

EXHIBIT D
U.S. Patent No. 12,301,627

As used herein, the term “Accused ’627 Defender Products” means:

- (a) Microsoft Defender XDR;
- (b) Any other products that utilize the libraries, applications, scripts, packages, or other modules that implement the functionality described below in a manner not materially different with respect to the claims charted below;
- (c) any other products that infringe the asserted claims for analogous reasons to those described below; and,
- (d) Microsoft products that practice one of more claims of the ’627 Patent.

This claim chart for the ’627 Patent covers all Accused ’627 Defender Products. The theory of infringement described below in connection with the Asserted Claims is analogous to the theory of infringement for all the Accused ’627 Defender Products.

I. Claim 1

A computer system comprising:
a hardware memory, wherein
the computer system is
configured to execute software
instructions stored on
nontransitory machine-readable
storage media comprising
software instructions that:

The Accused '627 Defender Products include a computer system comprising a hardware memory, wherein the computer system is configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that perform as discussed below.

For example, Defender XDR “is a cloud-based, unified, pre- and post-breach enterprise defense suite.”¹

Pilot and deploy Microsoft Defender XDR

Applies to:

- Microsoft Defender XDR

This series of articles steps you through the entire process of piloting the components of Microsoft Defender XDR in your production tenant so you can evaluate their features and capabilities and then completing the deployment across your organization.

An eXtended detection and response (XDR) solution is a step forward in cyber security because it takes the threat data from systems that were once isolated and unifies them so that you can see patterns and act on suspected cyberattacks faster.

Microsoft Defender XDR:

- Is an XDR solution that combines the information on cyberattacks for identities, endpoints, email, and cloud apps in one place. It leverages artificial intelligence (AI) and automation to automatically stop some types of attacks and remediate affected assets to a safe state.
- Is a cloud-based, unified, pre- and post-breach enterprise defense suite. It coordinates prevention, detection, investigation, and response across identities, endpoints, email, cloud apps, and their data.

Defender XDR “operates in Microsoft Azure data centers”:²

¹ Microsoft, *Pilot and deploy Microsoft Defender XDR*, available at <https://learn.microsoft.com/en-us/defender-xdr/pilot-deploy-overview> [hereinafter “*Deploy Defender XDR*”].

² Microsoft, *Data security and retention in Microsoft Defender XDR*, available at <https://learn.microsoft.com/en-us/defender-xdr/data-privacy?view=o365-worldwide> [hereinafter *Data in XDR*].

Data storage location

Microsoft Defender XDR operates in Microsoft Azure data centers in the following geographical regions:

- **European Union:** North Europe and West Europe
- **United Kingdom:** UK South and UK West
- **United States:** East US 2 and Central US
- **Australia:** Australia East and Australia Southeast
- **Switzerland:** Switzerland North and Switzerland West
- **India:** Central India and South India
- **UAE:** UAE North and UAE Central

These data centers “house[] thousands of powerful computers, or ‘servers,’” examples of which can be seen in the following photograph reproduced from Microsoft’s literature:³



³ Microsoft, *Microsoft Datacenters: Powering Our Daily Lives*, available at <https://datacenters.microsoft.com/WhatIsADatacenter/> [hereinafter *Microsoft Datacenters*].

store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges,

The software instructions of Defender XDR store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges.

For example, Defender XDR “use[s] interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. . . . The graphs generated in the Defender portal are composed of nodes and edges”:⁴

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

As another example, Microsoft’s documentation explains that the “graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes”:⁵


⁴ Microsoft, *Understanding graphs and visualizations in Microsoft Defender*, available at <https://learn.microsoft.com/en-us/defender-xdr/understand-graph-icons> [hereinafter *Understanding Graph Icons*].

⁵ Microsoft, *Investigate incidents in the Microsoft Defender portal*, available at <https://learn.microsoft.com/en-us/defender-xdr/investigate-incidents> [hereinafter *Investigate Incidents*].


The graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes.

From the graph, you can:

- Play the alerts and the nodes on the graph as they occurred over time to understand the chronology of the attack.



- Open an entity pane, allowing you to review the entity details and act on remediation actions, such as deleting a file or isolating a device.



- Highlight the alerts based on the entity to which they are related.
- Hunt for entity information of a device, file, IP address, URL, user, email, mailbox, or cloud resource.

	As another example, Microsoft’s documentation explains that a node of the graphs used by Defender XDR “pertains to an entity in your environment (for example, a device, user account, or IP address, among others)”. ⁶
--	--

⁶ *Understanding Graph Icons.*

Nodes

A **node** pertains to an entity in your environment (for example, a device, user account, or IP address, among others). Defender portal graphs usually depict nodes as any of the following circular icons:

Expand table

Icon	Node type	Entity type examples
	General	App service plan
	Compute	Device, virtual machine, Microsoft Azure Logic App
	Networking	Interface, public IP address, network security group
	Data	SQL data store, Azure Monitor Log Analytics workspace, storage account, Azure Event Hubs
	Containers	Kubernetes cluster
	Keys & secrets	Key vault
	DevOps	Azure DevOps repositories
	APIs	Cloud applications
	Identity & access	User account, Microsoft Entra ID service principal
	IoT	
	Certificate	
	IP address	
	Subscriptions	

As another example, Microsoft's documentation explains that nodes can be represented as an **ExposureGraphNode**s table of "organizational entities and their properties," which "include entities like devices, identities, user groups, and cloud assets such as virtual machines (VMs), storage, and containers. Each node corresponds to an individual entity and encapsulates information about its characteristics, attributes, and security related insights within the organizational structure":⁷

ExposureGraphNode

06/20/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Applies to:

- Microsoft Defender XDR
- Microsoft Security Exposure Management (public preview)

🔔 Important

Some information relates to prereleased product which may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided here.

The **ExposureGraphNode**s table in the **advanced hunting** schema contains organizational entities and their properties. These include entities like devices, identities, user groups, and cloud assets such as virtual machines (VMs), storage, and containers. Each node corresponds to an individual entity and encapsulates information about its characteristics, attributes, and security related insights within the organizational structure. Use this reference to construct queries that return information from this table.

As another example, edges can be represented as an **ExposureGraphEdge**s table of "relationships between entities and assets in the enterprise exposure graph":⁸

⁷ Microsoft, *ExposureGraphNode*s, available at <https://learn.microsoft.com/en-us/defender-xdr/advanced-hunting-exposuregraphnodes-table> [hereinafter *ExposureGraphNode*s].

⁸ Microsoft, *ExposureGraphEdge*s, available at <https://learn.microsoft.com/en-us/defender-xdr/advanced-hunting-exposuregraphedges-table> [hereinafter *ExposureGraphEdge*s].


	ExposureGraphEdges Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal Applies to: • Microsoft Defender XDR • Microsoft Security Exposure Management (public preview) Important Some information relates to prereleased product which may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided here. The <code>ExposureGraphEdges</code> table in the <code>advanced hunting</code> schema provides visibility into relationships between entities and assets in the enterprise exposure graph. This visibility can help uncover critical organizational assets and explore entity relationships and attack paths. Use this reference to construct queries that return information from this table. As another example, Defender XDR includes a “hunting graph,” which is “composed of nodes and edges to represent entities in your environment (for example, a device, user account, or IP address, among others) and their relationships or connection properties, respectively”:⁹
--	--

⁹ Microsoft, *Hunt for threats using the hunting graph (Preview)*, available at <https://learn.microsoft.com/en-us/defender-xdr/advanced-hunting-graph> [hereinafter *Hunting Graph*].

Hunting graph features

The interactive graphs generated in the hunting graph are composed of **nodes** and **edges** to represent entities in your environment (for example, a device, user account, or IP address, among others) and their relationships or connection properties, respectively. [Learn more about graphs and visualizations in Microsoft Defender](#)

The lower right-hand corner of the graph also has control buttons that let you **Zoom in** and **Zoom out**, and view the graph's **Layers**.



Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

wherein the first graph is a directed graph,

The first graph of the Accused '627 Defender Products is a directed graph.

For example, ExposureGraphEdges table schema includes information about “source nodes” and “target nodes.”¹⁰

For example, Microsoft’s documentation explains that the edges of Defender XDR’s graphs “indicate the relationship or connection properties between two nodes” and include “directional arrows.”¹¹

As another example, as discussed above, Defender XDR includes a “hunting graph,” which is depicted by Microsoft’s documentation as a directed graph:¹²



¹⁰ ExposureGraphEdges.

¹¹ Understanding Graph Icons.

¹² Hunting Graph.

	As another example, Microsoft’s documentation explains that users must take caution “to identify and input the correct start and end entities, as the generated graph will be directional.”¹³ <table><tr><th>Scenario</th><th>Description</th><th>Inputs</th></tr><tr><td>Paths between two entities</td><td> Provide two entities (nodes) to view the paths between them. Use this scenario if you want to discover if there’s a path leading from one entity to another. </td><td> • Start Entity • End Entity Note: Make sure to identify and input the correct start and end entities, as the generated graph will be directional. </td></tr></table> Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.	Scenario	Description	Inputs	Paths between two entities	Provide two entities (nodes) to view the paths between them. Use this scenario if you want to discover if there’s a path leading from one entity to another.	• Start Entity • End Entity Note: Make sure to identify and input the correct start and end entities, as the generated graph will be directional.
Scenario	Description	Inputs					
Paths between two entities	Provide two entities (nodes) to view the paths between them. Use this scenario if you want to discover if there’s a path leading from one entity to another.	• Start Entity • End Entity Note: Make sure to identify and input the correct start and end entities, as the generated graph will be directional.					
wherein the first plurality of entities comprises a plurality of accounts and a plurality of resources, and	The first plurality of entities of the Accused ’627 Defender Products comprises a plurality of accounts and a plurality of resources. For example, Microsoft’s documentation explains that Defender XDR stores graphs in which the nodes include accounts and resources: “[a] node pertains to an entity in your environment (for example, a device, user account, or IP address, among others.”¹⁴ Examples of node types and their icon representations are reproduced below:						

¹³ *Id.*

¹⁴ *Understanding Graph Icons.*

Nodes

A **node** pertains to an entity in your environment (for example, a device, user account, or IP address, among others). Defender portal graphs usually depict nodes as any of the following circular icons:

Expand table

Icon	Node type	Entity type examples
	General	App service plan
	Compute	Device, virtual machine, Microsoft Azure Logic App
	Networking	Interface, public IP address, network security group
	Data	SQL data store, Azure Monitor Log Analytics workspace, storage account, Azure Event Hubs
	Containers	Kubernetes cluster
	Keys & secrets	Key vault
	DevOps	Azure DevOps repositories
	APIs	Cloud applications
	Identity & access	User account, Microsoft Entra ID service principal
	IoT	
	Certificate	
	IP address	
	Subscriptions	

	As another example, as discussed above, Defender XDR includes a “hunting graph,” which includes nodes corresponding to “entities in your environment (for example, a device, user account, or IP address, among others).”¹⁵ Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities;	Each edge of the first plurality of edges of Defender XDR corresponds to a respective relationship between a respective pair of entities. For example, the edges that comprise Defender XDR’s graphs “indicate[] the relationship or connection properties between two nodes.”¹⁶ Examples of edge types and their icon representations are reproduced below:

