x SW FPGAs</td><td colspan="2">N/A</td></tr><tr><td>Rack height (U)</td><td colspan="3">3</td><td colspan="3">4</td><td colspan="3">2</td><td colspan="2">1</td></tr><tr><td>Integrated security</td><td colspan="11">TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks</td></tr></table>														Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240	Key attributes	Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple	Target workloads	In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads	Type of processor	4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors			1 x 2 nd Generation Intel® Xeon® Scalable processors			1 x Intel Xeon E-2200, Intel Core i3®, Intel Pentium®, or Intel Celeron® processor		Memory (DDR4 DIMM slots & max)	48 (15.36 TB)			24 (7.68 TB)			16 (1 TB)			4 (64 GB)		Disk drives up to:	24 x 2.5"			32 x 2.5" 16 x 3.5"			26 x 3.5" 16 x 3.5" + 10 x 2.5"			8 x 2.5" 4 x 3.5"		NVMe drives up to:	12			24			10			4		Gen4 PCIe slots up to:	13			8			5			2		Gen3 PCIe slots up to:	12			6			3			2		Accelerator support up to:	N/A			4 x DW GPUs or 4 x DW or 8 x SW FPGAs			2 x DW GPUs or 2 x SW or 4 x SW FPGAs			N/A		Rack height (U)	3			4			2			1		Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks																																																						
Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240																																																																																																																																																																																															
Key attributes	Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple																																																																																																																																																																																															
Target workloads	In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads																																																																																																																																																																																															
Type of processor	4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors			1 x 2 nd Generation Intel® Xeon® Scalable processors			1 x Intel Xeon E-2200, Intel Core i3®, Intel Pentium®, or Intel Celeron® processor																																																																																																																																																																																																
Memory (DDR4 DIMM slots & max)	48 (15.36 TB)			24 (7.68 TB)			16 (1 TB)			4 (64 GB)																																																																																																																																																																																																
Disk drives up to:	24 x 2.5"			32 x 2.5" 16 x 3.5"			26 x 3.5" 16 x 3.5" + 10 x 2.5"			8 x 2.5" 4 x 3.5"																																																																																																																																																																																																
NVMe drives up to:	12			24			10			4																																																																																																																																																																																																
Gen4 PCIe slots up to:	13			8			5			2																																																																																																																																																																																																
Gen3 PCIe slots up to:	12			6			3			2																																																																																																																																																																																																
Accelerator support up to:	N/A			4 x DW GPUs or 4 x DW or 8 x SW FPGAs			2 x DW GPUs or 2 x SW or 4 x SW FPGAs			N/A																																																																																																																																																																																																
Rack height (U)	3			4			2			1																																																																																																																																																																																																
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks																																																																																																																																																																																																									
1[a] a cooling fan:	The Accused Dell '320 Products include a cooling fan. See <i>supra</i> , claim 1[pre].																																																																																																																																																																																																									

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																											
	See, e.g., <table><tr><th>Feature</th><th colspan="2">Technical Specifications for XR8000r XR8610t and, XR8620t</th></tr><tr><td>Processor</td><td colspan="2">1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores</td></tr><tr><td>Memory</td><td colspan="2"> 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only </td></tr><tr><td>Storage</td><td colspan="2"> 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t </td></tr><tr><td>Chassis</td><td colspan="2"> - XR8000r </td></tr><tr><td>Sleds</td><td colspan="2"> - XR8610t and XR8620t </td></tr><tr><td>Hot swap Redundant Power Supplies</td><td> 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC </td><td> 1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V </td></tr><tr><td>Heating and Cooling Options</td><td colspan="2"> - Air cooling - Optional Heater Manager </td></tr><tr><td>Fans</td><td colspan="2">Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled</td></tr></table> Dell XR8610t and XR8620t Spec Sheet, 2	Feature	Technical Specifications for XR8000r XR8610t and, XR8620t		Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores		Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only		Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t		Chassis	XR8000r		Sleds	XR8610t and XR8620t		Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V	Heating and Cooling Options	- Air cooling - Optional Heater Manager		Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled	
Feature	Technical Specifications for XR8000r XR8610t and, XR8620t																											
Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores																											
Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only																											
Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t																											
Chassis	XR8000r																											
Sleds	XR8610t and XR8620t																											
Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V																										
Heating and Cooling Options	- Air cooling - Optional Heater Manager																											
Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																											

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thoughtful Thermal Design Changes High Performance Fans – New higher performance fans have been added to increase the amount of cool air pushed through the system, targeted at higher power semiconductors (CPUs, GPUs and NVMe drives). A 3-tier approach to fan performance is employed to address the trend of increasing total power. This includes standard fans and high performance silver and gold fans, which increase the airflow capability versus the previous-generation fans (see Figure 1).  Figure 1 – Graph illustrating the increased fan airflow (in CFM) for high performance silver and gold fans compared to baseline fans https://infohub.delltechnologies.com/en-us/p/next-generation-powered-edge-servers-thoughtful-thermal-design/ (“Next Gen PowerEdge Thermal Design” or CLOUDBYTE0005738), 1.
1[b] a first temperature sensor that detects a component temperature of the electronic component included in the ICT equipment;	The Accused Dell '320 Products include a first temperature sensor that detects a component temperature of the electronic component included in the ICT equipment. <i>See supra</i>, claim 1[a]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/powerededge-xr8000-technical-guide.pdf (Dell XR8610t and XR8620t Technical Guide or CLOUDBYTE0020907), 39 Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/powerededge-r760xd2-technical-guide.pdf (Dell R760xd2 Technical Guide or CLOUDBYTE0020666), 37

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.   Figure 2. Thermal chamber in the Dell performance testing lab https://infohub.delltechnologies.com/en-us/p/understanding-thermal-design-and-capabilities-for-the-poweredge-xr8000-server/ (“Dell XR8000 Thermal Design” or CLOUDBYTE0021009).

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed <table><thead><tr><th>Temperature (°C)</th><th>System Power consumption (W)</th><th>Fan speed (% of total)</th></tr></thead><tbody><tr><td>0</td><td>216</td><td>5</td></tr><tr><td>10</td><td>219</td><td>5</td></tr><tr><td>20</td><td>254</td><td>14</td></tr><tr><td>30</td><td>331</td><td>42</td></tr><tr><td>40</td><td>512</td><td>64</td></tr><tr><td>50</td><td>748</td><td>86</td></tr><tr><td>55</td><td>1014</td><td>100</td></tr></tbody></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	System Power consumption (W)	Fan speed (% of total)	0	216	5	10	219	5	20	254	14	30	331	42	40	512	64	50	748	86	55	1014	100
Temperature (°C)	System Power consumption (W)	Fan speed (% of total)																							
0	216	5																							
10	219	5																							
20	254	14																							
30	331	42																							
40	512	64																							
50	748	86																							
55	1014	100																							


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. https://infohub.delltechnologies.com/en-us/p/thermal-manage-features-and-benefits/ (“Thermal Manage” or CLOUDBYTE0005627).


