

MICROSOFT EXHIBIT 1045

Microsoft v. Qomplx

IPR2026-00183

EXHIBIT G
U.S. Patent No. 12,301,628

As used herein, the term “Accused ’628 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 ’628 Patent.

This claim chart for the ’628 Patent covers all Accused ’628 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 ’628 Defender Products.

I. Claim 1

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 '628 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 the Accused '628 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.

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 <table><tr><th>Icon</th><th>Node type</th><th>Entity type examples</th></tr><tr><td></td><td>General</td><td>App service plan</td></tr><tr><td></td><td>Compute</td><td>Device, virtual machine, Microsoft Azure Logic App</td></tr><tr><td></td><td>Networking</td><td>Interface, public IP address, network security group</td></tr><tr><td></td><td>Data</td><td>SQL data store, Azure Monitor Log Analytics workspace, storage account, Azure Event Hubs</td></tr><tr><td></td><td>Containers</td><td>Kubernetes cluster</td></tr><tr><td></td><td>Keys & secrets</td><td>Key vault</td></tr><tr><td></td><td>DevOps</td><td>Azure DevOps repositories</td></tr><tr><td></td><td>APIs</td><td>Cloud applications</td></tr><tr><td></td><td>Identity & access</td><td>User account, Microsoft Entra ID service principal</td></tr><tr><td></td><td>IoT</td><td></td></tr><tr><td></td><td>Certificate</td><td></td></tr><tr><td></td><td>IP address</td><td></td></tr><tr><td></td><td>Subscriptions</td><td></td></tr></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	
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 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”: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.  The screenshot shows a web interface titled 'Advanced hunting'. It features a search bar and a 'Selected Scenario' dropdown. The main area displays a directed graph with various nodes (e.g., 'User account', 'Device', 'IP address') and edges representing relationships. A red box in the bottom right corner highlights a control panel with buttons for 'Zoom in', 'Zoom out', and '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 '628 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, wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of	The first plurality of entities of the Accused ’628 Defender Products comprises a plurality of accounts and a plurality of resources, and each edge of the first plurality of edges of the Accused ’628 Defender Products corresponds to a respective relationship between a respective pair of entities of the first plurality of entities. For example, as discussed above, 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 Similarly, the edges that comprise Defender XDR’s graphs “indicate[] the relationship or connection properties between two nodes.”¹⁵						

¹³ *Id.*

¹⁴ *Understanding Graph Icons.*

¹⁵ *Id.*

entities of the first plurality of entities;

Examples of node and edge types and their icon representations are reproduced below:

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	

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

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 login to / Can login over the network to / Can remote interactive login to
	Runs on
	Provisions
	Identified as owner of
	Member of
	Is running
	Generic / Affects
	Created from / Used to create

	As another example, Defender XDR includes a “hunting graph,” as discussed above, which includes nodes corresponding to “entities in your environment (for example, a device, user account, or IP address, among others)” and includes “relationships or connection properties” between graph nodes.¹⁶ As another example, Microsoft’s documentation explains that nodes and edges can be represented as the ExposureGraphNodes and ExposureGraphEdges tables, discussed above, which “include entities like devices, identities, user groups and cloud assets such as virtual machines (VMs), storage, and containers,”¹⁷ and “relationships between entities and assets in the enterprise exposure graph.”¹⁸ 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 more entities of the first plurality of entities,	The software instructions of the Accused '628 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

¹⁶ *Hunting Graph*.

¹⁷ *ExposureGraphNode*s.

¹⁸ *ExposureGraphEdges*.

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*].

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

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

²⁰ *Id.*

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

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

²¹ *Id.*

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 nodes, wherein the first entity is not of the first plurality of entities;	Based on a first portion of the streaming data, the software instructions of the Accused '628 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. 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[ing] your network to discover unmanaged devices".²⁶

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

²⁶ Microsoft, *Device discovery overview*, available at <https://learn.microsoft.com/en-us/defender-endpoint/device-discovery> [hereinafter *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:²⁷



²⁷ Device Discovery.