¹⁵ *Hunting Graph.*

¹⁶ *Understanding Graph Icons.*

	Edges An edge indicates the relationship or connection properties between two nodes. The Defender portal graphs depicts an edge as lines or directional arrows that might have the following icons: Expand table <table><tr><th>Icon</th><th>Edge type</th></tr><tr><td></td><td>Contains</td></tr><tr><td></td><td>Routes traffic to</td></tr><tr><td></td><td>Has permission to / Has role on</td></tr><tr><td></td><td>Can authenticate as / Can authenticate to</td></tr><tr><td></td><td>Pushes</td></tr><tr><td></td><td>Maintains</td></tr><tr><td></td><td>Application</td></tr><tr><td></td><td>Moves data to</td></tr><tr><td></td><td>Exposed to internet</td></tr><tr><td></td><td>Can interactive logon to / Can logon over the network to / Can remote interactive logon to</td></tr><tr><td></td><td>Runs on</td></tr><tr><td></td><td>Provisions</td></tr><tr><td></td><td>Identified as owner of</td></tr><tr><td></td><td>Member of</td></tr><tr><td></td><td>Is running</td></tr><tr><td></td><td>Generic / Affects</td></tr><tr><td></td><td>Created from / Used to create</td></tr></table>	Icon	Edge type		Contains		Routes traffic to		Has permission to / Has role on		Can authenticate as / Can authenticate to		Pushes		Maintains		Application		Moves data to		Exposed to internet		Can interactive logon to / Can logon over the network to / Can remote interactive logon to		Runs on		Provisions		Identified as owner of		Member of		Is running		Generic / Affects		Created from / Used to create
Icon	Edge type																																				
	Contains																																				
	Routes traffic to																																				
	Has permission to / Has role on																																				
	Can authenticate as / Can authenticate to																																				
	Pushes																																				
	Maintains																																				
	Application																																				
	Moves data to																																				
	Exposed to internet																																				
	Can interactive logon to / Can logon over the network to / Can remote interactive logon to																																				
	Runs on																																				
	Provisions																																				
	Identified as owner of																																				
	Member of																																				
	Is running																																				
	Generic / Affects																																				
	Created from / Used to create																																				

	As another example, as discussed above, Defender XDR includes a “hunting graph,” which includes “relationships or connection properties” between graph nodes.¹⁷ As another example, Microsoft’s documentation explains that Defender XDR “provides visibility into relationships between entities and assets in the enterprise exposure graph”:¹⁸ ExposureGraphEdges 06/20/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal Applies to: • Microsoft Defender XDR • Microsoft Security Exposure Management (public preview) Important Some information relates to prereleased product which may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided here. The <code>ExposureGraphEdges</code> table in the advanced hunting schema provides visibility into relationships between entities and assets in the enterprise exposure graph. This visibility can help uncover critical organizational assets and explore entity relationships and attack paths. Use this reference to construct queries that return information from this table.
--	---

¹⁷ *Hunting Graph.*

¹⁸ *ExposureGraphEdges.*

	<table><tr><th>Column name</th><th>Data type</th><th>Description</th></tr><tr><td>EdgeId</td><td>string</td><td>Unique identifier for the relationship/edge</td></tr><tr><td>EdgeLabel</td><td>string</td><td>The edge label like "routes traffic to"</td></tr><tr><td>SourceNodeId</td><td>string</td><td>Node ID of the edge's source</td></tr><tr><td>SourceNodeName</td><td>string</td><td>Source node display name</td></tr><tr><td>SourceNodeLabel</td><td>string</td><td>Source node label</td></tr><tr><td>SourceNodeCategories</td><td>dynamic</td><td>Categories list of the source node in JSON format</td></tr><tr><td>TargetNodeId</td><td>string</td><td>Node ID of the edge's target</td></tr><tr><td>TargetNodeName</td><td>string</td><td>Display name of the target node</td></tr><tr><td>TargetNodeLabel</td><td>string</td><td>Target node label</td></tr><tr><td>TargetNodeCategories</td><td>dynamic</td><td>The categories list of the target node in JSON format</td></tr><tr><td>EdgeProperties</td><td>dynamic</td><td>Optional data relevant for the relationship between the nodes in JSON format</td></tr></table> Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.	Column name	Data type	Description	EdgeId	string	Unique identifier for the relationship/edge	EdgeLabel	string	The edge label like "routes traffic to"	SourceNodeId	string	Node ID of the edge's source	SourceNodeName	string	Source node display name	SourceNodeLabel	string	Source node label	SourceNodeCategories	dynamic	Categories list of the source node in JSON format	TargetNodeId	string	Node ID of the edge's target	TargetNodeName	string	Display name of the target node	TargetNodeLabel	string	Target node label	TargetNodeCategories	dynamic	The categories list of the target node in JSON format	EdgeProperties	dynamic	Optional data relevant for the relationship between the nodes in JSON format
Column name	Data type	Description																																			
EdgeId	string	Unique identifier for the relationship/edge																																			
EdgeLabel	string	The edge label like "routes traffic to"																																			
SourceNodeId	string	Node ID of the edge's source																																			
SourceNodeName	string	Source node display name																																			
SourceNodeLabel	string	Source node label																																			
SourceNodeCategories	dynamic	Categories list of the source node in JSON format																																			
TargetNodeId	string	Node ID of the edge's target																																			
TargetNodeName	string	Display name of the target node																																			
TargetNodeLabel	string	Target node label																																			
TargetNodeCategories	dynamic	The categories list of the target node in JSON format																																			
EdgeProperties	dynamic	Optional data relevant for the relationship between the nodes in JSON format																																			
receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities,	The software instructions of the Accused ’627 Defender Products receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities. For example, Defender XDR “collect[s] . . . signals that are displayed in the portal.” Two kinds of signals include alerts, which Microsoft describes as “[s]ignals that result from various threat detection																																				

activities,” and incidents, which Microsoft describes as “[c]ontainers that include collections of related alerts and tell the full story of an attack:¹⁹

Incidents and alerts in the Microsoft Defender portal

01/06/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

The Microsoft Defender portal brings together a unified set of security services to reduce your exposure to security threats, improve your organizational security posture, detect security threats, and investigate and respond to breaches. These services collect and produce signals that are displayed in the portal. The two main kinds of signals are:

Alerts: Signals that result from various threat detection activities. These signals indicate the occurrence of malicious or suspicious events in your environment.

Incidents: Containers that include collections of related alerts and tell the full story of an attack. The alerts in a single incident might come from all Microsoft security and compliance solutions, as well as from vast numbers of external solutions collected through Microsoft Sentinel and Microsoft Defender for Cloud.

Microsoft’s documentation explains that Defender “us[es] AI to continually monitor its telemetry sources”:²⁰

¹⁹ Microsoft, *Incidents and alerts in the Microsoft Defender portal*, available at <https://learn.microsoft.com/en-us/defender-xdr/incidents-overview> [hereinafter *Incidents and Alerts*].

²⁰ *Id.*

	Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms: As another example, “[a]lerts in the Microsoft Defender portal come from many sources. These sources include the many services that are part of Microsoft Defender XDR, as well as other services with varying degrees of integration with the Microsoft Defender portal. For example, when Microsoft Sentinel is onboarded to the Microsoft Defender portal, the correlation engine in the Defender portal has access to all the raw data ingested by Microsoft Sentinel, which you can find in Defender's Advanced hunting tables”:²¹
--	---

²¹ *Id.*

Alert sources and threat detection

Alerts in the Microsoft Defender portal come from many sources. These sources include the many services that are part of Microsoft Defender XDR, as well as other services with varying degrees of integration with the Microsoft Defender portal.

For example, when Microsoft Sentinel is **onboarded** to the Microsoft Defender portal, the correlation engine in the Defender portal has access to all the raw data ingested by Microsoft Sentinel, which you can find in Defender's **Advanced hunting** tables.

Microsoft Defender XDR itself also creates alerts. Defender XDR's unique correlation capabilities provide another layer of data analysis and threat detection for all the non-Microsoft solutions in your digital estate. These detections produce Defender XDR alerts, in addition to the alerts already provided by Microsoft Sentinel's analytics rules.

Within each of these sources, there are one or more threat detection mechanisms that produce alerts based on the rules defined in each mechanism.

For example, Microsoft Sentinel has at least four different engines that produce different types of alerts, each with its own rules.

As another example, Microsoft's documentation explains that "[a]lerts are signals that result from various threat detection activities. These signals are produced by the many security services that reside in the Microsoft Defender portal, and they indicate the occurrence of malicious or suspicious events in your environment":²²

²² Microsoft, *Investigate alerts in Microsoft Defender XDR*, available at <https://learn.microsoft.com/en-us/defender-xdr/investigate-alerts?tabs=settings> [hereinafter *Investigate Alerts*].

Investigate alerts in Microsoft Defender XDR

06/04/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal


Note

This article describes security alerts in Microsoft Defender XDR. However, you can use alert policies to send email notifications to yourself or other admins when users perform specific activities in Microsoft 365. For more information, see [Alert policies in the Microsoft Defender portal](#).

Alerts are signals that result from various threat detection activities. These signals are produced by the many security services that reside in the Microsoft Defender portal, and they indicate the occurrence of malicious or suspicious events in your environment.

These suspicious events are typically part of a broader attack story. In the Microsoft Defender portal, alerts represent individual pieces of evidence that Defender XDR correlates together to form [incidents](#). Incidents tell the whole attack story; however, analyzing alerts can be valuable when deeper analysis is required.

The **Alerts queue** shows the current set of alerts. You can view the entire alerts queue from **Incidents & alerts > Alerts** on the quick launch of the [Microsoft Defender portal](#). You can also see the alerts for each incident on the **incidents queue**, and on each individual incident's page, on the **Alerts** tab.




As another example, in Defender XDR, “[e]vent or activity data populates tables about alerts, security events, system events, and routine assessments. Advanced hunting receives this data almost immediately after the sensors that collect them successfully transmit them to the corresponding cloud services. For example, you can query event data from healthy sensors on workstations or domain

	controllers almost immediately after they're made available on Microsoft Defender for Endpoint and Microsoft Defender for Identity. . . . Advanced hunting data uses the UTC (Universal Time Coordinated) timezone. . . . Advanced hunting results are converted to the timezone set in Defender XDR.”²³ As another example, Microsoft’s documentation explains that Defender XDR includes “[a]ttack stories” with a “graph” that “shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went”.²⁴
--	--

²³ Microsoft, *Proactively hunt for threats with advanced hunting in Microsoft Defender*, available at <https://learn.microsoft.com/en-us/defender-xdr/advanced-hunting-overview> [hereinafter *Advanced Hunting*].

²⁴ *Investigate Incidents*.

	 Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
based on a first portion of the streaming data, identify a first entity that does not correspond to any of the first plurality of	Based on a first portion of the streaming data, the software instructions of the Accused '627 Defender Products identify a first entity that does not correspond to any of the first plurality of nodes, wherein the first entity is not of the first plurality of entities.

nodes, wherein the first entity is not of the first plurality of entities,	For example, Defender XDR includes a device inventory. Microsoft’s documentation explains that “[d]uring the onboarding process, the Devices list is gradually populated with devices as they begin to report sensor data”:²⁵ During the onboarding process, the Devices list is gradually populated with devices as they begin to report sensor data. Use this view to track your onboarded endpoints as they come online, or download the complete endpoint list as a CSV file for offline analysis. Further, Microsoft’s documentation explains that Defender XDR performs “device discovery” by “collect[ing], prob[ing], or scan[ning] your network to discover unmanaged devices”:²⁶ Device discovery overview 05/08/2025 • Applies to: Microsoft Defender for Endpoint Plan 2 Protecting your environment requires taking inventory of the devices that are in your network. However, mapping devices in a network can often be expensive, challenging, and time-consuming. Microsoft Defender for Endpoint provides a device discovery capability that helps you find unmanaged devices connected to your corporate network without the need for extra appliances or cumbersome process changes. Device discovery uses onboarded endpoints, in your network to collect, probe, or scan your network to discover unmanaged devices. The device discovery capability allows you to discover: • Enterprise endpoints (workstations, servers, and mobile devices) that aren't yet onboarded to Defender for Endpoint • Network devices like routers and switches • IoT devices like printers and cameras
---	---

²⁵ Microsoft, *Device inventory*, available at <https://learn.microsoft.com/en-us/defender-endpoint/machines-view-overview> (emphasis omitted)

²⁶ Microsoft, *Device discovery overview*, available at <https://learn.microsoft.com/en-us/defender-endpoint/device-discovery> [hereinafter *Device Discovery*].