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with thermal control mechanisms. It is divided into a 'Cold Aisle' on the left and a 'Hot Aisle' on the right. In the Cold Aisle, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points towards the server rack. Above the rack, three dots represent 'CFM#N', 'CFM#2', and 'CFM#1', with a bracket labeled '1 Airflow reporting (per server)'. A red arrow labeled 'Exhaust Temp Control (per server)' points away from the rack towards the Hot Aisle, with a bracket labeled '3'. A double-headed arrow labeled '2 Custom Delta-T' connects the inlet and exhaust airflows. Below the rack, a bracket labeled 'PCIe Thermal Control' encompasses two elements: '4 Inlet PCIe Temp' (with a fan icon) and '5 PCIe Airflow (LFM)' (with a blue arrow icon). A legend at the bottom left states 'CFM = Airflow in Cubic Feet/min'. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also has a checkbox for 'Set Maximum Exhaust Temperature Limit' (checked) with a dropdown set to 'Default, 70 °C (158.0 °F)'. Below that is a 'Target Exhaust Temperature Limit' set to 70°C (158.0°F). At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' dropdown set to 'Default' with a range of 25 - 100. 'Apply' and 'Reboot' buttons are at the bottom right. Thermal Manage.
1[c] a second temperature sensor that detects a temperature of an intake air; and	The Accused Dell '320 Products include a second temperature sensor that detects a temperature of an intake air. <i>See supra</i>, claim 1[b]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: • Optimized thermal design: The system layout is architected for optimum thermal design. • System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. • Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. • Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. • User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. • Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. • Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed The graph shows two data series: System Power consumption (W) in orange and Fan speed (% of total) in blue. The x-axis represents Temperature in degrees Celsius from 0 to 60. The y-axis represents Power Consumption in Watts from 0 to 1200. The orange line shows a significant increase in power consumption as temperature rises, starting at 216W at 0°C and reaching 1014W at 55°C. The blue line shows a much slower, steady increase in fan speed from 5% at 0°C to 100% at 55°C. <table><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. A circled '1' is next to the text 'Airflow reporting (per server)'. Server Rack: A central server rack is shown with a circled '2' at its base labeled 'Custom Delta-T'. Hot Aisle: On the right, a red arrow points away from the rack, labeled 'Exhaust Temp Control (per server)' with a circled '3'. PCIe Thermal Control: At the bottom, a bracket groups two elements: - A circled '4' next to 'Inlet PCIe Temp' with a small server component icon. - A circled '5' next to 'PCIe Airflow (LFM)' with a blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also has a 'Set Maximum Exhaust Temperature Limit' dropdown set to 'Default, 70 °C (158.0 °F)' and a 'Target Exhaust Temperature Limit' set to '70°C (158.0°F)'. At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' dropdown set to 'Default' with a range of 25 - 100. 'Apply' and 'Reboot' buttons are at the bottom right. Thermal Manage.
1[d] a microprocessor including:	The Accused Dell '320 Products include a microprocessor. <i>See supra</i>, claim 1[c]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																																																																																																																																																																																																	
	<table><tr><th>Feature</th><th colspan="2">Technical Specifications for XR8000r XR8610t and, XR8620t</th></tr><tr><td>Processor</td><td colspan="2">1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores</td></tr><tr><td>Memory</td><td colspan="2"> 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only </td></tr><tr><td>Storage</td><td colspan="2"> 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t </td></tr><tr><td>Chassis</td><td colspan="2"> - XR8000r </td></tr><tr><td>Sleds</td><td colspan="2"> - XR8610t and XR8620t </td></tr><tr><td>Hot swap Redundant Power Supplies</td><td> 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC </td><td> 1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V </td></tr><tr><td>Heating and Cooling Options</td><td colspan="2"> - Air cooling - Optional Heater Manager </td></tr><tr><td>Fans</td><td colspan="2">Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled</td></tr></table>	Feature	Technical Specifications for XR8000r XR8610t and, XR8620t		Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores		Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only		Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t		Chassis	XR8000r		Sleds	XR8610t and XR8620t		Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V	Heating and Cooling Options	- Air cooling - Optional Heater Manager		Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																																																																																																																																																																							
Feature	Technical Specifications for XR8000r XR8610t and, XR8620t																																																																																																																																																																																																	
Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores																																																																																																																																																																																																	
Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only																																																																																																																																																																																																	
Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t																																																																																																																																																																																																	
Chassis	XR8000r																																																																																																																																																																																																	
Sleds	XR8610t and XR8620t																																																																																																																																																																																																	
Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V																																																																																																																																																																																																
Heating and Cooling Options	- Air cooling - Optional Heater Manager																																																																																																																																																																																																	
Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																																																																																																																																																																																																	
	Dell XR8610t and XR8620t Spec Sheet), 2																																																																																																																																																																																																	
	<table><tr><th>Rack Server</th><th>R750</th><th>R750xa</th><th>R650</th><th>R752s</th><th>R652s</th><th>R751s</th><th>R651s</th><th>R750xs</th><th>R650xs</th><th>R450</th><th>R550</th><th>XR11</th><th>XR12</th><th>R350</th><th>R250</th></tr><tr><td>Key attributes</td><td>Outstanding performance for the most demanding workloads</td><td>Highly intensive GPU workloads</td><td>High scalability, optimized workload performance</td><td>Powerful performance and flexibility</td><td>Dense virtualization</td><td>Powerful performance and scalability</td><td>High density compute</td><td>Purpose-built 2U server for growing scale-out solutions</td><td>Purpose-built, full performance 1U server for dense, fast growing scale-out solutions</td><td>Value and density-focused, built for general purpose IT</td><td>Versatile, value-optimized, virtualization-ready, built for general purpose IT</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Powerful performance in 1U server for productivity and data intensive applications</td><td>Powerful compute for common business applications and streamlines productivity</td></tr><tr><td>Target workloads</td><td>Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments</td><td>AI, ML, or DL training or inferencing, HPC, and visualization environments</td><td>Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments</td><td>All flash SDS, VDI, and data analytics</td><td>HPC, Dense VDI, and Virtualization</td><td>SDS, Virtualization, and Data Analytics</td><td>Virtualization, HCI and NFV</td><td>Virtualization, medium VM density or VDI, and scale-out database workloads</td><td>Virtualization, cloud, scale-out database and high-performance compute workloads</td><td>Small IT infrastructure, light VM, small business specific workloads</td><td>Small IT infrastructure, light VM density, small business specific workloads</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads</td></tr><tr><td>Type of processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor</td><td colspan="2">2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor</td><td colspan="2">1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor</td><td colspan="2">1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor</td><td colspan="2">1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores</td><td></td></tr><tr><td>Memory (DDR4 DIMM slots & max)</td><td colspan="2">32 (8 TB)</td><td colspan="2">32 (4 TB)</td><td colspan="2">16 (2 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2">8 (1 TB)</td><td colspan="2">4 (128 GB)</td><td></td></tr><tr><td>Disk drives up to:</td><td>8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)</td><td>6 x 2.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)</td><td>4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)</td><td>26 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>12 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>24 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear) 8 x 3.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>16 x 2.5" 8 x 3.5"</td><td>4 x 2.5" 6 x 2.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)</td><td>4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)</td></tr><tr><td>NVMe drives up to:</td><td>24</td><td>8</td><td>12</td><td>24</td><td>12</td><td>24</td><td>10</td><td>8</td><td>10</td><td>N/A</td><td>4</td><td>6</td><td>N/A</td><td>N/A</td><td>N/A</td></tr><tr><td>Gen4 PCIe slots up to:</td><td>8</td><td>8</td><td>3</td><td>8</td><td>3</td><td>2</td><td>1</td><td>5</td><td>3</td><td>2</td><td>3</td><td>3</td><td>5</td><td>3</td><td>2</td></tr><tr><td>Gen3 PCIe slots up to:</td><td colspan="2">N/A</td><td colspan="2">N/A</td><td colspan="2">2</td><td colspan="2">1</td><td colspan="2">1</td><td colspan="2">N/A</td><td colspan="2">1</td><td colspan="2">N/A</td></tr><tr><td>Accelerator support up to:</td><td>2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW</td><td>4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW</td><td>3 x 75 W SW</td><td>3 x 300 W DW or 6 x 75 W SW</td><td>3 x SW or 6 x 75 W SW</td><td>4 x SW; 1 x DW; 1 x FPGA</td><td>1 x SW</td><td colspan="2">N/A</td><td colspan="2">N/A</td><td>2 x 75 W SW</td><td>2 x 75 W or 150 W SW or 2 x 300 W DW</td><td colspan="2">N/A</td></tr><tr><td>Rack height (U)</td><td>2</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td></tr><tr><td>Integrated security</td><td colspan="15">TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China Nation2, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks</td></tr></table>	Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250	Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity	Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and visualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads	Type of processor	2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores			Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)		16 (1 TB)		8 (1 TB)		4 (128 GB)			Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 12 x 3.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear) 8 x 3.5"	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A	4	6	N/A	N/A	N/A	Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2	Gen3 PCIe slots up to:	N/A		N/A		2		1		1		N/A		1		N/A		Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW or 6 x 75 W SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		N/A		2 x 75 W SW	2 x 75 W or 150 W SW or 2 x 300 W DW	N/A		Rack height (U)	2	2	1	2	1	2	1	2	1	1	2	1	2	1	1	Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China Nation2, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks														
Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250																																																																																																																																																																																			
Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity																																																																																																																																																																																			
Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and visualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads																																																																																																																																																																																			
Type of processor	2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores																																																																																																																																																																																					
Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)		16 (1 TB)		8 (1 TB)		4 (128 GB)																																																																																																																																																																																					
Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 12 x 3.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear) 8 x 3.5"	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)																																																																																																																																																																																			
NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A	4	6	N/A	N/A	N/A																																																																																																																																																																																			
Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2																																																																																																																																																																																			
Gen3 PCIe slots up to:	N/A		N/A		2		1		1		N/A		1		N/A																																																																																																																																																																																			
Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW or 6 x 75 W SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		N/A		2 x 75 W SW	2 x 75 W or 150 W SW or 2 x 300 W DW	N/A																																																																																																																																																																																				
Rack height (U)	2	2	1	2	1	2	1	2	1	1	2	1	2	1	1																																																																																																																																																																																			
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China Nation2, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks																																																																																																																																																																																																	
	PowerEdge Reference Guide, 1.																																																																																																																																																																																																	

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))												
	Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240	
	Key attributes	Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple	
	Target workloads	In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads	
	Type of processor	4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors						1 x Intel Xeon E-2200, Intel Core i3®, Intel Pentium®, or Intel Celeron® processor		
	Memory (DDR4 DIMM slots & max)	48 (15.36 TB)			24 (7.68 TB)			16 (1 TB)	24 (7.68 TB)		16 (1 TB)	4 (64 GB)	
	Disk drives up to:	24 x 2.5"	32 x 2.5"	26 x 2.5"	32 x 2.5" 18 x 3.5"	16 x 2.5" 8 x 3.5"	26 x 3.5" 16 x 3.5" + 10 x 2.5"	12 x 2.5" 4 x 3.5"	14 x 3.5"	10 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 2.5" 4 x 3.5"	
	NVMe drives up to:	12	4	24			N/A	10		N/A	4	N/A	
	Gen4 PCIe slots up to:	N/A											
	Gen3 PCIe slots up to:	13	12	6	8		5	3	5	2		2	
	Accelerator support up to:	N/A	4 x DW GPUs or 4 x DW or 8 x SW FPGAs	2 x DW GPUs or 2 x SW or DW FPGAs	3 x DW or 6 x SW GPUs or 3 x DW or 4 x SW FPGAs		N/A	1 x SW GPU or 1 x SW FPGA	N/A				
	Rack height (U)	3	4	2				1	2	1	1		
	Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks											
	PowerEdge Reference Guide, 2.												
1[e] a declination index value calculation unit that calculates an index value indicating a degree of declination of the component temperature of the electronic component based on a detection result of the first temperature sensor; and	The Accused Dell '320 Products includes a declination index value calculation unit that calculates an index value indicating a degree of declination of the component temperature of the electronic component based on a detection result of the first temperature sensor. <i>See supra</i>, claim 1[d]. <i>See, e.g.,</i> Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39												


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.   Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Results Perf vs Temp vs Fan speed <table><tr><td>Perf (Gb/s)</td><td>23.1</td><td>23.1</td><td>23.1</td><td>23.1</td><td>22.8</td><td>22.7</td><td>22.4</td></tr><tr><td>Temp (Celsius)</td><td>0</td><td>10</td><td>20</td><td>30</td><td>40</td><td>50</td><td>55</td></tr><tr><td>Fan speed (% of total)</td><td>5</td><td>5</td><td>14</td><td>42</td><td>64</td><td>86</td><td>100</td></tr></table> Perf (Gb/s) Temp (Celsius) Fan speed (% of total) Figure 3. Constant networking performance with varying temperature and fan speed Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C. Dell XR8000 Thermal Design.	Perf (Gb/s)	23.1	23.1	23.1	23.1	22.8	22.7	22.4	Temp (Celsius)	0	10	20	30	40	50	55	Fan speed (% of total)	5	5	14	42	64	86	100
Perf (Gb/s)	23.1	23.1	23.1	23.1	22.8	22.7	22.4																		
Temp (Celsius)	0	10	20	30	40	50	55																		
Fan speed (% of total)	5	5	14	42	64	86	100																		