	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 to any of the first plurality of edges;	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 '628 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, as noted above, 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

²⁸ *Incidents and Alerts.*

²⁹ *Device Discovery.*


	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. 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:
--	---

³⁰ *Investigate Incidents.*

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

³¹ *Id.*

	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. 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
--	---

³² Microsoft, *Alert correlation and incident merging in the Microsoft Defender portal*, available at <https://learn.microsoft.com/en-us/defender-xdr/alerts-incidents-correlation> [hereinafter *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 relationship, wherein the first node is not of the first	The software instructions of the Accused '628 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

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.
for an anomalous event associated with a node in the modified representation of the first graph, perform a first correlation using the modified representation of the first graph to identify a first plurality of correlated nodes, wherein each of the first plurality of correlated nodes corresponds to a respective event or resource, wherein each respective event or resource is associated with the anomalous event, and wherein each of the first plurality of correlated nodes is connected by a respective edge of a second plurality of edges to the node associated with the anomalous event in the	For an anomalous event associated with a node in the modified representation of the first graph, the software instructions of the Accused '628 Defender Products perform a first correlation using the modified representation of the first graph to identify a first plurality of correlated nodes, wherein each of the first plurality of correlated nodes corresponds to a respective event or resource, wherein each respective event or resource is associated with the anomalous event, and wherein each of the first plurality of correlated nodes is connected by a respective edge of a second plurality of edges to the node associated with the anomalous event in the modified representation of the first graph. For example, the Accused '628 Defender Products identify incidents, which are "collections of related alerts [that] tell the full story of an attack."³⁷ Microsoft's documentation explains that "correlation engines and algorithms in the Microsoft Defender portal automatically aggregate and correlate related alerts together to form incidents":³⁸

³⁷ *Incidents and Alerts.*

³⁸ *Id.*

modified representation of the first graph;

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

The Accused '628 Defender Products continuously monitor streaming data and “continue[] to monitor [incidents'] evolution, merging incidents together if situation warrants.”³⁹

The Accused '628 Defender Products associate incidents with entities. For example, “[e]ntities (assets etc.) follow the alerts they’re linked to.”⁴⁰ As another example, Microsoft’s documentation explains that entities are among the “common elements” of merged incidents:⁴¹

³⁹ *Alert Correlation.*

⁴⁰ *Id.*

⁴¹ *Incidents and Alerts.*

	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. Within a given incident, the Accused '628 Defender Products show an “incident graph,” including a set of correlated entities; the graph “shows the full scope of the attack,” including connections between “the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes”:⁴² 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. An example graph reproduced from Microsoft’s documentation is provided in the following image of the Defender Portal:⁴³
--	--

⁴² *Investigate Incidents.*

⁴³ *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.
for one or more of the first plurality of correlated nodes, perform a further correlation using the modified representation of the first graph to identify a second plurality of correlated nodes,	For one or more of the first plurality of correlated nodes, the software instructions of the Accused '628 Defender Products perform a further correlation using the modified representation of the first graph to identify a second plurality of correlated nodes, wherein each of the second plurality of correlated nodes is connected through a respective edge of a third plurality of edges to the respective node of the first plurality of correlated nodes in the modified representation of the first graph.

wherein each of the second plurality of correlated nodes is connected through a respective edge of a third plurality of edges to the respective node of the first plurality of correlated nodes in the modified representation of the first graph;

For example, as discussed above, the Accused '628 Defender Products identify incidents, which are “collections of related alerts [that] tell the full story of an attack,”⁴⁴ and monitor streaming data to “continue[] to monitor [incidents’] evolution, merging incidents together if situation warrants.”⁴⁵

As the streaming data is received and processed, the Accused '628 Defender Products merges incidents with commonalities, including based on common entities and events:⁴⁶

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.