	For example, the “Device Inventory” interface includes a summary of devices discovered in the last 7 days:²⁷  Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
based on a second portion of the streaming data, wherein the	Based on a second portion of the streaming data, wherein the second portion is not identical to the first portion, the software instructions of the Accused '627 Defender Products identify a first relationship

²⁷ *Id.*

second portion is not identical to the first portion, identify a first relationship between a pair of entities of the first plurality of entities that does not correspond to any of the first plurality of edges,	between a pair of entities of the first plurality of entities that does not correspond to any of the first plurality of edges. For example, Defender XDR “us[es] AI to continually monitor its telemetry sources” in order to “automatically aggregate and correlate related alerts”:²⁸ Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms: Further, Defender keeps track of “which onboarded device a discovered device was seen by,” allowing SeenBy queries:²⁹ By invoking the SeenBy function, in your advanced hunting query, you can get detail on which onboarded device a discovered device was seen by. This information can help determine the network location of each discovered device and subsequently, help to identify it in the network.
---	--

²⁸ *Incidents and Alerts.*


²⁹ *Device Discovery.*

As another example, when “view[ing] the blast radius of a single node,” a “new graph view loads showing the 8 top-rated attack paths” that “shows the potential path from the entry point to this target,”³⁰ based on the most recent device discovery information:

View blast radius graphs

After selecting an incident from the list in the **Incidents** page, a graph view is displayed showing the entities and assets involved in the incident.

Select a node to open the context menu, then select **View blast radius**. To view the blast radius of a single node in a group, use the **ungroup** toggle above the grid to present all nodes.



A new graph view loads showing the 8 top-rated attack paths. A full list of the paths is visible on the right side panel when selecting **View full blast radius list** above the graph. From the list of reachable targets, you can further explore the path by selecting one of the listed targets. The right panel shows the potential path from the entry point to this target. Some nodes may not have paths associated with them.

Similarly, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph

³⁰ Investigate Incidents.

showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions":³¹

Blast radius analysis

Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions.

① Note

Blast radius analysis extends and replaces Attack path analysis.

The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it's integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact.

As another example, Microsoft's documentation explains that "Defender's correlation engine" correlates incidents and alerts based on elements such as "Entities," which are "assets like users, devices, mailboxes, and others," based in part on "continu[ing] to detect commonalities and relationships":³²

³¹ *Id.*

³² *Alert Correlation.*

	Incident correlation and merging The Defender portal's correlation activities don't stop when incidents are created. Defender continues to detect commonalities and relationships between incidents and alerts across incidents. When multiple incidents are determined to be sufficiently alike, Defender merges the incidents into a single incident. Criteria for merging incidents Defender's correlation engine merges incidents when it recognizes common elements between alerts in separate incidents, based on its deep knowledge of the data and the attack behavior. Some of these elements include: • Entities—assets like users, devices, mailboxes, and others • Artifacts—files, processes, email senders, and others • Time frames • Sequences of events that point to multistage attacks—for example, a malicious email click event that follows closely on a phishing email detection. Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the modified representation of the first graph comprises a representation of a first node corresponding to the first entity and a representation of a first edge corresponding to the first	The software instructions of the Accused '627 Defender Products modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the modified representation comprises a representation of a first node corresponding to the first entity and a representation of a first edge corresponding to the first relationship, wherein the first node is not of the first plurality of nodes and the first edge is not of the first plurality of edges. For example, as Defender XDR discovers new entities and new relationships, it updates its graph representations so that these new entities and relationships are reflected in the user interface and in Defender XDR's analyses, such as "interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird's eye view of a

relationship, wherein the first node is not of the first plurality of nodes and the first edge is not of the first plurality of edges,

possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly”:³³

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird's eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

For example, “the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms,” such as “[l]ists of all the involved and impacted users, devices, and other resources,” a “visual representation of how all the players in the story interact, “[c]ollections of evidence supporting the attack story: bad actors' user accounts and device information and address, malicious files and processes, relevant threat intelligence, and so on”:³⁴

³³ *Understanding Graph Icons.*

³⁴ *Incidents and Alerts.*

Incidents and alerts in the Microsoft Defender portal

01/06/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

The Microsoft Defender portal brings together a unified set of security services to reduce your exposure to security threats, improve your organizational security posture, detect security threats, and investigate and respond to breaches. These services collect and produce signals that are displayed in the portal. The two main kinds of signals are:

Alerts: Signals that result from various threat detection activities. These signals indicate the occurrence of malicious or suspicious events in your environment.

Incidents: Containers that include collections of related alerts and tell the full story of an attack. The alerts in a single incident might come from all Microsoft security and compliance solutions, as well as from vast numbers of external solutions collected through Microsoft Sentinel and Microsoft Defender for Cloud.

Incidents for correlation and investigation

While you can investigate and mitigate the threats that individual alerts bring to your attention, by themselves these threats are isolated occurrences that don't tell you anything about a broader, complex attack story. You could search for, research, investigate, and correlate groups of alerts that belong together in a single attack story, but that will cost you lots of time, effort, and energy.

Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms:

- Timelines of alerts and the raw events on which they're based
- A list of the tactics that were used
- Lists of all the involved and impacted users, devices, and other resources
- A visual representation of how all the players in the story interact
- Logs of automatic investigation and response processes that Defender XDR initiated and completed
- Collections of evidence supporting the attack story: bad actors' user accounts and device information and address, malicious files and processes, relevant threat intelligence, and so on
- A textual summary of the attack story

Incidents also provide you with a framework for managing and documenting your investigations and threat response. For more information about incidents' functionality in this regard, see [Manage incidents in Microsoft Defender](#).

As another example, Microsoft's documentation explains that "Defender's correlation engine" correlates incidents and alerts based on elements such as entities, which are "assets like users, devices, mailboxes, and others," and it does so on a continuous basis:³⁵

Incident correlation and merging

The Defender portal's correlation activities don't stop when incidents are created. Defender continues to detect commonalities and relationships between incidents and alerts across incidents. When multiple incidents are determined to be sufficiently alike, Defender merges the incidents into a single incident.

Criteria for merging incidents

Defender's correlation engine merges incidents when it recognizes common elements between alerts in separate incidents, based on its deep knowledge of the data and the attack behavior. Some of these elements include:

- Entities—assets like users, devices, mailboxes, and others
- Artifacts—files, processes, email senders, and others
- Time frames
- Sequences of events that point to multistage attacks—for example, a malicious email click event that follows closely on a phishing email detection.

As another example, Microsoft's documentation explains that the "graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes":³⁶


³⁵ *Alert Correlation.*

³⁶ *Investigate Incidents.*


The graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes.

From the graph, you can:

- Play the alerts and the nodes on the graph as they occurred over time to understand the chronology of the attack.



- Open an entity pane, allowing you to review the entity details and act on remediation actions, such as deleting a file or isolating a device.



- Highlight the alerts based on the entity to which they are related.
- Hunt for entity information of a device, file, IP address, URL, user, email, mailbox, or cloud resource.

	Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
identify, based on the modified representation of the first graph, an attack path that could be involved in an attack involving the first entity, wherein identifying the attack path comprises:	The software instructions of the Accused '627 Defender Products identify, based on the modified representation of the first graph, an attack path that could be involved in an attack involving the first entity. For example, Defender XDR “use[s] interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly”:³⁷ Understanding graphs and visualizations in Microsoft Defender 09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal Microsoft Defender use interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly. The graphs generated in the Defender portal are composed of nodes and edges. This article enumerates and defines the commonly used icons for graph these elements. As another example, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph

³⁷ *Understanding Graph Icons.*

showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions."³⁸ As explained in Microsoft's documentation, blast radius analysis is "built on the Microsoft Sentinel data lake and graph infrastructure":

Blast radius analysis

Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions.

① Note

Blast radius analysis extends and replaces Attack path analysis.

The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it's integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact.

When "view[ing] the blast radius of a single node," a "new graph view loads showing the 8 top-rated attack paths" that "shows the potential path from the entry point to this target."³⁹


³⁸ *Investigate Incidents.*

³⁹ *Id.*

View blast radius graphs

After selecting an incident from the list in the **Incidents** page, a graph view is displayed showing the entities and assets involved in the incident.

Select a node to open the context menu, then select **View blast radius**. To view the blast radius of a single node in a group, use the **ungroup** toggle above the grid to present all nodes.



A new graph view loads showing the 8 top-rated attack paths. A full list of the paths is visible on the right side panel when selecting **View full blast radius list** above the graph. From the list of reachable targets, you can further explore the path by selecting one of the listed targets. The right panel shows the potential path from the entry point to this target. Some nodes may not have paths associated with them.

	Incidents > Multi-stage incident involving Privilege escalation on multiple endpoints reported by multiple sources Multi-stage incident involving Privilege escalation on multiple endpoints reported by multiple sources High Active baat15@ Critical asset VP Incident linked to case: 1116 - [Preview] Containers running in Azure should have vulnerability findings resolved. Learn more about cases Assets (588) Attack story Alerts (196) Activities Assets (588) Investigations (0) Evidence and Response (646) Summary Similar incidents (0) Alerts Play attack story Unpin all Close all Dec 13, 2024 5:10 PM Resolved Purview IRM (72a04e39a) Project Obsidian Ismat Bekarevich (DPS) Mar 7, 2025 12:49 PM Resolved Purview IRM (7943dd7f8) Data leaks quick policy - 1/4/2024 Ismat Bekarevich (DPS) Mar 21, 2025 8:15 AM Resolved Rare subscription-level operations in Azure 2 Users Mar 21, 2025 8:15 AM Resolved Suspicious granting of permissions to an account Nabil Hafsa (DFC) Mar 21, 2025 8:15 AM Resolved Suspicious Resource deployment 2 Users Mar 21, 2025 8:18 AM Resolved Privileged Accounts - Sign in Failure Spikes 4 Users Mar 22, 2025 8:15 AM Resolved Rare subscription-level operations in Azure Jai Castellanos (DPS) Mar 22, 2025 8:15 AM Resolved Suspicious granting of permissions to an account Mar 22, 2025 8:18 AM Resolved Privileged Accounts - Sign in Failure Spikes 2 Users Mar 23, 2025 8:15 AM Resolved Rare subscription-level operations in Azure Incident graph Layout View full blast radius list (10 paths) Communication Association Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element. <tr><td data-bbox="151 1284 592 1396">identifying a second entity that can be reached using the first entity, wherein the second</td><td data-bbox="592 1284 1942 1396">Identifying the attack path of the software instructions of the Accused '627 Defender Products comprises identifying a second entity that can be reached using the first entity, wherein the second</td></tr>	identifying a second entity that can be reached using the first entity, wherein the second	Identifying the attack path of the software instructions of the Accused '627 Defender Products comprises identifying a second entity that can be reached using the first entity, wherein the second
identifying a second entity that can be reached using the first entity, wherein the second	Identifying the attack path of the software instructions of the Accused '627 Defender Products comprises identifying a second entity that can be reached using the first entity, wherein the second		

entity corresponds to a second node, and the second node is related by one or more edges to the first node corresponding to the first entity in the modified representation of the first graph; and,

entity corresponds to a second node, and the second node is related by one or more edges to the first node corresponding to the first entity in the modified representation of the first graph.

For example, Defender XDR “use[s] interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly.”⁴⁰

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

As another example, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions”:⁴¹

⁴⁰ *Understanding Graph Icons.*

⁴¹ *Investigate Incidents.*

Blast radius analysis

Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions.

🕒 Note

Blast radius analysis extends and replaces Attack path analysis.

The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it's integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact.