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	 Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with thermal control mechanisms. It is divided into a 'Cold Aisle' on the left and a 'Hot Aisle' on the right. In the Cold Aisle, multiple servers (represented by dots) report airflow (CFM#1, CFM#2, ..., CFM#N) to a central point labeled '1 Airflow reporting (per server)'. This leads to 'Inlet Airflow (Sum of CFMs)' which enters the server rack. Inside the rack, a 'Custom Delta-T' is measured (labeled '2'). The air then exits into the Hot Aisle, where 'Exhaust Temp Control (per server)' is implemented (labeled '3'). Below the rack, 'PCIe Thermal Control' is shown, including 'Inlet PCIe Temp' (labeled '4') and 'PCIe Airflow (LFM)' (labeled '5'). A small icon of a PCIe card is shown. A legend at the bottom left states 'CFM = Airflow in Cubic Feet/min'. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
1[f] a control unit that controls the number of rotations of the cooling fan based on the index value calculated by the declination index value calculation unit and the temperature of the intake air by the second temperature sensor.	The Accused Dell '320 Products includes a control unit that controls the number of rotations of the cooling fan based on the index value calculated by the declination index value calculation unit and the temperature of the intake air by the second temperature sensor. <i>See supra</i>, claim 1[e]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	 Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. A circled '1' is next to the text 'Airflow reporting (per server)'. Server Rack: A central server rack is shown with a circled '2' at its base labeled 'Custom Delta-T'. Hot Aisle: On the right, a red arrow points away from the rack, labeled 'Exhaust Temp Control (per server)' with a circled '3'. PCIe Thermal Control: At the bottom, a bracket groups two elements: - A circled '4' next to 'Inlet PCIe Temp' with a small server component icon. - A circled '5' next to 'PCIe Airflow (LFM)' with a large blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
2. The ICT equipment according to claim 1, wherein the index value is a declined temperature per unit time of the component temperature of the electronic component.	In the Accused Dell '320 Products, the index value is a declined temperature per unit time of the component temperature of the electronic component. <i>See supra</i>, claim 1. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior. Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	 Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket labeled '1 Airflow reporting (per server)' groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. Server Rack: A central server rack is shown with a blue arrow labeled '2 Custom Delta-T' pointing into it. Hot Aisle: On the right, a red arrow labeled '3 Exhaust Temp Control (per server)' points away from the rack. PCIe Thermal Control: Below the rack, a bracket groups two elements: 4 Inlet PCIe Temp: Represented by a green fan icon. 5 PCIe Airflow (LFM): Represented by a blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
3. The ICT equipment according to claim 1, wherein the index value is an average value of declination rates of the component temperature of the electronic component in a plurality of measurement periods.	In the Accused Dell '320 Products, the index value is an average value of declination rates of the component temperature of the electronic component in a plurality of measurement periods. <i>See supra</i>, claim 1. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	 Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with thermal control mechanisms. It is divided into a 'Cold Aisle' on the left and a 'Hot Aisle' on the right. In the Cold Aisle, a server rack is shown with multiple fans labeled CFM#1, CFM#2, and CFM#N. A bracket labeled '1 Airflow reporting (per server)' encompasses these fans. A large blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points from the Cold Aisle into the server rack. A bracket labeled '2 Custom Delta-T' points to the server rack. In the Hot Aisle, a red arrow points away from the server rack, labeled '3 Exhaust Temp Control (per server)'. Below the server rack, a bracket labeled 'PCIe Thermal Control' encompasses two elements: '4 Inlet PCIe Temp' (with a fan icon) and '5 PCIe Airflow (LFM)' (with a blue arrow icon). A legend at the bottom left states 'CFM = Airflow in Cubic Feet/min'. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
4. The ICT equipment according to claim 1, wherein the index value is a difference between a highest value of the component temperature of the electronic component in a past certain period and a current component temperature.	In the Accused Dell '320 Products, the index value is a difference between a highest value of the component temperature of the electronic component in a past certain period and a current component temperature. <i>See supra</i>, claim 1. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed The graph shows two data series: 'Fan speed (% of total)' (blue line) and 'System Power consumption (W)' (orange line). The x-axis represents Temperature in degrees Celsius, ranging from 0 to 60. The y-axis represents Power Consumption in Watts, ranging from 0 to 1200. The fan speed increases linearly from 5% at 0°C to 100% at 55°C. The system power consumption increases exponentially from 216W at 0°C to 1014W at 55°C. <table><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with thermal control mechanisms. It is divided into a 'Cold Aisle' on the left and a 'Hot Aisle' on the right. In the Cold Aisle, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points towards the server rack. Above the rack, three dots represent 'CFM#N', 'CFM#2', and 'CFM#1'. A bracket labeled '1 Airflow reporting (per server)' encompasses these. A red arrow labeled '2 Custom Delta-T' points from the inlet airflow towards the rack. In the Hot Aisle, a red arrow labeled '3 Exhaust Temp Control (per server)' points away from the rack. Below the rack, a bracket labeled 'PCIe Thermal Control' encompasses two elements: '4 Inlet PCIe Temp' (with a fan icon) and '5 PCIe Airflow (LFM)' (with a blue arrow icon). A legend at the bottom left states 'CFM = Airflow in Cubic Feet/min'. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also shows 'Exhaust Temperature Limits' with a 'Set Maximum Exhaust Temperature Limit' (Default: 70°C / 158.0°F) and a 'Target Exhaust Temperature Limit' (70°C / 158.0°F). At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' set to 'Default' (Range 25 - 100). Buttons for 'Apply' and 'Reboot' are visible at the bottom right of the configuration panel. Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced. In addition and/or in the alternative to literally infringing this claim element, the Accused Dell '320 Products also infringe under the doctrine of equivalents. The Accused Dell '320 Products are at least equivalent to and insubstantially different from the claimed limitation where the index value is a difference between a highest value of the component temperature of the electronic component in a past certain period and a current component temperature to perform substantially the same function (e.g., identifying whether and to what the extent the current component needs to be cooled) in substantially the same way (e.g., comparing the highest component temperature value from the past to the current component temperature), to achieve substantially the same result (e.g., determining how fast to rotate the cooling fan to cool the component).

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
5. The ICT equipment according to claim 1, wherein the control unit controls the number of rotations of the cooling fan to be lower as the degree of declination of the component temperature of the electronic component, indicated by the index value calculated by the declination index value calculation unit, is larger.	In the Accused Dell '320 Products, the control unit controls the number of rotations of the cooling fan to be lower as the degree of declination of the component temperature of the electronic component, indicated by the index value calculated by the declination index value calculation unit, is larger. <i>See supra</i>, claim 1. <i>See, e.g.,</i> Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	 Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. A circled '1' is next to the text 'Airflow reporting (per server)'. Server Rack: A central server rack is shown with a circled '2' at its base labeled 'Custom Delta-T'. Hot Aisle: On the right, a red arrow points away from the rack, labeled 'Exhaust Temp Control (per server)' with a circled '3'. PCIe Thermal Control: At the bottom, a bracket groups two elements: - A circled '4' next to 'Inlet PCIe Temp' with a small server component icon. - A circled '5' next to 'PCIe Airflow (LFM)' with a blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also shows 'Exhaust Temperature Limits' with a 'Set Maximum Exhaust Temperature Limit' dropdown set to 'Default: 70 °C (158.0 °F)' and a 'Target Exhaust Temperature Limit' set to '70°C (158.0°F)'. At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' dropdown set to 'Default' with a range of 25 - 100. 'Apply' and 'Reboot' buttons are at the bottom right. Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
6. The ICT equipment according to claim 1, wherein if the index value is not less than a predetermined threshold, the control unit stops the cooling fan.	In the Accused Dell '320 Products, if the index value is not less than a predetermined threshold, the control unit stops the cooling fan. <i>See supra</i>, claim 1. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.   Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed <table border="1"><thead><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr></thead><tbody><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></tbody></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. A circled '1' is next to the text 'Airflow reporting (per server)'. Server Rack: A central server rack is shown with a circled '2' at its base labeled 'Custom Delta-T'. Hot Aisle: On the right, a red arrow points away from the rack. A circled '3' is next to the text 'Exhaust Temp Control (per server)'. PCIe Thermal Control: At the bottom, a bracket groups two elements: - A circled '4' next to 'Inlet PCIe Temp' with a small server component icon. - A circled '5' next to 'PCIe Airflow (LFM)' with a blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also shows 'Exhaust Temperature Limits' with a 'Set Maximum Exhaust Temperature Limit' (Default: 70°C / 158.0°F) and a 'Target Exhaust Temperature Limit' (70°C / 158.0°F). At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' set to 'Default' (Range 25 - 100). Buttons for 'Apply' and 'Reboot' are visible at the bottom right of the configuration panel. Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced. In addition and/or in the alternative to literally infringing this claim element, the Accused Dell '320 Products also infringe under the doctrine of equivalents. The Accused Dell '320 Products are at least equivalent to and insubstantially different from the claimed limitation where if the index value is not less than a predetermined threshold, the control unit stops the cooling fan to perform substantially the same function (e.g., identifying whether and to what extent the current component needs to be cooled) in substantially the same way (e.g., comparing the rate that the current component temperature value is changing to the predetermined threshold), to achieve substantially the same result (e.g., controlling the rotation of the cooling fan to cool the component).

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																											
7[pre] A control device for ICT (Information and Communication technology) equipment having a cooling fan, electronic component, a first temperature sensor for detecting a temperature of the electronic component, and a second temperature sensor that detects a temperature of an intake air, the control device comprising:	The Accused Dell '320 Products include a control device for ICT (Information and Communication technology) equipment having a cooling fan, electronic component, a first temperature sensor for detecting a temperature of the electronic component, and a second temperature sensor that detects a temperature of an intake air. <i>See supra</i>, claims 1[pre], 1[a], 1[b], and 1[c]. <i>See, e.g.,</i> <table><tr><th>Feature</th><th colspan="2">Technical Specifications for XR8000r XR8610t and, XR8620t</th></tr><tr><td>Processor</td><td colspan="2">1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores</td></tr><tr><td>Memory</td><td colspan="2"> • 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s • Supports registered ECC DDR5 DIMMs only </td></tr><tr><td>Storage</td><td colspan="2"> • 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t • 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t • BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t • Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t • Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t </td></tr><tr><td>Chassis</td><td colspan="2"> • XR8000r </td></tr><tr><td>Sleds</td><td colspan="2"> • XR8610t and XR8620t </td></tr><tr><td>Hot swap Redundant Power Supplies</td><td> • 1800 W Titanium 200—240 VAC or 240 HVDC • 1400 W Platinum 100—240 VAC or 240 HVDC </td><td> • 1400W DC/ -48 — (-60) V • 1100W DC/ -48 — (-60) V • 800W DC/ -48 — (-60) V </td></tr><tr><td>Heating and Cooling Options</td><td colspan="2"> • Air cooling • Optional Heater Manager </td></tr><tr><td>Fans</td><td colspan="2">Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR8620t sled</td></tr></table> Dell XR8610t and XR8620t Spec Sheet), , 2.	Feature	Technical Specifications for XR8000r XR8610t and, XR8620t		Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores		Memory	• 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s • Supports registered ECC DDR5 DIMMs only		Storage	• 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t • 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t • BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t • Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t • Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t		Chassis	• XR8000r		Sleds	• XR8610t and XR8620t		Hot swap Redundant Power Supplies	• 1800 W Titanium 200—240 VAC or 240 HVDC • 1400 W Platinum 100—240 VAC or 240 HVDC	• 1400W DC/ -48 — (-60) V • 1100W DC/ -48 — (-60) V • 800W DC/ -48 — (-60) V	Heating and Cooling Options	• Air cooling • Optional Heater Manager		Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR8620t sled	
Feature	Technical Specifications for XR8000r XR8610t and, XR8620t																											
Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores																											
Memory	• 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s • Supports registered ECC DDR5 DIMMs only																											
Storage	• 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t • 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t • BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t • Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t • Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t																											
Chassis	• XR8000r																											
Sleds	• XR8610t and XR8620t																											
Hot swap Redundant Power Supplies	• 1800 W Titanium 200—240 VAC or 240 HVDC • 1400 W Platinum 100—240 VAC or 240 HVDC	• 1400W DC/ -48 — (-60) V • 1100W DC/ -48 — (-60) V • 800W DC/ -48 — (-60) V																										
Heating and Cooling Options	• Air cooling • Optional Heater Manager																											
Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR8620t sled																											