In case of an incident merge, the entities and events of a “source incident” are migrated to a “target incident”:⁴⁷

⁴⁴ *Incidents and Alerts.*

⁴⁵ *Alert Correlation.*

⁴⁶ *Incidents and Alerts.*

⁴⁷ *Alert Correlation.*

	Details of the merge process When two or more incidents are merged, a new incident is <i>not</i> created to absorb them. Instead, the contents of one incident (the "source incident") are migrated into the other incident (the "target incident"), and the source incident is automatically closed. The source incident is no longer visible or available in the Defender portal, and any reference to it is redirected to the target incident. The source incident, though closed, remains accessible in Microsoft Sentinel in the Azure portal. Further, entities "follow the alerts they're linked to."⁴⁸ The Accused '628 Defender Products' correlation merge of incidents is an example of a further correlation. 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 representation of a second graph comprising representations of one or more of the first plurality of correlated nodes, representations of one or more of the second plurality of correlated nodes, representations of one or more of the second plurality	The software instructions of Accused '628 Defender Products generate a representation of a second graph comprising representations of one or more of the first plurality of correlated nodes, representations of one or more of the second plurality of correlated nodes, and representations of one or more of the second plurality of edges, and representations of one or more of the third plurality of edges, wherein one or more of the second plurality of edges together with one or more of the third plurality of edges represent one or more event flows that could be involved in a cybersecurity attack; and generate a report comprising information associated with the one or more event flows.

⁴⁸ *Id.*

of edges, and representations of one or more of the third plurality of edges, wherein one or more of the second plurality of edges together with one or more of the third plurality of edges represent one or more event flows that could be involved in a cybersecurity attack; and generate a report comprising information associated with the one or more event flows.	For example, the blast radius graph “provides a unique unified view of both prebreach and post-breach information on the incident page”:⁴⁹ 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. Microsoft’s documentation explains that the blast radius graphs can be used to “[i]nstantly see the compromised component at the center of the graph and the paths to potentially compromised targets.”⁵⁰ The blast radius is available from the incident graph page.⁵¹
--	---

⁴⁹ *Investigate Incidents*.

⁵⁰ *Id.*

⁵¹ *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. 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 4

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 identify the anomalous event.	The computer system of Claim 1 is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that identify the anomalous event. For example, Defender XDR provides “an incident’s details,” including “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.”⁵² In the screenshot from Microsoft’s documentation below, an “Incident details” pane is shown on the right side of the user interface.⁵³

⁵² *Id.*

⁵³ *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.

III. Claim 6

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 identify a portion of the modified representation of the first graph that matches a known attack pattern, wherein the portion of the modified representation of the first graph comprises the node associated with the anomalous event.	The computer system of Claim 1 is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that identify a portion of the modified representation of the first graph that matches a known attack pattern, wherein the portion of the modified representation of the first graph comprises the node associated with the anomalous event. For example, Defender XDR utilizes predictive analytics based on, among other things, “posture, activity, and scenario context to identify potential attack paths and targets,” using “threat intelligence, attacker behavior, [and] past incidents”:⁵⁴ How predictive shielding works Predictive shielding uses predictive analytics and real-time insights to dynamically identify emerging risks, and applies targeted protections. Predictive shielding integrates posture, activity, and scenario context to identify potential attack paths and targets, selectively hardening critical assets, or constraining attack paths just in time. This approach minimizes operational overhead and provides security teams with more time to respond. For example, predictive shielding can dynamically restrict access to sensitive data for devices identified as at-risk, reducing the need for broad, environment-wide restrictions. Predictive shielding relies on two pillars: • Prediction ◦ Involves analyzing threat intelligence, attacker behavior, past incidents, and organizational exposure. ◦ Defender uses this prediction data to identify emerging risks, to understand likely attack progression, and to infer risk on noncompromised assets. • Enforcement applies preventative protective controls to disrupt potential attack paths in real time. This dual approach ensures that protection is both precise and timely.

⁵⁴ Microsoft, *Predictive shielding in Microsoft Defender (Preview)*, available at <https://learn.microsoft.com/en-us/defender-xdr/shield-predict-threats> [hereinafter *Predictive Shielding*].