When “view[ing] the blast radius of a single node,” a “new graph view loads showing the 8 top-rated attack paths” that “shows the potential path from the entry point to this target”.⁴²

⁴² *Id.*


View blast radius graphs

After selecting an incident from the list in the **Incidents** page, a graph view is displayed showing the entities and assets involved in the incident.

Select a node to open the context menu, then select **View blast radius**. To view the blast radius of a single node in a group, use the **ungroup** toggle above the grid to present all nodes.



A new graph view loads showing the 8 top-rated attack paths. A full list of the paths is visible on the right side panel when selecting **View full blast radius list** above the graph. From the list of reachable targets, you can further explore the path by selecting one of the listed targets. The right panel shows the potential path from the entry point to this target. Some nodes may not have paths associated with them.

	 Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
identifying a third entity that can be reached using the second entity, wherein the third entity	Identifying the attack path of the software instructions of the Accused ’627 Defender Products comprises identifying a third entity that can be reached using the second entity, wherein the third entity

corresponds to a third node, and the third node is related by one or more edges to the second node in the modified representation of the first graph; and

corresponds to a third note, and the third node is related by one or more edges to the second node in the modified representation of the first graph.

For example, Defender XDR “use[s] interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly”:⁴³

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

As another example, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions.”⁴⁴

⁴³ *Understanding Graph Icons.*

⁴⁴ *Investigate Incidents.*

Blast radius analysis

Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions.

🕒 Note

Blast radius analysis extends and replaces Attack path analysis.

The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it's integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact.


When “view[ing] the blast radius of a single node,” a “new graph view loads showing the 8 top-rated attack paths” that “shows the potential path from the entry point to this target”.⁴⁵

⁴⁵ *Id.*


View blast radius graphs

After selecting an incident from the list in the **Incidents** page, a graph view is displayed showing the entities and assets involved in the incident.

Select a node to open the context menu, then select **View blast radius**. To view the blast radius of a single node in a group, use the **ungroup** toggle above the grid to present all nodes.



A new graph view loads showing the 8 top-rated attack paths. A full list of the paths is visible on the right side panel when selecting **View full blast radius list** above the graph. From the list of reachable targets, you can further explore the path by selecting one of the listed targets. The right panel shows the potential path from the entry point to this target. Some nodes may not have paths associated with them.



Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

generate a report comprising an identification of the first entity and at least one of the second entity and the third entity.

Defender XDR generates a report comprising an identification of the first entity and at least one of the second entity and the third entity.⁴⁶

For example, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions”:⁴⁷

Blast radius analysis

Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions.

① Note

Blast radius analysis extends and replaces Attack path analysis.

The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it’s integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact.

When “view[ing] the blast radius of a single node,” a “new graph view loads showing the 8 top-rated attack paths” that “shows the potential path from the entry point to this target.”⁴⁸

⁴⁶ *Id.*


⁴⁷ *Id.*

⁴⁸ *Id.*

View blast radius graphs

After selecting an incident from the list in the **Incidents** page, a graph view is displayed showing the entities and assets involved in the incident.


Select a node to open the context menu, then select **View blast radius**. To view the blast radius of a single node in a group, use the **ungroup** toggle above the grid to present all nodes.



A new graph view loads showing the 8 top-rated attack paths. A full list of the paths is visible on the right side panel when selecting **View full blast radius list** above the graph. From the list of reachable targets, you can further explore the path by selecting one of the listed targets. The right panel shows the potential path from the entry point to this target. Some nodes may not have paths associated with them.

	As another example, “[y]ou can view an incident's details on the right pane of an incident page. The incident details include incident assignment, ID, classification, categories, and first and last activity date and time. It also includes a description of the incident, impacted assets, active alerts, and where applicable, the related threats, recommendations, and disruption summary and impact. Here's an example of the incident details where the incident description is highlighted”: ⁴⁹
--	--

⁴⁹ *Id.*



Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

II. Claim 2

The computer system of claim 1, wherein identifying the attack path further comprises:	See above for an analysis of Claim 1.
identifying a first plurality of event flows that include a first anomalous event associated with the first entity.	Identifying the attack path of Claim 1 further comprises identifying a first plurality of event flows that include a first anomalous event associated with the first entity. For example, Defender's Attack Stories identifies a plurality of event flows that include a first anomalous event. In the example screenshot below, an anomalous event (for example, a suspicious sequence of exploration activities) is associated with a user and a workstation:⁵⁰

⁵⁰ *Investigate Incidents* (embedded *Attack stories* video at 0:40).

A screenshot of a security alert interface. At the top, it says 'Alerts' with a back arrow. Below that, it shows '0/26 Active alerts' with 'Unpin all' and 'Show all' options. There are three alerts listed: 1. A red dot icon, timestamp 'Mar 8, 2023 10:09 PM', status 'Resolved', title 'Suspicious URL clicked', source 'workstation6.seccxp.ninja', and user 'kdickens'. 2. An orange dot icon, timestamp 'Mar 8, 2023 10:09 PM', status 'Resolved', title 'User accessed link in ZAP-quarantined email', source 'kdickens'. 3. A yellow dot icon, timestamp 'Mar 8, 2023 10:10 PM', status 'Resolved', title 'Suspicious sequence of exploration activities', source 'workstation6.seccxp.ninja', and user 'Karla Dickens'. Each alert has 'Pin' and 'Unpin' icons on the right.

III. Claim 3

The system of claim 2, wherein identifying the attack path further comprises:	See above for an analysis of Claim 2.
identifying a point of origin for the first anomalous event based on the first plurality of event flows.	Identifying the attack path of Claim 2 further comprises identifying a point of origin for the first anomalous event based on the first plurality of event flows. For example, Defender's Attack Stories identifies a point of origin for the anomalous event. In the example screenshot below, the suspicious sequence of exploration activities is traced back to clicking a suspicious URL:⁵¹

⁵¹ *Id.*

A screenshot of a security alert interface. At the top, there's a header 'Alerts' with a back arrow. Below it, a status bar shows '0/26 Active alerts' and two buttons: 'Unpin all' and 'Show all'. The main area lists three alerts. Each alert has a colored circle (red, orange, yellow) on the left. The first alert is 'Suspicious URL clicked' at 'Mar 8, 2023 10:09 PM', status 'Resolved', from 'workstation6.seccxp.ninja' by 'kdickens'. The second alert is 'User accessed link in ZAP-quarantined email' at 'Mar 8, 2023 10:09 PM', status 'Resolved', from 'kdickens'. The third alert is 'Suspicious sequence of exploration activities' at 'Mar 8, 2023 10:10 PM', status 'Resolved', from 'workstation6.seccxp.ninja' by 'Karla Dickens'. Each alert has two blue icons (a pin and a magnifying glass) on the right.

IV. Claim 5

The computer system of claim 1	See above for an analysis of Claim 1.
wherein the modified representation of the first graph comprises a representation of a node corresponding to the first entity, wherein the first entity is identified based on active reconnaissance results.	The modified representation of the first graph of Claim 1 comprises a representation a node corresponding to the first entity, wherein the first entity is identified based on active reconnaissance results. For example, Defender XDR includes “Standard discovery,” a “device discovery capability that helps you find unmanaged devices connected to your corporate network.” Standard discovery “uses smart, active probing to discover additional information about observed devices to enrich existing device information.”⁵² For example, Standard discovery “uses various PowerShell scripts to actively probe devices in the network. Those PowerShell scripts are Microsoft signed and are executed from the following location: C:\ProgramData\Microsoft\Windows Defender Advanced Threat Protection\Downloads*.ps. For example, C:\ProgramData\Microsoft\Windows Defender Advanced Threat Protection\Downloads\UnicastScannerV1.1.0.ps1.”⁵³ Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

⁵² *Device Discovery*.

⁵³ Microsoft, *Configure device discovery in Defender for Endpoint*, available at <https://learn.microsoft.com/en-us/defender-endpoint/configure-device-discovery>.

V. Claim 6

The computer system of claim 1	See above for an analysis of Claim 1.
wherein the modified representation of the first graph comprises a representation of a node corresponding to the first entity, wherein the first entity is identified based on passive reconnaissance results.	The modified representation of the first graph of Claim 1 comprises a representation a node corresponding to the first entity, wherein the first entity is identified based on passive reconnaissance results. For example, Defender XDR includes “Basic discovery,” a “device discovery capability that helps you find unmanaged devices connected to your corporate network[.]”⁵⁴ Basic discovery “passively collect[s] events in your network and extract[s] device information from them,” us[ing] the SenseNDR.exe binary for passive network data collection.”⁵⁵ Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

⁵⁴ *Device Discovery*.

⁵⁵ *Id.*

VI. Claim 11

The computer system of claim 1,	See above for an analysis of Claim 1.
wherein the computer system is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that determine whether an event, of the events relating to one or more entities of the first plurality of entities, is anomalous, wherein determining whether the event is anomalous comprises: determining that the event relates to the first entity, determining at least one behavior pattern associated with the first entity, and comparing the event to the at least one behavior pattern.	The software instructions of the Accused '627 Defender Products determine whether an event, of the events relating to one or more entities of the first plurality of entities, is anomalous, wherein determining whether the event is anomalous comprises determining that the event relates to the first entity, determining at least one behavior pattern associated with the first entity, and comparing the event to the at least one behavior pattern. For example, Defender XDR “employ[s] large-scale learning algorithms to establish the normal behavior of common processes within an organization and worldwide and watch for when these processes show anomalous behaviors. These anomalous behaviors often indicate that extraneous code was introduced and is running in an otherwise trusted process.”⁵⁶ As another example, Microsoft Defender for Cloud Apps “detects anomalous behavior such as impossible-travel, credential access, and unusual downloading, file sharing, or mail forwarding activity and reports these to your security team.”⁵⁷

⁵⁶ Microsoft, *Investigate and respond using Microsoft Defender XDR*, available at <https://learn.microsoft.com/en-us/defender-xdr/pilot-deploy-investigate-respond> [hereinafter *Investigate and Respond*].

⁵⁷ *Deploy Defender XDR*.

Microsoft Defender XDR and an example cyber security attack

This diagram shows a common cyber-attack and the components of Microsoft Defender XDR that help detect and remediate it.

Mitigated by	
Microsoft Defender for Office 365	Microsoft Defender for Endpoint
Microsoft Defender for Identity	Microsoft Defender for Cloud Apps

The cyber-attack starts with a phishing email that arrives at the Inbox of an employee in your organization, who unknowingly opens the email attachment. This attachment installs malware, which can lead to a chain of attack attempts that can result in the theft of sensitive data.

In the illustration:

- **Exchange Online Protection**, part of Microsoft Defender for Office 365, can detect the phishing email and use mail flow rules (also known as transport rules) to make certain it never arrives in a user's Inbox.
- **Defender for Office 365** uses Safe Attachments to test the attachment and determine that it's harmful, so the mail that arrives either isn't actionable by the user, or policies prevent the mail from arriving at all.
- **Defender for Endpoint** detects device and network vulnerabilities that might otherwise be exploited for devices managed by your organization.
- **Defender for Identity** takes note of sudden on-premises user account changes like privilege escalation or high-risk lateral movement. It also reports on easily exploited identity issues like unconstrained Kerberos delegation, for correction by your security team.
- **Microsoft Defender for Cloud Apps** detects anomalous behavior such as impossible-travel, credential access, and unusual downloading, file sharing, or mail forwarding activity and reports these to your security team.

Defender for Cloud Apps “combin[es] multiple detection methods, including anomaly, behavioral analytics (UEBA), and rule-based activity detections, to provide a broad view of how your users use

	apps in your environment.”⁵⁸ Defender for Cloud Apps’ detection methods determine behavioral patterns associated with entities and compare events related to such entities to behavioral patterns. For example, Defender for Cloud Apps uses “user and entity behavioral analytics (UEBA) and machine learning (ML)” to “target[] numerous behavioral anomalies across your users and the machines and devices connected to your network.”⁵⁹ As another example, Defender for Cloud apps “[d]etects multiple file download activities in a single session with respect to the baseline learned,” and “[d]etects multiple administrative activities in a single session with respect to the baseline learned[.]”⁶⁰
--	---

⁵⁸ Microsoft, *Tutorial: Detect suspicious user activity with behavioral analytics (UEBA)*, available at <https://learn.microsoft.com/en-us/defender-cloud-apps/tutorial-suspicious-activity> [hereinafter *Detect Suspicious User Activity*].