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))															
	Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250
Key attributes		Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity
Target workloads		Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and virtualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POS/IOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads
Type of processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor			2 x 2 nd or 3 rd Generation AMD EPYC™ processors; up to 64 cores per processor			1 x 2 nd or 3 rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3 rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3 rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores
Memory (DDR4 DIMM slots & max)		32 (8 TB)		32 (4 TB)			16 (2 TB)		16 (1 TB)				8 (1 TB)		4 (128 GB)	
Disk drives up to:		8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	
NVMe drives up to:		24	8	12	24	12	24	10	8	10		N/A	4	6		N/A
Gen4 PCIe slots up to:		8	8	3	8	3	2	1	5	3	2	3	3	5	3	2
Gen3 PCIe slots up to:				N/A			2	1	1		N/A	1			N/A	
Accelerator support up to:		2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW			N/A		2 x 75 W SW	2 x 75 W or 150 W SW 2 x 300 W DW		N/A
Rack height (U)		2	2	1	2	1	2	1	2	1	1	2	1	2	1	1
Integrated security (or Datacenter), and System Erase on all racks		TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC9 Enterprise or Datacenter), and System Erase on all racks														
PowerEdge Reference Guide, 1.																
	Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240				
Key attributes		Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple				
Target workloads		In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads				
Type of processor		4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors						1 x Intel Xeon E-2200, Intel Core i3®, Intel Pentium®, or Intel Celeron® processor					
Memory (DDR4 DIMM slots & max)		48 (15.36 TB)			24 (7.68 TB)		16 (1 TB)	24 (7.68 TB)	16 (1 TB)		4 (64 GB)					
Disk drives up to:		24 x 2.5"	32 x 2.5"	26 x 2.5"	32 x 2.5" 16 x 3.5"	16 x 2.5" 8 x 3.5"	26 x 3.5" 16 x 3.5" + 10 x 2.5"	12 x 2.5" 4 x 3.5"	14 x 3.5"	10 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 2.5" 4 x 3.5"				
NVMe drives up to:		12	4	24			N/A	10	N/A	4		N/A				
Gen4 PCIe slots up to:		13	12	6	8		5	3	5	2	2					
Gen3 PCIe slots up to:																
Accelerator support up to:		N/A	4 x DW GPUs or 4 x DW or 8 x SW FPGAs	2 x DW GPUs or 2 x SW or DW FPGAs	3 x DW or 6 x SW GPUs or 3 x DW or 4 x SW FPGAs		N/A	1 x SW GPU or 1 x SW FPGA			N/A					
Rack height (U)		3	4	2			1	2	1	1						
Integrated security (or Datacenter), and System Erase on all racks		TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC9 Enterprise or Datacenter), and System Erase on all racks														
PowerEdge Reference Guide, 2.																

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																
	Thoughtful Thermal Design Changes High Performance Fans – New higher performance fans have been added to increase the amount of cool air pushed through the system, targeted at higher power semiconductors (CPUs, GPUs and NVMe drives). A 3-tier approach to fan performance is employed to address the trend of increasing total power. This includes standard fans and high performance silver and gold fans, which increase the airflow capability versus the previous-generation fans (see Figure 1). 1U System Airflow vs. Fan Type <table border="1"><thead><tr><th>Fan Type</th><th>STD</th><th>High Perf Silver</th><th>High Perf Gold</th></tr></thead><tbody><tr><td>R6515 10x2.5"</td><td>~82</td><td>~95</td><td>~115</td></tr><tr><td>R650 10x2.5"</td><td>~80</td><td>~90</td><td>~110</td></tr><tr><td>R650s 10x2.5"</td><td>~72</td><td>~85</td><td>~100</td></tr></tbody></table> Figure 1 – Graph illustrating the increased fan airflow (in CFM) for high performance silver and gold fans compared to baseline fans Next Gen PowerEdge Thermal Design, 1. Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.	Fan Type	STD	High Perf Silver	High Perf Gold	R6515 10x2.5"	~82	~95	~115	R650 10x2.5"	~80	~90	~110	R650s 10x2.5"	~72	~85	~100
Fan Type	STD	High Perf Silver	High Perf Gold														
R6515 10x2.5"	~82	~95	~115														
R650 10x2.5"	~80	~90	~110														
R650s 10x2.5"	~72	~85	~100														

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior. Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed <table><thead><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr></thead><tbody><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></tbody></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. A circled '1' is next to the text 'Airflow reporting (per server)'. Server Rack: A central server rack is shown with a circled '2' at its base labeled 'Custom Delta-T'. Hot Aisle: On the right, a red arrow points away from the rack, labeled 'Exhaust Temp Control (per server)' with a circled '3'. PCIe Thermal Control: At the bottom, a bracket groups two elements: - A circled '4' next to 'Inlet PCIe Temp' with a small server component icon. - A circled '5' next to 'PCIe Airflow (LFM)' with a blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also shows 'Exhaust Temperature Limits' with a 'Set Maximum Exhaust Temperature Limit' dropdown set to 'Default, 70 °C (158.0 °F)' and a 'Target Exhaust Temperature Limit' set to '70°C (158.0°F)'. At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' dropdown set to 'Default' with a range of 25 - 100. 'Apply' and 'Reboot' buttons are at the bottom right. Thermal Manage.
7[a] a microprocessor including:	The Accused Dell '320 Products include a microprocessor. <i>See supra</i>, claims 1[d] and 7[pre]. <i>See, e.g.,</i>


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																																																																																																																																																																																																			
	<table><tr><th>Feature</th><th colspan="2">Technical Specifications for XR8000r XR8610t and, XR8620t</th></tr><tr><td>Processor</td><td colspan="2">1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores</td></tr><tr><td>Memory</td><td colspan="2"> 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only </td></tr><tr><td>Storage</td><td colspan="2"> 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t </td></tr><tr><td>Chassis</td><td colspan="2"> - XR8000r </td></tr><tr><td>Sleds</td><td colspan="2"> - XR8610t and XR8620t </td></tr><tr><td>Hot swap Redundant Power Supplies</td><td> 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC </td><td> 1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V </td></tr><tr><td>Heating and Cooling Options</td><td colspan="2"> - Air cooling - Optional Heater Manager </td></tr><tr><td>Fans</td><td colspan="2">Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled</td></tr></table>	Feature	Technical Specifications for XR8000r XR8610t and, XR8620t		Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores		Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only		Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t		Chassis	XR8000r		Sleds	XR8610t and XR8620t		Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V	Heating and Cooling Options	- Air cooling - Optional Heater Manager		Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																																																																																																																																																																									
Feature	Technical Specifications for XR8000r XR8610t and, XR8620t																																																																																																																																																																																																			
Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores																																																																																																																																																																																																			
Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only																																																																																																																																																																																																			
Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t																																																																																																																																																																																																			
Chassis	XR8000r																																																																																																																																																																																																			
Sleds	XR8610t and XR8620t																																																																																																																																																																																																			
Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V																																																																																																																																																																																																		
Heating and Cooling Options	- Air cooling - Optional Heater Manager																																																																																																																																																																																																			
Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																																																																																																																																																																																																			
	Dell XR8610t and XR8620t Spec Sheet), 2.																																																																																																																																																																																																			
	<table><tr><th>Rack Server</th><th>R750</th><th>R750xa</th><th>R650</th><th>R752s</th><th>R652s</th><th>R751s</th><th>R651s</th><th>R750xs</th><th>R650xs</th><th>R450</th><th>R550</th><th>XR11</th><th>XR12</th><th>R350</th><th>R250</th></tr><tr><td>Key attributes</td><td>Outstanding performance for the most demanding workloads</td><td>Highly intensive GPU workloads</td><td>High scalability, optimized workload performance</td><td>Powerful performance and flexibility</td><td>Dense virtualization</td><td>Powerful performance and scalability</td><td>High density compute</td><td>Purpose-built 2U server for growing scale-out solutions</td><td>Purpose-built, full performance 1U server for dense, fast growing scale-out solutions</td><td>Value and density-focused, built for general purpose IT</td><td>Versatile, value-optimized, virtualization-ready, built for general purpose IT</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Powerful performance in 1U server for productivity and data intensive applications</td><td>Powerful compute for common business applications and streamlines productivity</td></tr><tr><td>Target workloads</td><td>Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments</td><td>AI, ML, or DL training or inferencing, HPC, and virtualization environments</td><td>Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments</td><td>All flash SDS, VDI, and data analytics</td><td>HPC, Dense VDI, and Virtualization</td><td>SDS, Virtualization, and Data Analytics</td><td>Virtualization, HCI and NFV</td><td>Virtualization, medium VM density or VDI, and scale-out database workloads</td><td>Virtualization, cloud, scale-out database and high-performance compute workloads</td><td>Small IT infrastructure, light VM, small business specific workloads</td><td>Small IT infrastructure, light VM density, small business specific workloads</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads</td></tr><tr><td>Type of processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor</td><td colspan="2">2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor</td><td colspan="2">1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor</td><td colspan="2">1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor</td><td colspan="2">1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores</td><td colspan="2">1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores</td></tr><tr><td>Memory (DDR4 DIMM slots & max)</td><td colspan="2">32 (8 TB)</td><td colspan="2">32 (4 TB)</td><td colspan="2">16 (2 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2">8 (1 TB)</td><td colspan="2">4 (128 GB)</td><td colspan="2">4 (128 GB)</td></tr><tr><td>Disk drives up to:</td><td>8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)</td><td>6 x 2.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)</td><td>4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)</td><td>26 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>12 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>24 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear) 8 x 3.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>16 x 2.5" 8 x 3.5"</td><td>4 x 2.5" 6 x 2.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)</td><td>4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)</td></tr><tr><td>NVMe drives up to:</td><td>24</td><td>8</td><td>12</td><td>24</td><td>12</td><td>24</td><td>10</td><td>8</td><td>10</td><td>N/A</td><td>4</td><td>6</td><td>N/A</td><td>N/A</td><td>N/A</td></tr><tr><td>Gen4 PCIe slots up to:</td><td>8</td><td>8</td><td>3</td><td>8</td><td>3</td><td>2</td><td>1</td><td>5</td><td>3</td><td>2</td><td>3</td><td>3</td><td>5</td><td>3</td><td>2</td></tr><tr><td>Gen3 PCIe slots up to:</td><td colspan="2">N/A</td><td colspan="2">N/A</td><td colspan="2">2</td><td colspan="2">1</td><td colspan="2">1</td><td colspan="2">N/A</td><td colspan="2">1</td><td colspan="2">N/A</td></tr><tr><td>Accelerator support up to:</td><td>2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW</td><td>4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW</td><td>3 x 75 W SW</td><td>3 x 300 W DW or 6 x 75 W SW</td><td>3 x SW</td><td>4 x SW; 1 x DW; 1 x FPGA</td><td>1 x SW</td><td colspan="2">N/A</td><td colspan="2">2 x 75 W SW</td><td>2 x 75 W or 150 W SW</td><td>2 x 300 W DW</td><td colspan="2">N/A</td></tr><tr><td>Rack height (U)</td><td>2</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td></tr><tr><td>Integrated security</td><td colspan="15">TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks</td></tr></table>	Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250	Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity	Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and virtualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads	Type of processor	2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores		Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)		16 (1 TB)		8 (1 TB)		4 (128 GB)		4 (128 GB)		Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 12 x 3.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear) 8 x 3.5"	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A	4	6	N/A	N/A	N/A	Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2	Gen3 PCIe slots up to:	N/A		N/A		2		1		1		N/A		1		N/A		Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		2 x 75 W SW		2 x 75 W or 150 W SW	2 x 300 W DW	N/A		Rack height (U)	2	2	1	2	1	2	1	2	1	1	2	1	2	1	1	Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks														
Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250																																																																																																																																																																																					
Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity																																																																																																																																																																																					
Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and virtualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads																																																																																																																																																																																					
Type of processor	2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores																																																																																																																																																																																					
Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)		16 (1 TB)		8 (1 TB)		4 (128 GB)		4 (128 GB)																																																																																																																																																																																					
Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 12 x 3.5" 4 x 3.5"	24 x 2.5" 12 x 3.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear) 8 x 3.5"	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 3.5"	4 x 2.5" 6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)																																																																																																																																																																																					
NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A	4	6	N/A	N/A	N/A																																																																																																																																																																																					
Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2																																																																																																																																																																																					
Gen3 PCIe slots up to:	N/A		N/A		2		1		1		N/A		1		N/A																																																																																																																																																																																					
Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		2 x 75 W SW		2 x 75 W or 150 W SW	2 x 300 W DW	N/A																																																																																																																																																																																						
Rack height (U)	2	2	1	2	1	2	1	2	1	1	2	1	2	1	1																																																																																																																																																																																					
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks																																																																																																																																																																																																			
	PowerEdge Reference Guide, 1.																																																																																																																																																																																																			