	As another example, Defender XDR uses, among other things, “known attacker tools and tactics”.⁵⁵ Prediction logic Prediction allows organizations to identify assets at risk and apply tailored protections in real time. Prediction focuses on emerging risks rather than static prevention, which minimizes operational friction and ensures that security measures are applied precisely where needed. For example, if a specific attacker tool is detected, predictive shielding can infer the next likely target based on past attack patterns. Defender uses multiple layers of insight to make accurate predictions: • Threat intelligence aligns observed activity with known attacker tools and tactics. • Learnings from past incidents are used to recognize statistical patterns, and extrapolate the most probable next steps. • Organizational exposure data is used to map how the environment is structured—which assets and identities are connected, which permissions these identities have, which vulnerabilities or misconfigurations exist, and how risk can propagate across them. Defender XDR continually updates the “organization’s exposure graph,” to identify “potential attack paths,” “the blast radius,” and “paths attackers are most likely to take, factoring in past behaviors, asset characteristics, and environmental vulnerabilities.”⁵⁶
--	---


⁵⁵ *Id.*

⁵⁶ *Id.*

	Graph-based logic Graph-based prediction logic bridges the gap between pre-breach and post-breach systems, providing a unified view of attacker activity across the organizational topology. This unified view includes the organization's assets, connections, and vulnerabilities. Graph-based logic combines live activity data with the structural map of the environment. This integration allows Defender to dynamically adjust protections based on the most critical vulnerabilities, enabling real-time prioritization of defenses and stopping attackers before they reach critical assets. The process involves three key stages: 1. Defender overlays post-breach activity onto the organization's exposure graph, creating a comprehensive view of potential attack paths. 2. Defender identifies the blast radius—the related assets that the identified activity might affect. 3. Reasoning models predict paths attackers are most likely to take, factoring in past behaviors, asset characteristics, and environmental vulnerabilities. This dynamic understanding allows Defender to move beyond reactive responses, enabling just-in-time protection that stops attackers before they reach critical assets. For example, in the screenshot from Microsoft's documentation shown below, the user interface displays, along with an incident graph for a "[h]ands-on keyboard attack," a list of various MITRE tactics and techniques, including "Group Policy Modification (T1484.001), Safe Mode Boot (T1562.009)," and "Local Account (T1087.001)":⁵⁷
--	--

⁵⁷ Microsoft, *Manage predictive shielding in Microsoft Defender (Preview)*, available at <https://learn.microsoft.com/en-us/defender-xdr/shield-predict-threats-manage> [hereinafter *Manage Predictive Shielding*].

Exhibit G ('628 Patent)
November 24, 2025



	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.
--	---

IV. Claim 9

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, an activity, an event, a group, or a service.	At least one of entity of the first plurality of entities of Claim 1 is at least one of a user, a place, a device, a resource, an activity, an event, a group, or a service. For example, Microsoft's documentation explains that graph nodes pertain to a wide variety of entity types in an environment, including, for example, "General" nodes that can pertain to, for example, entities like "App service plan[s]"; "Compute" nodes that can pertain to "Device, virtual machine, [or] Microsoft Azure Logic App" entities; "Networking" nodes that can pertain to "Interface, public IP address, [or] network security group" entities; "Data" nodes that can pertain to entities like "SQL data store, Azure Monitor Log Analytics workspace, storage account, [or] Azure Event Hubs"; "Containers" nodes that can pertain to entities like a "Kubernetes cluster"; "Keys & secrets" nodes that can pertain to entities like a "Key vault"; "DevOps" nodes that can pertain to entities like "Azure DevOps repositories"; "APIs" nodes that can pertain to entities like "Cloud applications"; and "Identity & access" nodes that can pertain to entities like "User account, [or] Microsoft Entra ID service principal."⁵⁸ Other node types include "IoT," "Certificate," "IP address," and "Subscriptions" nodes.⁵⁹ 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.