⁵⁹ Microsoft, *Create Defender for Cloud Apps anomaly detection policies*, available at <https://learn.microsoft.com/en-us/defender-cloud-apps/anomaly-detection-policy> [hereinafter *Anomaly Detection Policies*].

⁶⁰ *Detect Suspicious User Activity*.

Phase 2: Tune anomaly detection policies

Several built-in anomaly detection policies are available in Defender for Cloud Apps that are preconfigured for common security use cases. You should take some time to familiarize yourself with the more popular detections, such as:

- **Impossible travel**
Activities from the same user in different locations within a period that is shorter than the expected travel time between the two locations.
- **Activity from infrequent country**
Activity from a location that wasn't recently or never visited by the user.
- **Malware detection**
Scans files in your cloud apps and runs suspicious files through Microsoft's threat intelligence engine to determine whether they're associated with known malware.
- **Ransomware activity**
File uploads to the cloud that might be infected with ransomware.
- **Activity from suspicious IP addresses**
Activity from an IP address that has been identified as risky by Microsoft Threat Intelligence.
- **Suspicious inbox forwarding**
Detects suspicious inbox forwarding rules set on a user's inbox.
- **Unusual multiple file download activities**
Detects multiple file download activities in a single session with respect to the baseline learned, which could indicate an attempted breach.
- **Unusual administrative activities**
Detects multiple administrative activities in a single session with respect to the baseline learned, which could indicate an attempted breach.

As another example, the “impossible travel” detection “uses a machine-learning algorithm” that “learns a new user's activity pattern” and “identifies unusual and impossible user activity between two locations.” Similarly, the “infrequent country” detection “stores information about previous locations used by the user. An alert is triggered when an activity occurs from a location that wasn't recently or never visited by the user.”⁶¹

As another example, the “unusual activities (by user)” detection “look for activities within a single session with respect to the baseline learned, which could indicate on a breach attempt. These

⁶¹ *Anomaly Detection Policies.*

	detections leverage a machine-learning algorithm that profiles the users log on pattern and reduces false positives. These detections are part of the heuristic anomaly detection engine that profiles your environment and triggers alerts with respect to a baseline that was learned on your organization's activity.”⁶² Unusual activities (by user) These detections identify users who perform: • Unusual multiple file download activities • Unusual file share activities • Unusual file deletion activities • Unusual impersonated activities • Unusual administrative activities • Unusual Power BI report sharing activities (preview) • Unusual multiple VM creation activities (preview) • Unusual multiple storage deletion activities (preview) • Unusual region for cloud resource (preview) Note As part of ongoing improvements to Defender for Cloud Apps alert threat protection capabilities, the policy with the title “Suspicious file access activity (by user)” has been disabled, migrated to the new dynamic model and renamed to Suspicious file access indicative of lateral movement and Suspicious file access from untrusted ISP and user agent with malicious IP indicator. If you previously configured governance actions or email notifications for this policy, you can re-enable it at any time in the Microsoft Defender portal > Cloud Apps > Policy management page. These policies look for activities within a single session with respect to the baseline learned, which could indicate on a breach attempt. These detections leverage a machine-learning algorithm that profiles the users log on pattern and reduces false positives. These detections are part of the heuristic anomaly detection engine that profiles your environment and triggers alerts with respect to a baseline that was learned on your organization's activity. Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents.
--	--

⁶² *Id.*

	For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
--	--

VII. Claim 12

The computer system of claim 11,	See above for an analysis of Claim 11.
wherein comparing the event to the at least one behavior pattern comprises using a threshold.	Comparing the event to the at least one behavior pattern of Claim 11 comprises using a threshold. For example, the Microsoft Defender for Cloud Apps component of Defender XDR uses tuning and thresholds to determine whether to flag an event as anomalous. Microsoft Defender for Cloud Apps allows users to set “dynamic thresholds” for anomaly detection.⁶³ Next, you want to tune your policies. The following policies can be fine-tuned by setting filters, dynamic thresholds (UEBA) to help train their detection models, and suppressions to reduce common false positive detections: • Anomaly detection • Cloud discovery anomaly detection • Rule-based activity detection As another example, the “impossible travel” detection discussed above contains a “sensitivity slider” to “determine[] the level of suppressions applied to anomalous behavior before triggering an impossible travel alert[,]” offering “Low,” “Medium,” and “High” sensitivity levels.⁶⁴

⁶³ *Detect Suspicious User Activity.*

⁶⁴ *Id.*

3. **Tune sensitivity of impossible travel** Configure the **sensitivity slider** that determines the level of suppressions applied to anomalous behavior before triggering an impossible travel alert. For example, organizations interested in high fidelity should consider increasing the sensitivity level. On the other hand, if your organization has many users that travel, consider lowering the sensitivity level to suppress activities from a user's common locations learned from previous activities. You can choose from the following sensitivity levels:

- **Low:** System, tenant, and user suppressions
- **Medium:** System and user suppressions
- **High:** Only system suppressions

As another example, Defender for Cloud Apps offers the ability to tune “the volume of activity required before the detection raises an alert.”⁶⁵

Phase 4: Tune rule-based detection (activity) policies

Rule-based detection policies give you the ability to complement anomaly detection policies with organization-specific requirements. We recommend creating rules-based policies using one of our Activity policy templates (go to **Control > Templates** and set the **Type** filter to **Activity policy**) and then **configuring them** to detect behaviors that aren't normal for your environment. For example, for some organization that don't have any presence in a particular country/region, it may make sense to create a policy that detects the anomalous activities from that country/region and alert on them. For others, who have large branches in that country/region, activities from that country/region would be normal and it wouldn't make sense to detect such activities.

1. Tune activity volume

Choose the volume of activity required before the detection raises an alert. Using our **country/region** example, if you have no presence in a country/region, even a single activity is significant and warrants an alert. However, a single sign-in failure could be human error and only of interest if there are many failures in a short period.

2. Tune activity filters

Set the filters you require to detect the type of activity you want to alert on. For example, to detect activity from a country/region, use the **Location** parameter.

3. Tune alerts

To prevent alert fatigue, set the **daily alert limit**.

Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be

⁶⁵ *Id.*

	insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
--	--

VIII. Claim 13

The system of claim 11, wherein comparing the event to the at least one behavior pattern further comprises:	See above for an analysis of Claim 11.
determining that the first entity is associated with a resource, wherein the resource is a sensitive resource.	Comparing the event to the at least one behavior pattern of Claim 11 further comprises determining that the first entity is associated with a resource, wherein the resource is a sensitive resource. For example, Defender XDR's blast radius analysis "generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions":⁶⁶ Blast radius analysis Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions. Blast radius analysis also "helps analysts better assess the risk to highly regarded targets, and understand the business impact."⁶⁷

⁶⁶ *Investigate Incidents.*

⁶⁷ *Id.*

As another example, Microsoft's documentation explains that a node can have a "[c]ritical asset" indicator associated with it, which "[i]ndicates that an entity is classified as business-critical or valuable[.]"⁶⁸


A node might also have any of the following indicators around it:

- **Critical asset** - Indicates that an entity is classified as business-critical or valuable, as identified in the *critical asset management* in Microsoft Security Exposure Management. This indicator appears as a golden crown 👑. The nodes representing critical assets also have a golden halo surrounding them.
- **Vulnerability** - Indicates that at least one vulnerability was detected on the entity. This indicator appears as a red bug 🐛.
- **Explore connected assets** - Indicates that the node can expand the hunting graph further beyond the initial results. Expanding the graph lets you explore other relationships the selected entity has with the other ones. This indicator appears as a blue plus sign ➕.
- **Discovery source** - Indicates the entity's data source. This indicator appears as the icon of the Defender product protecting the entity in blue (for example, 🛡️ for Microsoft Defender for Endpoint, or ☁️ for Microsoft Defender for Cloud).

In the example from Microsoft's documentation shown below, multiple nodes are identified as sensitive resources.⁶⁹

⁶⁸ *Understanding Graph Icons.*

⁶⁹ *Investigate Incidents.*



Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

IX. Claim 14

The computer system of claim 1,	See above for an analysis of Claim 1.
wherein the computer system comprises a plurality of physical computing machines.	The computer system of Claim 1 comprises a plurality of physical computing machines. For example, Defender XDR uses Microsoft Azure’s distributed computing cluster that consists of numerous physical computing machines in numerous geographic regions.⁷⁰ These data centers “house[] thousands of powerful computers, or ‘servers.’”⁷¹  Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

⁷⁰ *Data in XDR.*

⁷¹ *Microsoft Datacenters.*

X. Claim 16

The computer system of claim 1,	See above for an analysis of Claim 1.
wherein at least one entity of the first plurality of entities is at least one of a user, a place, a device, a resource, a group, or a service.	At least one entity of the first plurality of entities of claim 1 is at least one of a user, a place, a device, a resource, a group, or a service. For example, Microsoft’s documentation explains that a node of the graphs used by Defender XDR “pertains to an entity in your environment (for example, a device, user account, or IP address, among others)”:⁷²

⁷² *Understanding Graph Icons.*

Nodes

A **node** pertains to an entity in your environment (for example, a device, user account, or IP address, among others). Defender portal graphs usually depict nodes as any of the following circular icons:

Expand table

Icon	Node type	Entity type examples
	General	App service plan
	Compute	Device, virtual machine, Microsoft Azure Logic App
	Networking	Interface, public IP address, network security group
	Data	SQL data store, Azure Monitor Log Analytics workspace, storage account, Azure Event Hubs
	Containers	Kubernetes cluster
	Keys & secrets	Key vault
	DevOps	Azure DevOps repositories
	APIs	Cloud applications
	Identity & access	User account, Microsoft Entra ID service principal
	IoT	
	Certificate	
	IP address	
	Subscriptions	

	Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
--	---

XI. Claim 18

A computer system comprising:

a hardware memory, wherein the computer system is configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that:

The Accused '627 Defender Products include a computer system comprising a hardware memory, wherein the computer system is configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that perform as discussed below.

For example, Defender XDR “is a cloud-based, unified, pre- and post-breach enterprise defense suite.”⁷³

Pilot and deploy Microsoft Defender XDR

Applies to:

- Microsoft Defender XDR

This series of articles steps you through the entire process of piloting the components of Microsoft Defender XDR in your production tenant so you can evaluate their features and capabilities and then completing the deployment across your organization.

An eXtended detection and response (XDR) solution is a step forward in cyber security because it takes the threat data from systems that were once isolated and unifies them so that you can see patterns and act on suspected cyberattacks faster.

Microsoft Defender XDR:

- Is an XDR solution that combines the information on cyberattacks for identities, endpoints, email, and cloud apps in one place. It leverages artificial intelligence (AI) and automation to automatically stop some types of attacks and remediate affected assets to a safe state.
- Is a cloud-based, unified, pre- and post-breach enterprise defense suite. It coordinates prevention, detection, investigation, and response across identities, endpoints, email, cloud apps, and their data.

Defender XDR “operates in Microsoft Azure data centers”.⁷⁴

⁷³ *Deploy Defender XDR.*

⁷⁴ *Data in XDR.*

Data storage location

Microsoft Defender XDR operates in Microsoft Azure data centers in the following geographical regions:

- **European Union:** North Europe and West Europe
- **United Kingdom:** UK South and UK West
- **United States:** East US 2 and Central US
- **Australia:** Australia East and Australia Southeast
- **Switzerland:** Switzerland North and Switzerland West
- **India:** Central India and South India
- **UAE:** UAE North and UAE Central

These data centers “house[] thousands of powerful computers, or ‘servers,’” examples of which can be seen in the following photograph reproduced from Microsoft’s literature:⁷⁵



⁷⁵ *Microsoft Datacenters.*

store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges,

The software instructions of the Accused '627 Defender Products store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges, wherein the first graph is a directed graph.