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																																																																																																																																																	
	<table><tr><th>Rack Server</th><th>R940</th><th>R940xa</th><th>R840</th><th>R740xd</th><th>R740</th><th>R740xd2</th><th>R640</th><th>R540</th><th>R440</th><th>R340</th><th>R240</th></tr><tr><td>Key attributes</td><td>Powerful performance</td><td>Extreme acceleration</td><td>Turbocharge data analytics</td><td>Scalable storage performance</td><td>Optimal application performance</td><td>Enterprise content server</td><td>Performance and density</td><td>Balanced and adaptable</td><td>Scale-out computing</td><td>Accelerate business growth</td><td>Compute made simple</td></tr><tr><td>Target workloads</td><td>In-memory databases</td><td>GPU database acceleration and machine learning</td><td>Data-intensive workloads, HFT, and dense virtualization</td><td>SDS, service providers, and big data servers</td><td>VDI and cloud workloads</td><td>Media streaming and SDS</td><td>Dense scale-out data center computing and storage</td><td>Mail messaging and virtualization</td><td>HPC, web tech, and scale-out infrastructure</td><td>ROBO productivity and data-intensive applications</td><td>Small business and service provider workloads</td></tr><tr><td>Type of processor</td><td colspan="3">4 x 2nd Generation Intel® Xeon® Scalable processors</td><td colspan="3">2 x 2nd Generation Intel® Xeon® Scalable processors</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>Memory (DDR4 DIMM slots & max)</td><td>48 (15.36 TB)</td><td></td><td></td><td>24 (7.68 TB)</td><td></td><td>16 (1 TB)</td><td>24 (7.68 TB)</td><td>16 (1 TB)</td><td></td><td>4 (64 GB)</td><td></td></tr><tr><td>Disk drives up to:</td><td>24 x 2.5"</td><td>32 x 2.5"</td><td>26 x 2.5"</td><td>32 x 2.5" 16 x 3.5"</td><td>16 x 2.5" 8 x 3.5"</td><td>26 x 3.5" 16 x 3.5" + 10 x 2.5"</td><td>12 x 2.5" 4 x 3.5"</td><td>14 x 3.5"</td><td>10 x 2.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>4 x 2.5" 4 x 3.5"</td></tr><tr><td>NVMe drives up to:</td><td>12</td><td>4</td><td>24</td><td></td><td></td><td>N/A</td><td>10</td><td>N/A</td><td>4</td><td></td><td>N/A</td></tr><tr><td>Gen4 PCIe slots up to:</td><td>13</td><td>12</td><td>6</td><td>8</td><td></td><td>5</td><td>3</td><td>5</td><td>2</td><td>2</td><td></td></tr><tr><td>Gen3 PCIe slots up to:</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>Accelerator support up to:</td><td>N/A</td><td>4 x DW GPUs or 4 x DW or 8 x SW FPGAs</td><td>2 x DW GPUs or 2 x SW or DW FPGAs</td><td>3 x DW or 6 x SW GPUs or 3 x DW or 4 x SW FPGAs</td><td></td><td>N/A</td><td>1 x SW GPU or 1 x SW FPGA</td><td></td><td></td><td>N/A</td><td></td></tr><tr><td>Rack height (U)</td><td>3</td><td>4</td><td>2</td><td></td><td></td><td></td><td>1</td><td>2</td><td>1</td><td>1</td><td></td></tr><tr><td>Integrated security</td><td colspan="11">TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC9 Enterprise or Datacenter), and System Erase on all racks.</td></tr></table>	Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240	Key attributes	Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple	Target workloads	In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads	Type of processor	4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors								Memory (DDR4 DIMM slots & max)	48 (15.36 TB)			24 (7.68 TB)		16 (1 TB)	24 (7.68 TB)	16 (1 TB)		4 (64 GB)		Disk drives up to:	24 x 2.5"	32 x 2.5"	26 x 2.5"	32 x 2.5" 16 x 3.5"	16 x 2.5" 8 x 3.5"	26 x 3.5" 16 x 3.5" + 10 x 2.5"	12 x 2.5" 4 x 3.5"	14 x 3.5"	10 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 2.5" 4 x 3.5"	NVMe drives up to:	12	4	24			N/A	10	N/A	4		N/A	Gen4 PCIe slots up to:	13	12	6	8		5	3	5	2	2		Gen3 PCIe slots up to:												Accelerator support up to:	N/A	4 x DW GPUs or 4 x DW or 8 x SW FPGAs	2 x DW GPUs or 2 x SW or DW FPGAs	3 x DW or 6 x SW GPUs or 3 x DW or 4 x SW FPGAs		N/A	1 x SW GPU or 1 x SW FPGA			N/A		Rack height (U)	3	4	2				1	2	1	1		Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC9 Enterprise or Datacenter), and System Erase on all racks.											PowerEdge Reference Guide, 2.
Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240																																																																																																																																							
Key attributes	Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple																																																																																																																																							
Target workloads	In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads																																																																																																																																							
Type of processor	4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors																																																																																																																																														
Memory (DDR4 DIMM slots & max)	48 (15.36 TB)			24 (7.68 TB)		16 (1 TB)	24 (7.68 TB)	16 (1 TB)		4 (64 GB)																																																																																																																																								
Disk drives up to:	24 x 2.5"	32 x 2.5"	26 x 2.5"	32 x 2.5" 16 x 3.5"	16 x 2.5" 8 x 3.5"	26 x 3.5" 16 x 3.5" + 10 x 2.5"	12 x 2.5" 4 x 3.5"	14 x 3.5"	10 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 2.5" 4 x 3.5"																																																																																																																																							
NVMe drives up to:	12	4	24			N/A	10	N/A	4		N/A																																																																																																																																							
Gen4 PCIe slots up to:	13	12	6	8		5	3	5	2	2																																																																																																																																								
Gen3 PCIe slots up to:																																																																																																																																																		
Accelerator support up to:	N/A	4 x DW GPUs or 4 x DW or 8 x SW FPGAs	2 x DW GPUs or 2 x SW or DW FPGAs	3 x DW or 6 x SW GPUs or 3 x DW or 4 x SW FPGAs		N/A	1 x SW GPU or 1 x SW FPGA			N/A																																																																																																																																								
Rack height (U)	3	4	2				1	2	1	1																																																																																																																																								
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC9 Enterprise or Datacenter), and System Erase on all racks.																																																																																																																																																	
7[b] a declination index value calculation unit that calculates an index value indicating a degree of declination of a component temperature of the electronic component included in ICT equipment, based on a detection result of the first temperature sensor that detects the component temperature of the electronic component; and	The Accused Dell '320 Products includes a declination index value calculation unit that calculates an index value indicating a degree of declination of a component temperature of the electronic component included in ICT equipment, based on a detection result of the first temperature sensor that detects the component temperature of the electronic component. <i>See supra</i>, claims 1[e] and 7[a]. <i>See, e.g.,</i> Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39.																																																																																																																																																	