⁵⁸ *Understand Graph Icons.*

⁵⁹ *Id.*

V. Claim 11

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 '628 Defender Products are computer systems 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, as discussed above, Defender XDR “is a cloud-based, unified, pre- and post-breach enterprise defense suite”⁶⁰ that “operates in Microsoft Azure data centers” across a variety of geographical regions:⁶¹ 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:⁶²
---	---

⁶⁰ *Deploy Defender XDR.*

⁶¹ *Data in XDR.*

⁶² *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 '628 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. For example, as discussed above, 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.”⁶³ Microsoft’s documentation further 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.”⁶⁴ 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,

⁶³ *Understanding Graph Icons.*

⁶⁴ *Investigate Incidents.*

	among others).”⁶⁵ As another example, Microsoft’s documentation explains that nodes can be represented as an <code>ExposureGraphNode</code>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”.⁶⁶ 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.
--	---

⁶⁵ *Understanding Graph Icons.*

⁶⁶ *ExposureGraphNode*s.

	As another example, edges can be represented as an ExposureGraphEdges table of “relationships between entities and assets in the enterprise exposure graph”:⁶⁷ 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 ExposureGraphEdges 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.
--	--

⁶⁷ ExposureGraphEdges.

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”:⁶⁸

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.


⁶⁸ *Hunting Graph.*

	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 '628 Defender Products is a directed graph. For example, as noted above, the ExposureGraphEdges table schema includes information about “source nodes” and “target nodes.”⁶⁹ 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.”⁷⁰ As another example, as discussed above, Defender XDR includes a “hunting graph,” which is depicted by Microsoft’s documentation as a directed graph:⁷¹

⁶⁹ ExposureGraphEdges.

⁷⁰ Understanding Graph Icons.

⁷¹ 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.”⁷²

⁷² *Hunting Graph.*

	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, wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities of the first plurality of entities,	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 first plurality of entities of the Accused '628 Defender Products comprises a plurality of accounts and a plurality of resources, and each edge of the first plurality of edges of the Accused '628 Defender Products corresponds to a respective relationship between a respective pair of entities of the first plurality of entities. For example, as discussed above, 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.)"⁷³ Similarly, the edges that comprise Defender XDR's graphs "indicate[] the relationship or connection properties between two nodes."⁷⁴ Examples of node and edge types and their icon representations are reproduced below:		

⁷³ *Understanding Graph Icons.*

⁷⁴ *Id.*

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	

Edges	
Expand 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 login to / Can login over the network to / Can remote interactive login 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 nodes corresponding to “entities in your environment (for example, a device, user account, or IP address, among others)” and includes “relationships or connection properties” between graph nodes.⁷⁵ As another example, the <code>ExposureGraphNode</code>s and <code>ExposureGraphEdges</code> tables, discussed above, “include entities like devices, identities, user groups and cloud assets such as virtual machines (VMs), storage, and containers,”⁷⁶ and “relationships between entities and assets in the enterprise exposure graph.”⁷⁷ 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 representation of a first node of the first plurality of nodes includes information about a first security vulnerability associated with a first entity, wherein the first node corresponds to the first entity, and	A representation of a first node of the first plurality of nodes of the Accused '628 Defender Products includes information about a first security vulnerability associated with a first entity, wherein the first node corresponds to the first entity. For example, Microsoft’s documentation explains that a node can have a “[v]ulnerability” indicator associated with it, which “[i]ndicates that at least one vulnerability was detected on the entity.”⁷⁸

⁷⁵ *Hunting Graph.*

⁷⁶ *ExposureGraphNode*s.

⁷⁷ *ExposureGraphEdges.*

⁷⁸ *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). 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 security vulnerability is identified based on a database of known security vulnerabilities;	The first security vulnerability is identified based on a database of known security vulnerabilities. For example, Defender XDR uses “threat intelligence, breach likelihood predictions, business contexts, and device assessments to quickly prioritize the biggest vulnerabilities in your organization,” using “critical details including related CVEs,” and “vulnerabilities being exploited in the wild.”⁷⁹

⁷⁹ Microsoft, *What is Microsoft Defender Vulnerability Management*, available at <https://learn.microsoft.com/en-us/defender-vulnerability-management/defender-vulnerability-management> [hereinafter *Vulnerability Management*].