For example, Defender XDR “use[s] interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. . . . The graphs generated in the Defender portal are composed of nodes and edges”:⁷⁶

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

As another example, Microsoft’s documentation explains that the “graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes”:⁷⁷


⁷⁶ *Understanding Graph Icons.*

⁷⁷ *Investigate Incidents.*


The graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes.

From the graph, you can:

- Play the alerts and the nodes on the graph as they occurred over time to understand the chronology of the attack.



- Open an entity pane, allowing you to review the entity details and act on remediation actions, such as deleting a file or isolating a device.



- Highlight the alerts based on the entity to which they are related.
- Hunt for entity information of a device, file, IP address, URL, user, email, mailbox, or cloud resource.

	As another example, Microsoft’s documentation explains that a node of the graphs used by Defender XDR “pertains to an entity in your environment (for example, a device, user account, or IP address, among others)”: ⁷⁸
--	---

⁷⁸ *Understanding Graph Icons.*

Nodes

A **node** pertains to an entity in your environment (for example, a device, user account, or IP address, among others). Defender portal graphs usually depict nodes as any of the following circular icons:

Expand table

Icon	Node type	Entity type examples
	General	App service plan
	Compute	Device, virtual machine, Microsoft Azure Logic App
	Networking	Interface, public IP address, network security group
	Data	SQL data store, Azure Monitor Log Analytics workspace, storage account, Azure Event Hubs
	Containers	Kubernetes cluster
	Keys & secrets	Key vault
	DevOps	Azure DevOps repositories
	APIs	Cloud applications
	Identity & access	User account, Microsoft Entra ID service principal
	IoT	
	Certificate	
	IP address	
	Subscriptions	

	As another example, Microsoft’s documentation explains that nodes can be represented as an ExposureGraphNodes table of “organizational entities and their properties,” which “include entities like devices, identities, user groups, and cloud assets such as virtual machines (VMs), storage, and containers. Each node corresponds to an individual entity and encapsulates information about its characteristics, attributes, and security related insights within the organizational structure”:⁷⁹ ExposureGraphNode 06/20/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal Applies to: • Microsoft Defender XDR • Microsoft Security Exposure Management (public preview) Important Some information relates to prereleased product which may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided here. The <code>ExposureGraphNode</code>s table in the advanced hunting schema contains organizational entities and their properties. These include entities like devices, identities, user groups, and cloud assets such as virtual machines (VMs), storage, and containers. Each node corresponds to an individual entity and encapsulates information about its characteristics, attributes, and security related insights within the organizational structure. Use this reference to construct queries that return information from this table. As another example, edges can be represented as an ExposureGraphEdges table of “relationships between entities and assets in the enterprise exposure graph”:⁸⁰
--	---

⁷⁹ *ExposureGraphNode*s.

⁸⁰ *ExposureGraphEdge*s.

	ExposureGraphEdges Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal Applies to: • Microsoft Defender XDR • Microsoft Security Exposure Management (public preview) Important Some information relates to prereleased product which may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided here. The <code>ExposureGraphEdges</code> table in the advanced hunting schema provides visibility into relationships between entities and assets in the enterprise exposure graph. This visibility can help uncover critical organizational assets and explore entity relationships and attack paths. Use this reference to construct queries that return information from this table.
--	--

As another example, Defender XDR includes a “hunting graph,” which is “composed of nodes and edges to represent entities in your environment (for example, a device, user account, or IP address, among others) and their relationships or connection properties, respectively”:⁸¹

Hunting graph features

The interactive graphs generated in the hunting graph are composed of **nodes** and **edges** to represent entities in your environment (for example, a device, user account, or IP address, among others) and their relationships or connection properties, respectively. [Learn more about graphs and visualizations in Microsoft Defender](#)

The lower right-hand corner of the graph also has control buttons that let you **Zoom in** and **Zoom out**, and view the graph's **Layers**.




⁸¹ *Hunting Graph.*

	Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
wherein the first graph is a directed graph,	The first graph of the Accused '627 Defender Products is a directed graph. For example, <code>ExposureGraphEdges</code> table schema includes information about “source nodes” and “target nodes.”⁸² For example, Microsoft’s documentation explains that the edges of Defender XDR’s graphs “indicate the relationship or connection properties between two nodes” and include “directional arrows.”⁸³ As another example, as discussed above, Defender XDR includes a “hunting graph,” which is depicted by Microsoft’s documentation as a directed graph.⁸⁴

⁸² *ExposureGraphEdges*.

⁸³ *Understanding Graph Icons*.

⁸⁴ *Hunting Graph*.



As another example, Microsoft’s documentation explains that users must take caution “to identify and input the correct start and end entities, as the generated graph will be directional.”⁸⁵

Scenario	Description	Inputs
Paths between two entities	Provide two entities (nodes) to view the paths between them. Use this scenario if you want to discover if there’s a path leading from one entity to another.	• Start Entity • End Entity Note: Make sure to identify and input the correct start and end entities, as the generated graph will be directional.

⁸⁵ *Hunting Graph.*

	Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
wherein the first plurality of entities comprises a plurality of accounts and a plurality of resources, and	The first plurality of entities of the Accused '627 Defender Products comprises a plurality of accounts and a plurality of resources. For example, Microsoft’s documentation explains that Defender XDR stores graphs in which the nodes include accounts and resources: “[a] node pertains to an entity in your environment (for example, a device, user account, or IP address, among others.”⁸⁶ Examples of node types and their icon representations are reproduced below:

⁸⁶ *Understanding Graph Icons.*

Nodes

A **node** pertains to an entity in your environment (for example, a device, user account, or IP address, among others). Defender portal graphs usually depict nodes as any of the following circular icons:

Expand table

Icon	Node type	Entity type examples
	General	App service plan
	Compute	Device, virtual machine, Microsoft Azure Logic App
	Networking	Interface, public IP address, network security group
	Data	SQL data store, Azure Monitor Log Analytics workspace, storage account, Azure Event Hubs
	Containers	Kubernetes cluster
	Keys & secrets	Key vault
	DevOps	Azure DevOps repositories
	APIs	Cloud applications
	Identity & access	User account, Microsoft Entra ID service principal
	IoT	
	Certificate	
	IP address	
	Subscriptions	

	As another example, as discussed above, Defender XDR includes a “hunting graph,” which includes nodes corresponding to “entities in your environment (for example, a device, user account, or IP address, among others).”⁸⁷ Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
Wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities;	Each edge of the first plurality of edges of the Accused '627 Defender Products corresponds to a respective relationship between a respective pair of entities. For example, the edges that comprise Defender XDR’s graphs “indicate[] the relationship or connection properties between two nodes.”⁸⁸ Examples of edge types and their icon representations are reproduced below:

⁸⁷ *Hunting Graph.*

⁸⁸ *Understanding Graph Icons.*

	Edges An edge indicates the relationship or connection properties between two nodes. The Defender portal graphs depicts an edge as lines or directional arrows that might have the following icons: Expand table <table><tr><th>Icon</th><th>Edge type</th></tr><tr><td></td><td>Contains</td></tr><tr><td></td><td>Routes traffic to</td></tr><tr><td></td><td>Has permission to / Has role on</td></tr><tr><td></td><td>Can authenticate as / Can authenticate to</td></tr><tr><td></td><td>Pushes</td></tr><tr><td></td><td>Maintains</td></tr><tr><td></td><td>Application</td></tr><tr><td></td><td>Moves data to</td></tr><tr><td></td><td>Exposed to internet</td></tr><tr><td></td><td>Can interactive logon to / Can logon over the network to / Can remote interactive logon to</td></tr><tr><td></td><td>Runs on</td></tr><tr><td></td><td>Provisions</td></tr><tr><td></td><td>Identified as owner of</td></tr><tr><td></td><td>Member of</td></tr><tr><td></td><td>Is running</td></tr><tr><td></td><td>Generic / Affects</td></tr><tr><td></td><td>Created from / Used to create</td></tr></table>	Icon	Edge type		Contains		Routes traffic to		Has permission to / Has role on		Can authenticate as / Can authenticate to		Pushes		Maintains		Application		Moves data to		Exposed to internet		Can interactive logon to / Can logon over the network to / Can remote interactive logon to		Runs on		Provisions		Identified as owner of		Member of		Is running		Generic / Affects		Created from / Used to create
Icon	Edge type																																				
	Contains																																				
	Routes traffic to																																				
	Has permission to / Has role on																																				
	Can authenticate as / Can authenticate to																																				
	Pushes																																				
	Maintains																																				
	Application																																				
	Moves data to																																				
	Exposed to internet																																				
	Can interactive logon to / Can logon over the network to / Can remote interactive logon to																																				
	Runs on																																				
	Provisions																																				
	Identified as owner of																																				
	Member of																																				
	Is running																																				
	Generic / Affects																																				
	Created from / Used to create																																				

	As another example, as discussed above, Defender XDR includes a “hunting graph,” which includes “relationships or connection properties” between graph nodes.⁸⁹ As another example, Microsoft’s documentation explains that Defender XDR “provides visibility into relationships between entities and assets in the enterprise exposure graph”:⁹⁰ ExposureGraphEdges 06/20/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal Applies to: • Microsoft Defender XDR • Microsoft Security Exposure Management (public preview) Important Some information relates to prereleased product which may be substantially modified before it's commercially released. Microsoft makes no warranties, express or implied, with respect to the information provided here. The <code>ExposureGraphEdges</code> table in the advanced hunting schema provides visibility into relationships between entities and assets in the enterprise exposure graph. This visibility can help uncover critical organizational assets and explore entity relationships and attack paths. Use this reference to construct queries that return information from this table.
--	--

⁸⁹ *Hunting Graph.*

⁹⁰ *ExposureGraphEdges.*

	<table><tr><th>Column name</th><th>Data type</th><th>Description</th></tr><tr><td>EdgeId</td><td>string</td><td>Unique identifier for the relationship/edge</td></tr><tr><td>EdgeLabel</td><td>string</td><td>The edge label like "routes traffic to"</td></tr><tr><td>SourceNodeId</td><td>string</td><td>Node ID of the edge's source</td></tr><tr><td>SourceNodeName</td><td>string</td><td>Source node display name</td></tr><tr><td>SourceNodeLabel</td><td>string</td><td>Source node label</td></tr><tr><td>SourceNodeCategories</td><td>dynamic</td><td>Categories list of the source node in JSON format</td></tr><tr><td>TargetNodeId</td><td>string</td><td>Node ID of the edge's target</td></tr><tr><td>TargetNodeName</td><td>string</td><td>Display name of the target node</td></tr><tr><td>TargetNodeLabel</td><td>string</td><td>Target node label</td></tr><tr><td>TargetNodeCategories</td><td>dynamic</td><td>The categories list of the target node in JSON format</td></tr><tr><td>EdgeProperties</td><td>dynamic</td><td>Optional data relevant for the relationship between the nodes in JSON format</td></tr></table>	Column name	Data type	Description	EdgeId	string	Unique identifier for the relationship/edge	EdgeLabel	string	The edge label like "routes traffic to"	SourceNodeId	string	Node ID of the edge's source	SourceNodeName	string	Source node display name	SourceNodeLabel	string	Source node label	SourceNodeCategories	dynamic	Categories list of the source node in JSON format	TargetNodeId	string	Node ID of the edge's target	TargetNodeName	string	Display name of the target node	TargetNodeLabel	string	Target node label	TargetNodeCategories	dynamic	The categories list of the target node in JSON format	EdgeProperties	dynamic	Optional data relevant for the relationship between the nodes in JSON format
Column name	Data type	Description																																			
EdgeId	string	Unique identifier for the relationship/edge																																			
EdgeLabel	string	The edge label like "routes traffic to"																																			
SourceNodeId	string	Node ID of the edge's source																																			
SourceNodeName	string	Source node display name																																			
SourceNodeLabel	string	Source node label																																			
SourceNodeCategories	dynamic	Categories list of the source node in JSON format																																			
TargetNodeId	string	Node ID of the edge's target																																			
TargetNodeName	string	Display name of the target node																																			
TargetNodeLabel	string	Target node label																																			
TargetNodeCategories	dynamic	The categories list of the target node in JSON format																																			
EdgeProperties	dynamic	Optional data relevant for the relationship between the nodes in JSON format																																			
identify, based on the representation of the first graph, a first plurality of attack paths comprising a first entity of the first plurality of entities, wherein each attack path of the first plurality of attack paths targets a second entity of the first plurality of entities,	The software instructions of the Accused '627 Defender Products identify, based on the representation of the first graph, a first plurality of attack paths comprising a first entity of the first plurality of entities, wherein each attack path of the first plurality of attack paths targets a second entity of the first plurality of entities. For example, Defender XDR “use[s] interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird’s eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly”:⁹¹																																				