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed The graph shows two data series: Fan speed (% of total) represented by a blue line with circular markers, and System Power consumption (W) represented by an orange line with circular markers. The x-axis represents Temperature in degrees Celsius, ranging from 0 to 60. The y-axis represents Power Consumption in Watts, ranging from 0 to 1200. Both metrics increase as temperature rises, with power consumption showing a much steeper increase at higher temperatures. <table><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a vertical stack of servers is shown. A bracket labeled '1 Airflow reporting (per server)' points to the servers. To the right of the servers, a vertical list of 'CFM#N', 'CFM#2', and 'CFM#1' is shown. A large blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points upwards from the bottom left towards the servers. Server Rack: A central server rack is shown with a bracket labeled '2 Custom Delta-T' pointing to it. Hot Aisle: On the right, a red arrow points upwards, labeled '3 Exhaust Temp Control (per server)'. PCIe Thermal Control: At the bottom, a bracket labeled 'PCIe Thermal Control' points to two elements: '4 Inlet PCIe Temp' (with a fan icon) and '5 PCIe Airflow (LFM)' (with a horizontal blue arrow icon). CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
7[c] a control unit that controls the number of rotations of a cooling fan based on the index value calculated by the declination index value calculation unit and the temperature of the intake air by the second temperature sensor.	The Accused Dell '320 Products includes a control unit that controls the number of rotations of a cooling fan based on the index value calculated by the declination index value calculation unit and the temperature of the intake air by the second temperature sensor. <i>See supra</i>, claims 1[f] and 7[b]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed <table><thead><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr></thead><tbody><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></tbody></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and numbered callouts: 1 Airflow reporting (per server): A bracket groups three dots labeled CFM#N, CFM#2, and CFM#1. 2 Custom Delta-T: A double-headed arrow between the inlet and exhaust airflows. 3 Exhaust Temp Control (per server): A red arrow pointing up from the hot aisle. 4 Inlet PCIe Temp: A callout pointing to a server component. 5 PCIe Airflow (LFM): A callout pointing to a server component. Other labels include: Cold Aisle, Hot Aisle, Inlet Airflow (Sum of CFMs), and PCIe Thermal Control. A legend at the bottom states: CFM = Airflow in Cubic Feet/min. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it shows 'Current Temperature Readings' for 'System Inlet Temperature' (22°C / 71.6°F) and 'System Exhaust Temperature' (25°C / 77.0°F). It also has a checkbox for 'Set Maximum Exhaust Temperature Limit' (checked) with a dropdown set to 'Default: 70 °C (158.0 °F)'. Below that is a 'Target Exhaust Temperature Limit' set to 70°C (158.0°F). At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' dropdown set to 'Default' with a range of 25 - 100. 'Apply' and 'Reboot' buttons are at the bottom right. Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
8[pre] A temperature controlling method comprising:	The Accused Dell '320 Products perform a temperature controlling method. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																																																																																																																																																																																																	
	<table><tr><th>Feature</th><th colspan="2">Technical Specifications for XR8000r XR8610t and, XR8620t</th></tr><tr><td>Processor</td><td colspan="2">1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores</td></tr><tr><td>Memory</td><td colspan="2"> 8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only </td></tr><tr><td>Storage</td><td colspan="2"> 2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t </td></tr><tr><td>Chassis</td><td colspan="2"> - XR8000r </td></tr><tr><td>Sleds</td><td colspan="2"> - XR8610t and XR8620t </td></tr><tr><td>Hot swap Redundant Power Supplies</td><td> 1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC </td><td> 1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V </td></tr><tr><td>Heating and Cooling Options</td><td colspan="2"> - Air cooling - Optional Heater Manager </td></tr><tr><td>Fans</td><td colspan="2">Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled</td></tr></table>	Feature	Technical Specifications for XR8000r XR8610t and, XR8620t		Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores		Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only		Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t		Chassis	XR8000r		Sleds	XR8610t and XR8620t		Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V	Heating and Cooling Options	- Air cooling - Optional Heater Manager		Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																																																																																																																																																																							
Feature	Technical Specifications for XR8000r XR8610t and, XR8620t																																																																																																																																																																																																	
Processor	1 x 4th Generation Intel® Xeon® Scalable processor with optional vRAN boost up to 32 cores or 1 x 5th Generation Intel® Xeon® Scalable processor up to 16 cores																																																																																																																																																																																																	
Memory	8 x DDR5 DIMM slots, supports RDIMM total 512 GB with max speeds up to 5600 MT/s - Supports registered ECC DDR5 DIMMs only																																																																																																																																																																																																	
Storage	2 x M.2 2280 or 22110 direct connect NVMe drives with dual M.2 NVMe direct riser module (non-RAID) in XR8620t 2 x M.2 2280 SSDs on ROR-N1 (RAISER) with RAID 0/1 in XR8620t - BOSS N1 with 2 x 2280 NVMe SSDs in XR8620t and XR8610t - Single M.2 riser with 1 x 2280 M.2 SSD in XR8610t and XR8620t - Dual M.2 riser (non-RAID) with 2 x M.2 2280 or 22110 M.2 SSDs in XR8610t																																																																																																																																																																																																	
Chassis	XR8000r																																																																																																																																																																																																	
Sleds	XR8610t and XR8620t																																																																																																																																																																																																	
Hot swap Redundant Power Supplies	1800 W Titanium 200—240 VAC or 240 HVDC 1400 W Platinum 100—240 VAC or 240 HVDC	1400W DC/ -48 — (-60) V 1100W DC/ -48 — (-60) V 800W DC/ -48 — (-60) V																																																																																																																																																																																																
Heating and Cooling Options	- Air cooling - Optional Heater Manager																																																																																																																																																																																																	
Fans	Up to 4 standard cabled cooling fans with reverse air-flow installed in XR8610t sled and up to 8 Standard (STD) cabled cooling fans with reverse air-flow installed in XR86620t sled																																																																																																																																																																																																	
	Dell XR8610t and XR8620t Spec Sheet), 2.																																																																																																																																																																																																	
	<table><tr><th>Rack Server</th><th>R750</th><th>R750xa</th><th>R650</th><th>R752s</th><th>R652s</th><th>R751s</th><th>R651s</th><th>R750xs</th><th>R650xs</th><th>R450</th><th>R550</th><th>XR11</th><th>XR12</th><th>R350</th><th>R250</th></tr><tr><td>Key attributes</td><td>Outstanding performance for the most demanding workloads</td><td>Highly intensive GPU workloads</td><td>High scalability, optimized workload performance</td><td>Powerful performance and flexibility</td><td>Dense virtualization</td><td>Powerful performance and scalability</td><td>High density compute</td><td>Purpose-built 2U server for growing scale-out solutions</td><td>Purpose-built, full performance 1U server for dense, fast growing scale-out solutions</td><td>Value and density-focused, built for general purpose IT</td><td>Versatile, value-optimized, virtualization-ready, built for general purpose IT</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Edge-centric, short depth and rugged with reverse mounting options</td><td>Powerful performance in 1U server for productivity and data intensive applications</td><td>Powerful compute for common business applications and streamlines productivity</td></tr><tr><td>Target workloads</td><td>Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments</td><td>AI, ML, or DL training or inferencing, HPC, and visualization environments</td><td>Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments</td><td>All flash SDS, VDI, and data analytics</td><td>HPC, Dense VDI, and Virtualization</td><td>SDS, Virtualization, and Data Analytics</td><td>Virtualization, HCI and NFV</td><td>Virtualization, medium VM density or VDI, and scale-out database workloads</td><td>Virtualization, cloud, scale-out database and high-performance compute workloads</td><td>Small IT infrastructure, light VM, small business specific workloads</td><td>Small IT infrastructure, light VM density, small business specific workloads</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)</td><td>Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads</td><td>Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads</td></tr><tr><td>Type of processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor</td><td colspan="2">2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor</td><td colspan="2">1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor</td><td colspan="2">2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor</td><td colspan="2">1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor</td><td colspan="2">1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores</td><td></td></tr><tr><td>Memory (DDR4 DIMM slots & max)</td><td colspan="2">32 (8 TB)</td><td colspan="2">32 (4 TB)</td><td colspan="2">16 (2 TB)</td><td colspan="2">16 (1 TB)</td><td colspan="2"></td><td colspan="2">8 (1 TB)</td><td colspan="2">4 (128 GB)</td><td></td></tr><tr><td>Disk drives up to:</td><td>8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)</td><td>6 x 2.5" 8 x 2.5"</td><td>4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)</td><td>26 x 2.5" 12 x 3.5" 4 x 3.5"</td><td>12 x 2.5" 4 x 3.5"</td><td>24 x 2.5" 12 x 3.5"</td><td>8 x 2.5" 4 x 3.5"</td><td>8 x 2.5" 4 x 3.5" 24 x 2.5" 10 x 2.5" 12 x 3.5" 8 x 3.5"</td><td>4 x 3.5" 8 x 2.5" 4 x 3.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>16 x 2.5" 8 x 3.5"</td><td>4 x 2.5"</td><td>6 x 2.5"</td><td>4 x 3.5" 8 x 2.5"</td><td>4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)</td></tr><tr><td>NVMe drives up to:</td><td>24</td><td>8</td><td>12</td><td>24</td><td>12</td><td>24</td><td>10</td><td>8</td><td>10</td><td colspan="2">N/A</td><td>4</td><td>6</td><td colspan="2">N/A</td></tr><tr><td>Gen4 PCIe slots up to:</td><td>8</td><td>8</td><td>3</td><td>8</td><td>3</td><td>2</td><td>1</td><td>5</td><td>3</td><td>2</td><td>3</td><td>3</td><td>5</td><td>3</td><td>2</td></tr><tr><td>Gen3 PCIe slots up to:</td><td colspan="2"></td><td colspan="2">N/A</td><td colspan="2">2</td><td colspan="2">1</td><td colspan="2">1</td><td colspan="2">N/A</td><td colspan="2">1</td><td colspan="2">N/A</td></tr><tr><td>Accelerator support up to:</td><td>2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW</td><td>4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW</td><td>3 x 75 W SW</td><td>3 x 300 W DW or 6 x 75 W SW</td><td>3 x SW</td><td>4 x SW; 1 x DW; 1 x FPGA</td><td>1 x SW</td><td colspan="2">N/A</td><td colspan="2">N/A</td><td>2 x 75 W SW</td><td>2 x 75 W or 150 W SW or 2 x 300 W DW</td><td colspan="2">N/A</td></tr><tr><td>Rack height (U)</td><td>2</td><td></td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>1</td></tr><tr><td>Integrated security</td><td colspan="15">TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks</td></tr></table>	Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250	Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity	Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and visualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads	Type of processor	2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores			Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)				8 (1 TB)		4 (128 GB)			Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 4 x 3.5"	24 x 2.5" 12 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5" 24 x 2.5" 10 x 2.5" 12 x 3.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 4 x 3.5"	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 3.5"	4 x 2.5"	6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)	NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A		4	6	N/A		Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2	Gen3 PCIe slots up to:			N/A		2		1		1		N/A		1		N/A		Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		N/A		2 x 75 W SW	2 x 75 W or 150 W SW or 2 x 300 W DW	N/A		Rack height (U)	2		1	2	1	2	1	2	1	1	2	1	2	1	1	Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks														
Rack Server	R750	R750xa	R650	R752s	R652s	R751s	R651s	R750xs	R650xs	R450	R550	XR11	XR12	R350	R250																																																																																																																																																																																			
Key attributes	Outstanding performance for the most demanding workloads	Highly intensive GPU workloads	High scalability, optimized workload performance	Powerful performance and flexibility	Dense virtualization	Powerful performance and scalability	High density compute	Purpose-built 2U server for growing scale-out solutions	Purpose-built, full performance 1U server for dense, fast growing scale-out solutions	Value and density-focused, built for general purpose IT	Versatile, value-optimized, virtualization-ready, built for general purpose IT	Edge-centric, short depth and rugged with reverse mounting options	Edge-centric, short depth and rugged with reverse mounting options	Powerful performance in 1U server for productivity and data intensive applications	Powerful compute for common business applications and streamlines productivity																																																																																																																																																																																			
Target workloads	Database and analytics, HPC, traditional corporate IT, VDI, AI, or ML environments	AI, ML, or DL training or inferencing, HPC, and visualization environments	Mixed workload standardization, database and analytics, HFT, traditional corporate IT, VDI, HPC, AI, or ML environments	All flash SDS, VDI, and data analytics	HPC, Dense VDI, and Virtualization	SDS, Virtualization, and Data Analytics	Virtualization, HCI and NFV	Virtualization, medium VM density or VDI, and scale-out database workloads	Virtualization, cloud, scale-out database and high-performance compute workloads	Small IT infrastructure, light VM, small business specific workloads	Small IT infrastructure, light VM density, small business specific workloads	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Telco/5G (MEC, CDN, vRAN), Military, Retail (Analytics - video surveillance/POSIOT aggregation)	Small mid-sized businesses, remote office/branch office, collaboration and sharing, data analytics and virtualization workloads	Small mid-sized businesses, remote office/branch office, collaboration and sharing, mail/messaging and file/print workloads																																																																																																																																																																																			
Type of processor	2 x 3rd Generation Intel® Xeon® Scalable processors; up to 40 cores per processor		2 x 2nd or 3rd Generation AMD EPYC™ processors; up to 64 cores per processor		1 x 2nd or 3rd Generation AMD EPYC™ processor; up to 64 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 32 cores per processor		2 x 3rd Generation Intel® Xeon® Scalable processors; up to 24 cores per processor		1 x 3rd Generation Intel® Xeon® Scalable processors; up to 36 cores per processor		1 x Intel Xeon E-2300 series processors with up to 8 cores or 1 x Intel Pentium processor with up to 2 cores																																																																																																																																																																																					
Memory (DDR4 DIMM slots & max)	32 (8 TB)		32 (4 TB)		16 (2 TB)		16 (1 TB)				8 (1 TB)		4 (128 GB)																																																																																																																																																																																					
Disk drives up to:	8 x 2.5" 16 x 2.5" 24 x 2.5" 12 x 3.5" 2 x 2.5" or 4 x 2.5" (rear)	6 x 2.5" 8 x 2.5"	4 x 3.5" 8 x 2.5" 10 x 2.5" 2 x 2.5" (rear)	26 x 2.5" 12 x 3.5" 4 x 3.5"	12 x 2.5" 4 x 3.5"	24 x 2.5" 12 x 3.5"	8 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5" 24 x 2.5" 10 x 2.5" 12 x 3.5" 8 x 3.5"	4 x 3.5" 8 x 2.5" 4 x 3.5"	4 x 3.5" 8 x 2.5"	16 x 2.5" 8 x 3.5"	4 x 2.5"	6 x 2.5"	4 x 3.5" 8 x 2.5"	4 x 3.5" 2 x 3.5" (cabled) 4 x 3.5" (cabled)																																																																																																																																																																																			
NVMe drives up to:	24	8	12	24	12	24	10	8	10	N/A		4	6	N/A																																																																																																																																																																																				
Gen4 PCIe slots up to:	8	8	3	8	3	2	1	5	3	2	3	3	5	3	2																																																																																																																																																																																			
Gen3 PCIe slots up to:			N/A		2		1		1		N/A		1		N/A																																																																																																																																																																																			
Accelerator support up to:	2 x 300 W DW or 4 x 150 W SW or 6 x 75 W SW	4 x 150 W SW or 4 x 300 W DW or 2 x 75 W SW	3 x 75 W SW	3 x 300 W DW or 6 x 75 W SW	3 x SW	4 x SW; 1 x DW; 1 x FPGA	1 x SW	N/A		N/A		2 x 75 W SW	2 x 75 W or 150 W SW or 2 x 300 W DW	N/A																																																																																																																																																																																				
Rack height (U)	2		1	2	1	2	1	2	1	1	2	1	2	1	1																																																																																																																																																																																			
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks																																																																																																																																																																																																	
	PowerEdge Reference Guide, 1.																																																																																																																																																																																																	