	Risk-based intelligent prioritization Defender Vulnerability Management uses Microsoft's threat intelligence, breach likelihood predictions, business contexts, and device assessments to quickly prioritize the biggest vulnerabilities in your organization. A single view of prioritized recommendations from multiple security feeds, along with critical details including related CVEs and exposed devices, helps you quickly remediate the biggest vulnerabilities on your most critical assets. Risk-based intelligent prioritization: • Focuses on emerging threats - Dynamically aligns the prioritization of security recommendations with vulnerabilities currently being exploited in the wild and emerging threats that pose the highest risk. • Pinpoints active breaches - Correlates vulnerability management and EDR insights to prioritize vulnerabilities being exploited in an active breach within the organization. • Protects high-value assets - Identifies exposed devices with business-critical applications, confidential data, or high-value users. As another example, Defender XDR uses “known Common Vulnerabilities and Exposures (CVEs)” listed “by their CVE ID.”⁸⁰ Vulnerabilities in my organization The <i>Weaknesses</i> page⁸⁰ in Microsoft Defender Vulnerability Management lists known Common Vulnerabilities and Exposures (CVEs) by their CVE ID. CVE IDs are unique IDs assigned to publicly disclosed cybersecurity vulnerabilities that affect software, hardware, and firmware. They provide organizations with a standard way to identify and track vulnerabilities, and helps them understand, prioritize, and address these vulnerabilities in their organization. CVEs are tracked in a public registry accessed from https://www.cve.org/. Defender Vulnerability Management uses endpoint sensors to scan and detect for these and other vulnerabilities in an organization. Applies to: • Microsoft Defender Vulnerability Management • Microsoft Defender for Endpoint Plan 2 • Microsoft Defender XDR • Microsoft Defender for Servers Plan 1 & 2
--	---

⁸⁰ Microsoft, *Vulnerabilities in my organization*, available at <https://learn.microsoft.com/en-us/defender-vulnerability-management/tvm-weaknesses> [hereinafter *TVM Weaknesses*].

	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 more entities of the first plurality of entities,	The software instructions of the Accused '628 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, as discussed above, Defender XDR "collect[s] . . . signals that are displayed in the portal," including 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."⁸¹ 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.

⁸¹ *Incidents and Alerts.*

	Further, as explained above, Defender “us[es] AI to continually monitor its telemetry sources”:⁸² 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,” “includ[ing] 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”:⁸³
--	---

⁸² *Id.*

⁸³ *Incidents and Alerts.*

	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".⁸⁴
--	--

⁸⁴ *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.”⁸⁵

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”:⁸⁶



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

⁸⁵ *Advanced Hunting.*

⁸⁶ *Investigate Incidents.*

	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 second entity that does not correspond to any of the first plurality of nodes, wherein the second entity is not of the first plurality of entities,	Based on a first portion of the streaming data, the software instructions of the Accused '628 Defender Products identify a second entity that does not correspond to any of the first plurality of nodes, wherein the second entity is not of the first plurality of entities. For example, as noted above, Defender XDR includes a device inventory which “is gradually populated with devices as they begin to report sensor data”:⁸⁷ 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[ing] your network to discover unmanaged devices”:⁸⁸

⁸⁷ *Device Inventory.*

⁸⁸ *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:⁸⁹



⁸⁹ *Device Discovery.*

	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 a second security vulnerability associated with the second entity,	The software instructions of the Accused '628 Defender Products identify a second security vulnerability associated with the second entity. For example, as noted above, Defender XDR uses "threat intelligence, breach likelihood predictions, business contexts, and device assessments to quickly prioritize the biggest vulnerabilities in your organization," using "critical details including related CVEs," and "vulnerabilities being exploited in the wild."⁹⁰ Microsoft's documentation further explains that Defender "uses endpoint sensors to scan and detect for [CVEs] and other vulnerabilities in an organization."⁹¹ 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	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 '628 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.