⁹¹ *Understanding Graph Icons.*

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird's eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

As another example, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions.”⁹² As explained in Microsoft’s documentation, blast radius analysis is “built on the Microsoft Sentinel data lake and graph infrastructure”:

⁹² *Investigate Incidents.*


	Blast radius analysis Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions. Note Blast radius analysis extends and replaces Attack path analysis. The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it's integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact. When “view[ing] the blast radius of a single node,” a “new graph view loads showing the 8 top-rated attack paths” that “shows the potential path from the entry point to this target.”⁹³
--	---

⁹³ *Id.*


View blast radius graphs

After selecting an incident from the list in the **Incidents** page, a graph view is displayed showing the entities and assets involved in the incident.

Select a node to open the context menu, then select **View blast radius**. To view the blast radius of a single node in a group, use the **ungroup** toggle above the grid to present all nodes.



A new graph view loads showing the 8 top-rated attack paths. A full list of the paths is visible on the right side panel when selecting **View full blast radius list** above the graph. From the list of reachable targets, you can further explore the path by selecting one of the listed targets. The right panel shows the potential path from the entry point to this target. Some nodes may not have paths associated with them.

	 Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
Receive streaming data comprising time-stamped data about events relating to one or	The software instructions of the Accused ’627 Defender Products receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities.

more entities of the first plurality of entities,

For example, Defender XDR “collect[s] . . . signals that are displayed in the portal.” Two kinds of signals include alerts, which Microsoft describes as “[s]ignals that result from various threat detection activities,” and incidents, which Microsoft describes as “[c]ontainers that include collections of related alerts and tell the full story of an attack:⁹⁴

Incidents and alerts in the Microsoft Defender portal

01/06/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

The Microsoft Defender portal brings together a unified set of security services to reduce your exposure to security threats, improve your organizational security posture, detect security threats, and investigate and respond to breaches. These services collect and produce signals that are displayed in the portal. The two main kinds of signals are:

Alerts: Signals that result from various threat detection activities. These signals indicate the occurrence of malicious or suspicious events in your environment.

Incidents: Containers that include collections of related alerts and tell the full story of an attack. The alerts in a single incident might come from all Microsoft security and compliance solutions, as well as from vast numbers of external solutions collected through Microsoft Sentinel and Microsoft Defender for Cloud.

Microsoft’s documentation explains that Defender “us[es] AI to continually monitor its telemetry sources”.⁹⁵

⁹⁴ *Incidents and Alerts.*

⁹⁵ *Id.*

Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms:

As another example, “[a]lerts in the Microsoft Defender portal come from many sources. These sources include the many services that are part of Microsoft Defender XDR, as well as other services with varying degrees of integration with the Microsoft Defender portal. For example, when Microsoft Sentinel is onboarded to the Microsoft Defender portal, the correlation engine in the Defender portal has access to all the raw data ingested by Microsoft Sentinel, which you can find in Defender’s Advanced hunting tables”:⁹⁶

⁹⁶ *Id.*

Alert sources and threat detection

Alerts in the Microsoft Defender portal come from many sources. These sources include the many services that are part of Microsoft Defender XDR, as well as other services with varying degrees of integration with the Microsoft Defender portal.

For example, when Microsoft Sentinel is [onboarded](#) to the Microsoft Defender portal, the correlation engine in the Defender portal has access to all the raw data ingested by Microsoft Sentinel, which you can find in Defender's **Advanced hunting** tables.

Microsoft Defender XDR itself also creates alerts. Defender XDR's unique correlation capabilities provide another layer of data analysis and threat detection for all the non-Microsoft solutions in your digital estate. These detections produce Defender XDR alerts, in addition to the alerts already provided by Microsoft Sentinel's analytics rules.

Within each of these sources, there are one or more threat detection mechanisms that produce alerts based on the rules defined in each mechanism.

For example, Microsoft Sentinel has at least four different engines that produce different types of alerts, each with its own rules.

As another example, Microsoft's documentation explains that "[a]lerts are signals that result from various threat detection activities. These signals are produced by the many security services that reside in the Microsoft Defender portal, and they indicate the occurrence of malicious or suspicious events in your environment".⁹⁷

⁹⁷ *Investigate Alerts.*

Investigate alerts in Microsoft Defender XDR

06/04/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

① **Note**

This article describes security alerts in Microsoft Defender XDR. However, you can use alert policies to send email notifications to yourself or other admins when users perform specific activities in Microsoft 365. For more information, see [Alert policies in the Microsoft Defender portal](#).

Alerts are signals that result from various threat detection activities. These signals are produced by the many security services that reside in the Microsoft Defender portal, and they indicate the occurrence of malicious or suspicious events in your environment.

These suspicious events are typically part of a broader attack story. In the Microsoft Defender portal, alerts represent individual pieces of evidence that Defender XDR correlates together to form [incidents](#). Incidents tell the whole attack story; however, analyzing alerts can be valuable when deeper analysis is required.


The **Alerts queue** shows the current set of alerts. You can view the entire alerts queue from **Incidents & alerts > Alerts** on the quick launch of the [Microsoft Defender portal](#). You can also see the alerts for each incident on the **incidents queue**, and on each individual incident's page, on the **Alerts** tab.

[illegible]

As another example, in Defender XDR, “[e]vent or activity data populates tables about alerts, security events, system events, and routine assessments. Advanced hunting receives this data almost immediately after the sensors that collect them successfully transmit them to the corresponding cloud services. For example, you can query event data from healthy sensors on workstations or domain controllers almost immediately after they’re made available on Microsoft Defender for Endpoint and Microsoft Defender for Identity. . . . Advanced hunting data uses the UTC (Universal Time

Coordinated) timezone. . . . Advanced hunting results are converted to the timezone set in Defender XDR.”⁹⁸

As another example, Microsoft’s documentation explains that Defender XDR includes “[a]ttack stories” with a “graph” that “shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went”:⁹⁹



Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents.

⁹⁸ Advanced Hunting.
⁹⁹ Investigate Incidents.

	For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
based on a first portion of the streaming data, identify a third entity that does not correspond to any of the first plurality of nodes, wherein the third entity is not of the first plurality of entities,	Based on a first portion of the streaming data, the software instructions of the Accused '627 Defender Products identify a third entity that does not correspond to any of the first plurality of nodes, wherein the third entity is not of the first plurality of entities. For example, Defender XDR includes a device inventory. Microsoft's documentation explains that "[d]uring the onboarding process, the Devices list is gradually populated with devices as they begin to report sensor data":¹⁰⁰ During the onboarding process, the Devices list is gradually populated with devices as they begin to report sensor data. Use this view to track your onboarded endpoints as they come online, or download the complete endpoint list as a CSV file for offline analysis. Further, Microsoft's documentation explains that Defender XDR performs "device discovery" by "collect[ing], prob[ing], or scan[ning] your network to discover unmanaged devices":¹⁰¹

¹⁰⁰ *Device Inventory.*

¹⁰¹ *Device Discovery.*

	Device discovery overview 05/08/2025 • Applies to: Microsoft Defender for Endpoint Plan 2 Protecting your environment requires taking inventory of the devices that are in your network. However, mapping devices in a network can often be expensive, challenging, and time-consuming. Microsoft Defender for Endpoint provides a device discovery capability that helps you find unmanaged devices connected to your corporate network without the need for extra appliances or cumbersome process changes. Device discovery uses onboarded endpoints, in your network to collect, probe, or scan your network to discover unmanaged devices. The device discovery capability allows you to discover: • Enterprise endpoints (workstations, servers, and mobile devices) that aren't yet onboarded to Defender for Endpoint • Network devices like routers and switches • IoT devices like printers and cameras For example, the “Device Inventory” interface includes a summary of devices discovered in the last 7 days:¹⁰²
--	---

¹⁰² *Id.*

	Devices discovered in the last 7 days Recently discovered devices in your organization: 20  View recently discovered devices Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
based on a second portion of the streaming data, wherein the second portion is not identical to the first portion, identify a first relationship between a pair of entities of the first plurality of entities that does not correspond	Based on a second portion of the streaming data, wherein the second portion is not identical to the first portion, the software instructions of the Accused '627 Defender Products identify a first relationship between a pair of entities of the first plurality of entities that does not correspond to any of the first plurality of edges. For example, Defender XDR “us[es] AI to continually monitor its telemetry sources” in order to “automatically aggregate and correlate related alerts”:¹⁰³

¹⁰³ *Incidents and Alerts.*

to any of the first plurality of edges,

Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms:

Further, Defender keeps track of “which onboarded device a discovered device was seen by,” allowing **SeenBy** queries:¹⁰⁴

By invoking the **SeenBy** function, in your advanced hunting query, you can get detail on which onboarded device a discovered device was seen by. This information can help determine the network location of each discovered device and subsequently, help to identify it in the network.

Similarly, Defender XDR performs “[b]last radius analysis,” which is “an advanced graph visualization integrated into incident investigation experience” that “generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions”:¹⁰⁵

¹⁰⁴ *Device Discovery.*

¹⁰⁵ *Investigate Incidents.*

	Blast radius analysis Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user's permissions. ① Note Blast radius analysis extends and replaces Attack path analysis. The blast radius graph provides a unique unified view of both prebreach and post-breach information on the incident page. During an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph. Because it's integrated into the incident graph, the blast radius graph helps security teams better understand the scope of the security incident quicker and enhance their defensive measures to reduce the likelihood of widespread damage. Blast radius analysis helps analysts better assess the risk to highly regarded targets, and understand the business impact. As another example, Microsoft's documentation explains that "Defender's correlation engine" correlates incidents and alerts based on elements such as "Entities," which are "assets like users, devices, mailboxes, and others," based in part on "continu[ing] to detect commonalities and relationships":¹⁰⁶
--	---

¹⁰⁶ Alert Correlation.