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))											
	Rack Server	R940	R940xa	R840	R740xd	R740	R740xd2	R640	R540	R440	R340	R240
Key attributes	Powerful performance	Extreme acceleration	Turbocharge data analytics	Scalable storage performance	Optimal application performance	Enterprise content server	Performance and density	Balanced and adaptable	Scale-out computing	Accelerate business growth	Compute made simple	
Target workloads	In-memory databases	GPU database acceleration and machine learning	Data-intensive workloads, HFT, and dense virtualization	SDS, service providers, and big data servers	VDI and cloud workloads	Media streaming and SDS	Dense scale-out data center computing and storage	Mail messaging and virtualization	HPC, web tech, and scale-out infrastructure	ROBO productivity and data-intensive applications	Small business and service provider workloads	
Type of processor	4 x 2 nd Generation Intel® Xeon® Scalable processors			2 x 2 nd Generation Intel® Xeon® Scalable processors			16 (1 TB)			4 (64 GB)		
Memory (DDR4 DIMM slots & max)	48 (15.36 TB)			24 (7.68 TB)		26 x 3.5"	24 (7.68 TB)	16 (1 TB)				
Disk drives up to:	24 x 2.5"	32 x 2.5"	26 x 2.5"	32 x 2.5" 16 x 3.5"	16 x 2.5" 8 x 3.5"	26 x 3.5" 16 x 3.5" + 10 x 2.5"	12 x 2.5" 4 x 3.5"	14 x 3.5"	10 x 2.5" 4 x 3.5"	8 x 2.5" 4 x 3.5"	4 x 2.5" 4 x 3.5"	
NVMe drives up to:	12	4	24	N/A			10	N/A			N/A	
Gen4 PCIe slots up to:	13	12	6	8	5			3	5	2	2	
Gen3 PCIe slots up to:												
Accelerator support up to:	N/A	4 x DW GPUs or 4 x DW or 8 x SW FPGAs	2 x DW GPUs or 2 x SW or DW FPGAs	3 x DW or 6 x SW GPUs or 3 x DW or 4 x SW FPGAs	N/A			1 x SW GPU or 1 x SW FPGA	N/A			
Rack height (U)	3	4	2				1	2	1	1		
Integrated security	TPM 1.2/2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ, Cryptographically Signed Firmware, Chassis Intrusion Alert, and Secure Boot being standard security on all racks. Integrated security features such as Silicon Root of Trust, System Lockdown (requires iDRAC Enterprise or Datacenter), and System Erase on all racks.											

PowerEdge Reference Guide, 2.

Thermal

Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see [Environmental Specifications](#)). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Dell XR8610t and XR8620t Technical Guide, 39.

Multi Vector Cooling

An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow.

This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced.

Dell XR8610t and XR8620t Technical Guide, 39.

PowerEdge Reference Guide, 2.

Thermal

Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see [Environmental Specifications](#)). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption.

Dell XR8610t and XR8620t Technical Guide, 39.

Multi Vector Cooling

An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm ‘learn’ this information and the card is automatically cooled with the proper airflow.

This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced.

Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))										
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges. <table border="0"> <tr> <td data-bbox="762 513 951 602">1. Reliability</td><td data-bbox="951 513 1671 602"> • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. </td></tr> <tr> <td data-bbox="762 602 951 699">2. Performance</td><td data-bbox="951 602 1671 699"> • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. </td></tr> <tr> <td data-bbox="762 699 951 805">3. Efficiency</td><td data-bbox="951 699 1671 805"> • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fan speeds while meeting the above Reliability and Performance tenets. </td></tr> <tr> <td data-bbox="762 805 951 902">4. Management</td><td data-bbox="951 805 1671 902"> • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. </td></tr> <tr> <td data-bbox="762 902 951 1000">5. Forward Compatibility</td><td data-bbox="951 902 1671 1000"> • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. </td></tr> </table> Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.	1. Reliability	• Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life.	2. Performance	• Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations.	3. Efficiency	• 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fan speeds while meeting the above Reliability and Performance tenets.	4. Management	• System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads.	5. Forward Compatibility	• Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced.
1. Reliability	• Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life.										
2. Performance	• Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations.										
3. Efficiency	• 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fan speeds while meeting the above Reliability and Performance tenets.										
4. Management	• System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads.										
5. Forward Compatibility	• Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced.										

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	 Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points into the rack. Above it, a bracket groups three dots labeled 'CFM#N', 'CFM#2', and 'CFM#1'. A circled '1' is next to the text 'Airflow reporting (per server)'. Server Rack: A central server rack is shown with a circled '2' at its base labeled 'Custom Delta-T'. Hot Aisle: On the right, a red arrow points away from the rack, labeled 'Exhaust Temp Control (per server)' with a circled '3'. PCIe Thermal Control: At the bottom, a bracket groups two elements: - A circled '4' next to 'Inlet PCIe Temp' with a small server component icon. - A circled '5' next to 'PCIe Airflow (LFM)' with a blue arrow pointing right. CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage.
8[a] by a declination index value calculation unit, calculating an index value indicating a degree of declination of a component temperature of an electronic component included in ICT (Information and Communication Technology) equipment, based on a detection result of a first temperature sensor that detects the component	The Accused Dell '320 Products, by a declination index value calculation unit, calculate an index value indicating a degree of declination of a component temperature of an electronic component included in ICT (Information and Communication Technology) equipment, based on a detection result of a first temperature sensor that detects the component temperature of the electronic component and a second temperature sensor that detects a temperature of an intake air. <i>See supra</i>, claims 1[e], 7[b], and 8[pre]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
temperature of the electronic component and a second temperature sensor that detects a temperature of an intake air; and	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior. Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results





Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed <table border="1"><thead><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr></thead><tbody><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></tbody></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.


'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a vertical stack of servers is shown. A bracket labeled '1 Airflow reporting (per server)' points to the servers. To the right of the servers, a vertical list of dots is labeled 'CFM#N', 'CFM#2', and 'CFM#1' from top to bottom. A large blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points upwards from the bottom left towards the servers. Server Rack: A central server rack is shown with a bracket labeled '2 Custom Delta-T' pointing to it. Hot Aisle: On the right, a red arrow points upwards and is labeled '3 Exhaust Temp Control (per server)'. PCIe Thermal Control: At the bottom, a bracket labeled 'PCIe Thermal Control' points to two elements: '4 Inlet PCIe Temp' (with a fan icon) and '5 PCIe Airflow (LFM)' (with a horizontal blue arrow icon). CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.
8[c] by a control unit, controlling the number of rotations of a cooling fan based on the index value calculated by the declination index value calculation unit and the temperature of the intake air by the second temperature sensor.	The Accused Dell '320 Products, by a control unit, control the number of rotations of a cooling fan based on the index value calculated by the declination index value calculation unit and the temperature of the intake air by the second temperature sensor. <i>See supra</i>, claims 1[f], 7[c], and 8[b]. <i>See, e.g.,</i>

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Thermal management of the platform helps delivers high performance for the right amount of cooling to components at the lowest fan speeds across a wide range of ambient temperatures from -5°C to 55°C (23°F to 131°F) and to extended ambient temperature range from -20°C to 65°C (see Environmental Specifications). It might be reflected in lower overall power consumption (fans, platform, cooling/heating, data center power consumption, etc.) and greater acoustical versatility. PowerEdge servers have an extensive collection of sensors that automatically track thermal activity, which helps regulate temperature thereby reducing server noise and power consumption. Dell XR8610t and XR8620t Technical Guide, 39. Multi Vector Cooling An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow (LFM – Linear Feet per Minute) requirement specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8610t and XR8620t Technical Guide, 39.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal design Thermal management of the platform helps deliver high performance with the right amount of cooling to components, while maintaining the lowest fan speeds possible. This is done across a wide range of ambient temperatures from 10°C to 35°C (50°F to 95°F) and to extended ambient temperature ranges.  1. Reliability • Component hardware reliability remains the top thermal priority. • System thermal architectures and thermal control algorithms are designed to ensure there are no tradeoffs in system level hardware life. 2. Performance • Performance and uptime are maximized through the development of cooling solutions that meet the needs of even the densest of hardware configurations. 3. Efficiency • 15G servers are designed with an efficient thermal solution to minimize power and airflow consumption, and/or acoustics for acoustical deployments. • Dell's advanced thermal control algorithms enable minimization of system fans speeds while meeting the above Reliability and Performance tenets. 4. Management • System management settings are provided such that customers have options to customize for their unique hardware, environments, and/or workloads. 5. Forward Compatibility • Forward compatibility means that thermal controls and thermal architecture solutions are robust to scale to new components that historically would have otherwise required firmware updates to ensure proper cooling. • The frequency of required firmware updates is thus reduced. Figure 31. Thermal design characteristics Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	The thermal design of the PowerEdge R760xd2 reflects the following: ● Optimized thermal design: The system layout is architected for optimum thermal design. ● System component placement and layout are designed to provide maximum airflow coverage to critical components with minimum expense of fan power. ● Comprehensive thermal management: The thermal control system regulates the fan speed based on several different responses from all system-component temperature sensors, as well as inventory for system configurations. Temperature monitoring includes components such as processors, DIMMs, chipset, the inlet air ambient, hard disk drives, and OCP. ● Open and closed loop thermal fan speed control: Open loop thermal control uses system configuration to determine fan speed based on inlet air ambient temperature. Closed loop thermal control method uses feedback temperatures to dynamically determine proper fan speeds. ● User-configurable settings: With the understanding and realization that every customer has unique set of circumstances or expectations from the system, in this generation of servers, we have introduced limited user- configurable settings residing in the iDRAC BIOS setup screen. For more information, see the Dell PowerEdge R760xd2 Installation and Service Manual at PowerEdge Manuals and "Advanced Thermal Control: Optimizing across Environments and Power Goals" on Dell.com. ● Cooling redundancy: The PowerEdge R760xd2 allows N+1 fan redundancy, allowing continuous operation with one fan failure in the system. ● Environmental Specifications: The optimized thermal management makes the PowerEdge R760xd2 reliable under a wide range of operating environments. Dell R760xd2 Technical Guide, 37.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal chamber testing For the purpose of conducting this test, we placed a 2U XR8000 inside the thermal chamber in our test lab. While in the thermal chamber, we ran the iPerf3 workload on the system for more than eight hours, stressing the system from 5-20%. We measured power consumption and fan speed using iDRAC at 10-degree intervals of Celsius temperature from 0C to 55C. The iPerf3 throughput measured for 1GB, 10GB, and 25GB seemed consistent across the entire temperature range, with no impact on performance as temperature increased. The fan speed and power consumption increased with temperature, which is the expected behavior.  Figure 2. Thermal chamber in the Dell performance testing lab Dell XR8000 Thermal Design.

**'320 Patent Claim²
(P.R. 3-1(a))**

**Accused Instrumentalities And Where Each Claim Element Is Found³
(P.R. 3-1(b)-(c))**

Results




Figure 3. Constant networking performance with varying temperature and fan speed

Figure 3 shows that as the temperature and fan speed increases, the iPerf3 throughput stays the same. Fan speed is only 14% for temperatures near 20°C.


Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))																								
	Power consumption and Fan Speed <table><thead><tr><th>Temperature (°C)</th><th>Fan speed (% of total)</th><th>System Power consumption (W)</th></tr></thead><tbody><tr><td>0</td><td>5</td><td>216</td></tr><tr><td>10</td><td>5</td><td>219</td></tr><tr><td>20</td><td>14</td><td>254</td></tr><tr><td>30</td><td>42</td><td>331</td></tr><tr><td>40</td><td>64</td><td>512</td></tr><tr><td>50</td><td>86</td><td>748</td></tr><tr><td>55</td><td>100</td><td>1014</td></tr></tbody></table> Figure 4. Power consumption and fan speed Figure 4 shows that as temperature increases, Chassis power consumption for the system increases. It is 254W at 20°C. Dell XR8000 Thermal Design.	Temperature (°C)	Fan speed (% of total)	System Power consumption (W)	0	5	216	10	5	219	20	14	254	30	42	331	40	64	512	50	86	748	55	100	1014
Temperature (°C)	Fan speed (% of total)	System Power consumption (W)																							
0	5	216																							
10	5	219																							
20	14	254																							
30	42	331																							
40	64	512																							
50	86	748																							
55	100	1014																							

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	A deep dive into the results The consistent performance with increasing temperature and power can be attributed to several design considerations when designing and building these edge/telecom servers: • RAF: The Reverse Airflow option offered in these servers is carried from Dell's innovation in Multi-Vector Cooling technology. While most of the innovations for MVC center around optimizing thermal controls and management, the physical cooling hardware and its architecture layout help. XR servers are shallower, which can mean less airflow impedance, resulting in more airflow. • Fans: XR servers are designed with high-performance fans, which have significantly increased airflow performance over previous fan generations. • Layout: The T-shape system motherboard layout, along with PSUs that are located at each corner of the chassis, allows improved airflow balancing and system cooling, and consequently, improved system cooling efficiency. This layout also improves PSU cooling due to reduced risk from high pre-heat coming from CPU heat sinks. The streamlined airflow helps with PCIe cooling and enables support for PCIe Gen5 adapters. • Smaller PSU Form Factor: In the 1U systems, a new, narrower, 60mm form factor PSU is implemented to increase the exhaust path space. • XR servers usually support CPUs with higher TCASE requirements. TCASE stands for Case Temperature and is the maximum temperature allowed at the processor Integrated Heat Spreader (IHS)[2]. For more details about the design considerations used for edge servers, see the blog Computing on the Edge—Other Design Considerations for the Edge. Dell XR8000 Thermal Design. iDRAC To best supplement the improved cooling hardware, the PowerEdge engineering team carried on the key features from the previous generation of PowerEdge servers to deliver autonomous thermal solutions capable of cooling next-generation PowerEdge servers. An iDRAC feature in XR8000 detects Dell PCIe cards and automatically delivers the correct airflow to the slot to cool that card. When non-Dell PCIe cards are detected, the customer is given the option to enter the airflow requirement (LFM – Linear Feet per Minute) as specified by the card manufacturer. iDRAC and the fan algorithm 'learn' this information and the card is automatically cooled with the proper airflow. This feature saves power by not having to run the fans to cool the worst-case card in the system. Noise is also reduced. Dell XR8000 Thermal Design.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Thermal Manage – Feature Overview Thermal Manage allows customers to customize the thermal operation of their PowerEdge servers with the following benefits: • Optimize server-related power and cooling efficiencies across their datacenters. • Integrates seamlessly with OpenManage Enterprise Power Manager for optimized management experience. • Provides a state-of-the-art PCIe cooling management dashboard. Represented in the following diagram (See figure 1) and listed below is a summary of the features and its utilities. 1. System Airflow Consumption: Displays the real-time system airflow consumption (in CFM), allowing airflow balancing at rack and datacenter level. 2. Custom Delta-T: Limit air temperature rise from inlet air to exhaust to right-size your infrastructure level cooling. 3. Exhaust Temperature Control: Specify the temperature limit of the air exiting the server to match your datacenter needs. 4. Custom PCIe inlet temperature: Choose the right input inlet temperature to match 3rd party device requirements. 5. PCIe airflow settings: Provides a comprehensive PCIe device cooling view of the server and allows cooling customization of 3rd party cards. Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Server Rack + Datacenter Thermal Control The diagram illustrates a server rack in a datacenter with the following components and processes: Cold Aisle: On the left, a vertical stack of servers is shown. A bracket labeled '1 Airflow reporting (per server)' points to the servers. To the right of the servers, a vertical list of dots is labeled 'CFM#N', 'CFM#2', and 'CFM#1' from top to bottom. A large blue arrow labeled 'Inlet Airflow (Sum of CFMs)' points upwards from the bottom left towards the servers. Server Rack: A central server rack is shown with a bracket labeled '2 Custom Delta-T' pointing to it. Hot Aisle: On the right, a red arrow points upwards and is labeled '3 Exhaust Temp Control (per server)'. PCIe Thermal Control: At the bottom, a bracket labeled 'PCIe Thermal Control' points to two items: '4 Inlet PCIe Temp' (with a small fan icon) and '5 PCIe Airflow (LFM)' (with a large blue arrow pointing right). CFM = Airflow in Cubic Feet/min Thermal Manage.

'320 Patent Claim ² (P.R. 3-1(a))	Accused Instrumentalities And Where Each Claim Element Is Found ³ (P.R. 3-1(b)-(c))
	Using the real-time airflow telemetry, a datacenter can create a good balance of airflow delivery vs. airflow demand at the server. A reduction in CFM also can be monetized on a dollar/CFM basis. • In an example analysis using a 17 KW rack, a drop in CFM by 10% could result in capital savings (CRAC costs of \$257/rack) and an annual operational savings of \$93 per rack based on the typical energy cost and data center efficiencies assumed. • However, the greater benefit is the potential ability to fit more racks on the floor (or more servers in a rack), if airflow balancing is achieved by closely matching the server/rack airflow consumption. iDRAC Thermal Manage features require an iDRAC Datacenter license. Here is an image from the iDRAC GUI showing the thermal telemetry and customization options:  The screenshot shows the 'Cooling Configuration' section of the iDRAC GUI. It includes settings for 'Automatic Fan Speed Calculation' (set to 'Default Thermal Profile Settings (Maximum Performance)'), 'Fan Speed Offset' (set to 'Off'), and 'Thresholds'. Under 'Thresholds', it displays 'Current Temperature Readings' for System Inlet (22°C / 71.6°F) and System Exhaust (25°C / 77.0°F). It also shows 'Exhaust Temperature Limits' with a 'Set Maximum Exhaust Temperature Limit' dropdown set to 'Default, 70 °C (158.0 °F)' and a 'Target Exhaust Temperature Limit' set to '70°C (158.0°F)'. At the bottom, there is a 'Minimum Fan Speed in PWM (% of Max)' dropdown set to 'Default' with a range of 25-100. 'Apply' and 'Reboot' buttons are at the bottom right. Thermal Manage. This is a software limitation that may require access to the underlying source code. Accordingly, pursuant to Paragraph 3(a)(i) of the Discovery Order, Cloud Byte may supplement these contentions with additional information after Dell's source code has been produced.