	Incident correlation and merging The Defender portal's correlation activities don't stop when incidents are created. Defender continues to detect commonalities and relationships between incidents and alerts across incidents. When multiple incidents are determined to be sufficiently alike, Defender merges the incidents into a single incident. Criteria for merging incidents Defender's correlation engine merges incidents when it recognizes common elements between alerts in separate incidents, based on its deep knowledge of the data and the attack behavior. Some of these elements include: • Entities—assets like users, devices, mailboxes, and others • Artifacts—files, processes, email senders, and others • Time frames • Sequences of events that point to multistage attacks—for example, a malicious email click event that follows closely on a phishing email detection. Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the modified representation of the first graph comprises a representation of a node corresponding to the third entity and an edge corresponding to the first relationship, wherein	The software instructions of the Accused '627 Defender Products modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the modified representation of the first graph comprises a representation of a node corresponding to the third entity and an edge corresponding to the first relationship, wherein the node is not of the first plurality of nodes and the edge is not of the first plurality of edges. For example, as Defender XDR discovers new entities and new relationships, it updates its graph representations so that these new entities and relationships are reflected in the user interface and in Defender XDR's analyses, such as "interactive graphs to visualize attack paths, blast radius, and relationships between entities in your environment. These visualizations provide a bird's eye view of

the node is not of the first plurality of nodes and the edge is not of the first plurality of edges, and

a possible threat or attack, letting you and your security operations (SOC) team to investigate and hunt them quickly”:¹⁰⁷

Understanding graphs and visualizations in Microsoft Defender

09/30/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

Microsoft Defender use interactive graphs to visualize attack paths, [blast radius](#), and relationships between entities in your environment. These visualizations provide a bird's eye view of a possible threat or attack, letting you and your security operations (SOC) team to investigate and [hunt](#) them quickly.

The graphs generated in the Defender portal are composed of [nodes](#) and [edges](#). This article enumerates and defines the commonly used icons for graph these elements.

For example, “the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms,” such as “[l]ists of all the involved and impacted users, devices, and other resources,” a “visual representation of how all the players in the story interact, “[c]ollections of evidence supporting the attack story: bad actors' user accounts and device information and address, malicious files and processes, relevant threat intelligence, and so on”:¹⁰⁸

¹⁰⁷ *Understanding Graph Icons.*

¹⁰⁸ *Incidents and Alerts.*

Incidents and alerts in the Microsoft Defender portal

01/06/2025 • Applies to: Microsoft Defender XDR, Microsoft Sentinel in the Microsoft Defender portal

The Microsoft Defender portal brings together a unified set of security services to reduce your exposure to security threats, improve your organizational security posture, detect security threats, and investigate and respond to breaches. These services collect and produce signals that are displayed in the portal. The two main kinds of signals are:

Alerts: Signals that result from various threat detection activities. These signals indicate the occurrence of malicious or suspicious events in your environment.

Incidents: Containers that include collections of related alerts and tell the full story of an attack. The alerts in a single incident might come from all Microsoft security and compliance solutions, as well as from vast numbers of external solutions collected through Microsoft Sentinel and Microsoft Defender for Cloud.

Incidents for correlation and investigation

While you can investigate and mitigate the threats that individual alerts bring to your attention, by themselves these threats are isolated occurrences that don't tell you anything about a broader, complex attack story. You could search for, research, investigate, and correlate groups of alerts that belong together in a single attack story, but that will cost you lots of time, effort, and energy.

Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms:

- Timelines of alerts and the raw events on which they're based
- A list of the tactics that were used
- Lists of all the involved and impacted users, devices, and other resources
- A visual representation of how all the players in the story interact
- Logs of automatic investigation and response processes that Defender XDR initiated and completed
- Collections of evidence supporting the attack story: bad actors' user accounts and device information and address, malicious files and processes, relevant threat intelligence, and so on
- A textual summary of the attack story

Incidents also provide you with a framework for managing and documenting your investigations and threat response. For more information about incidents' functionality in this regard, see [Manage incidents in Microsoft Defender](#).

As another example, Microsoft’s documentation explains that “Defender’s correlation engine” correlates incidents and alerts based on elements such as entities, which are “assets like users, devices, mailboxes, and others,” and it does so on a continuous basis.¹⁰⁹

Incident correlation and merging

The Defender portal's correlation activities don't stop when incidents are created. Defender continues to detect commonalities and relationships between incidents and alerts across incidents. When multiple incidents are determined to be sufficiently alike, Defender merges the incidents into a single incident.

Criteria for merging incidents

Defender's correlation engine merges incidents when it recognizes common elements between alerts in separate incidents, based on its deep knowledge of the data and the attack behavior. Some of these elements include:

- Entities—assets like users, devices, mailboxes, and others
- Artifacts—files, processes, email senders, and others
- Time frames
- Sequences of events that point to multistage attacks—for example, a malicious email click event that follows closely on a phishing email detection.

As another example, Microsoft’s documentation explains that the “graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes”:¹¹⁰


¹⁰⁹ *Alert Correlation.*

¹¹⁰ *Investigate Incidents.*


The graph shows the full scope of the attack, how the attack spread through your network over time, where it started, and how far the attacker went. It connects the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes.

From the graph, you can:

- Play the alerts and the nodes on the graph as they occurred over time to understand the chronology of the attack.



- Open an entity pane, allowing you to review the entity details and act on remediation actions, such as deleting a file or isolating a device.



- Highlight the alerts based on the entity to which they are related.
- Hunt for entity information of a device, file, IP address, URL, user, email, mailbox, or cloud resource.

	Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
identify, based on the modified representation of the first graph, a second plurality of attack paths comprising the first entity, wherein each attack path of the second plurality of attack paths targets the second entity, and wherein an attack path of the second plurality of attack paths comprises the third entity.	The software instructions of Defender XDR identify, based on the modified representation of the first graph, a second plurality of attack paths comprising the first entity, wherein each attack path of the second plurality of attack paths targets the second entity, and wherein an attack path of the second plurality of attack paths comprises the third entity. For example, Microsoft's documentation explains that Defender "us[es] AI to continually monitor its telemetry sources":¹¹¹ Instead, the correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents that represent these larger attack stories. Defender identifies multiple signals as belonging to the same attack story, using AI to continually monitor its telemetry sources and add more evidence to already open incidents. Incidents contain all the alerts deemed to be related to each other and to the overall attack story, and present the story in various forms: As another example, Microsoft's documentation explains that "Defender's correlation engine" correlates incidents and alerts based on elements such as "Entities," which are "assets like users, devices, mailboxes, and others," based in part on "continu[ing] to detect commonalities and relationships":¹¹²

¹¹¹ *Incidents and Alerts.*

¹¹² *Alert Correlation.*

Incident correlation and merging

The Defender portal's correlation activities don't stop when incidents are created. Defender continues to detect commonalities and relationships between incidents and alerts across incidents. When multiple incidents are determined to be sufficiently alike, Defender merges the incidents into a single incident.

Criteria for merging incidents

Defender's correlation engine merges incidents when it recognizes common elements between alerts in separate incidents, based on its deep knowledge of the data and the attack behavior. Some of these elements include:

- Entities—assets like users, devices, mailboxes, and others
- Artifacts—files, processes, email senders, and others
- Time frames
- Sequences of events that point to multistage attacks—for example, a malicious email click event that follows closely on a phishing email detection.

As Defender continually monitors the network and updates its graph, Microsoft's documentation explains that “[d]uring an incident investigation, analysts can see the current impact of a breach and the possible future impact in one consolidated graph,” including the “8 top-rated attack paths.”¹¹³

Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

¹¹³ *Investigate Incidents.*

XII. Claim 22

The computer system of claim 18,	See above for an analysis of Claim 18.
wherein a portion of the representation of the first graph is derived from reconnaissance data received by the computer system from a plurality of computer systems,	A portion of the representation of the first graph of Claim 18 is derived from reconnaissance data received by the computer system from a plurality of computer systems. For example, Defender XDR includes a “device discovery capability that helps you find unmanaged devices connected to your corporate network Device discovery uses onboarded endpoints, in your network, to collect, probe, or scan your network to discover unmanaged devices. The device discovery capability allows you to discover: Enterprise endpoints (workstations, servers, and mobile devices) that aren't yet onboarded to Defender for Endpoint[;] Network devices like routers and switches[;] IoT devices like printers and cameras[.].”¹¹⁴ Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
wherein a first computer system of the plurality of computer systems performs passive reconnaissance, and	A first computer system of the plurality of computer systems performs passive reconnaissance. For example, Defender XDR “passively collect[s] events in your network and extract[s] device information from them. Basic discovery uses the SenseNDR.exe binary for passive network data collection and no network traffic is initiated. Endpoints extract data from all network traffic seen by an onboarded device.”¹¹⁵

¹¹⁴ *Device Discovery.*

¹¹⁵ *Id.*

	Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
wherein a second computer system of the plurality of computer systems performs active reconnaissance.	A second computer system of the plurality of computer systems performs active reconnaissance. For example, Defender XDR includes a "[n]etwork device discovery" capability in which a "designated Microsoft Defender for Endpoint device is used on each network segment to perform periodic authenticated scans of preconfigured network devices. Once discovered, vulnerability management capabilities in Defender for Endpoint provide integrated workflows to secure discovered switches, routers, WLAN controllers, firewalls, and VPN gateways.... These types of devices require an agentless approach where a remote scan obtains the necessary information from the devices. Depending on the network topology and characteristics, a single device or a few devices onboarded to Microsoft Defender for Endpoint performs authenticated scans of network devices using SNMP (read-only)."¹¹⁶ Based on Qomplx's analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.

¹¹⁶ Microsoft, *Network device discovery and vulnerability management*, available at <https://learn.microsoft.com/en-us/defender-endpoint/network-devices>.

XIII. Claim 27

The computer system of claim 18,	See above for an analysis of Claim 18.
wherein the representation of the first graph comprises an identification of the third entity as a sensitive resource, and	The representation of the first graph of Claim 18 comprises an identification of the third entity as a sensitive resource. For example, Defender XDR’s blast radius analysis “generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions”:¹¹⁷ Blast radius analysis Blast radius analysis is an advanced graph visualization integrated into incident investigation experience. Built on the Microsoft Sentinel data lake and graph infrastructure, it generates an interactive graph showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions. Blast radius analysis also “helps analysts better assess the risk to highly regarded targets, and understand the business impact.”¹¹⁸ As another example, Microsoft’s documentation explains that a node can have a “[c]ritical asset” indicator associated with it, which “[i]ndicates that an entity is classified as business-critical or valuable[.]”¹¹⁹


¹¹⁷ *Investigate Incidents.*

¹¹⁸ *Id.*

¹¹⁹ *Understanding Graph Icons.*

	A node might also have any of the following indicators around it: • Critical asset - Indicates that an entity is classified as business-critical or valuable, as identified in the critical asset management in Microsoft Security Exposure Management. This indicator appears as a golden crown . The nodes representing critical assets also have a golden halo surrounding them. • Vulnerability - Indicates that at least one vulnerability was detected on the entity. This indicator appears as a red bug . • Explore connected assets - Indicates that the node can expand the hunting graph further beyond the initial results. Expanding the graph lets you explore other relationships the selected entity has with the other ones. This indicator appears as a blue plus sign . • Discovery source - Indicates the entity's data source. This indicator appears as the icon of the Defender product protecting the entity in blue (for example,  for Microsoft Defender for Endpoint, or  for Microsoft Defender for Cloud). In the example from Microsoft's documentation shown below, multiple nodes are identified as sensitive resources.¹²⁰
--	---

¹²⁰ *Investigate Incidents.*

	
wherein the computer system is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that generate a report that identifies	Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element. The computer system of Claim 18 is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that generate a report that identifies the third entity as a sensitive resource.

the third entity as a sensitive resource.	For example, Microsoft’s documentation explains that alerts in Defender XDR “can have system tags and/or custom tags with certain color backgrounds” that can “identify” among other things “[c]ritical assets involved in the incident.”¹²¹ An alert can have system tags and/or custom tags with certain color backgrounds. Custom tags use the white background while system tags typically use red or black background colors. System tags identify the following in an incident: • A type of attack, like ransomware or credential phishing • Automatic actions, like automatic investigation and response and automatic attack disruption • Defender Experts handling an incident • Critical assets involved in the incident Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted above. Any differences between the charted material and this claim element are expected to be insubstantial, in which case the element is expected to be practiced under the Doctrine of Equivalents. For example, the technology described above performs substantially the same function in substantially the same way to achieve substantially the same result as this claim element.
--	--

¹²¹ *Investigate Alerts* (emphasis omitted).