⁹⁰ *Vulnerability Management.*

⁹¹ *TVM Weaknesses.*

that does not correspond to any of the first plurality of edges,	For example, as noted above, Defender XDR “us[es] AI to continually monitor its telemetry sources” in order to “automatically aggregate and correlate related alerts”:⁹² 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:⁹³ 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. 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,”⁹⁴ based on the most recent device discovery information:
---	--

⁹² *Incidents and Alerts.*


⁹³ *Device Discovery.*

⁹⁴ *Investigate Incidents.*

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 showing possible propagation paths from the selected node to predefined critical targets scoped to the user’s permissions”:⁹⁵

⁹⁵ *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. 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”:⁹⁶ 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.
--	---

⁹⁶ Alert Correlation.

	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 second node corresponding to the second entity and a representation of a first edge corresponding to the first relationship, wherein the second node is not of the first plurality of nodes and the first edge is not of the first plurality of edges, and wherein the modified representation of the first graph includes information about the second security vulnerability associated with the second entity,	The software instructions of the Accused '628 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 second node corresponding to the second entity and a representation of a first edge corresponding to the first relationship, wherein the second node is not of the first plurality of nodes and the first edge is not of the first plurality of edges, and wherein the modified representation of the first graph includes information about the second security vulnerability associated with the second entity. For example, as explained above, Defender updates its graph representations as it discovers new entities and relationships so that they 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. . . . provid[ing] 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."⁹⁷ Microsoft's documentation further explains that Defender "automatically aggregate[s] and correlate[s] 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'

⁹⁷ *Understanding Graph Icons.*

	user accounts and device information and address, malicious files and processes, relevant threat intelligence, and so on.”⁹⁸ 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.”⁹⁹ Further, as explained above, a node represented on the graph representation can have a “[v]ulnerability” indicator associated with it, which “[i]ndicates that at least one vulnerability was detected on the entity.”¹⁰⁰ 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).
--	---

⁹⁸ *Incidents and Alerts.*

⁹⁹ *Investigate Incidents.*

¹⁰⁰ *Understanding Graph Icons.*

	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.
for a first event associated with a node in the modified representation of the first graph, perform a first correlation using the modified representation of the first graph to identify a first plurality of correlated nodes, wherein each of the first plurality of correlated nodes corresponds to a respective event or resource, wherein each respective event or resource is associated with the first event, and wherein each of the first plurality of correlated nodes is connected by a respective edge of a second plurality of edges to the node associated with the first event in the modified representation of the first graph,	For a first event associated with a node in the modified representation of the first graph, the software instructions of the Accused '628 Defender Products perform a first correlation using the modified representation of the first graph to identify a first plurality of correlated nodes, wherein each of the first plurality of correlated nodes corresponds to a respective event or resource, wherein each respective event or resource is associated with the first event, and wherein each of the first plurality of correlated nodes is connected by a respective edge of a second plurality of edges to the node associated with the first event in the modified representation of the first graph. For example, as explained above, the Accused '628 Defender Products identify incidents, "collections of related alerts [that] tell the full story of an attack," and use "correlation engines and algorithms in the Microsoft Defender portal [to] automatically aggregate and correlate related alerts together to form incidents".¹⁰¹

¹⁰¹ *Incidents and Alerts.*

	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 The Accused '628 Defender Products continuously monitor streaming data and “monitor [incidents’] evolution, merging incidents together if situation warrants.”¹⁰² The Accused '628 Defender Products associate incidents with entities. For example, “[e]ntities (assets etc.) follow the alerts they’re linked to.”¹⁰³ As another example, Microsoft’s documentation explains that entities are among the “common elements” of merged incidents:¹⁰⁴
--	---

¹⁰² *Alert Correlation.*


¹⁰³ *Id.*

¹⁰⁴ *Incidents and Alerts.*

	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. Within a given incident, the Accused '628 Defender Products show an “incident graph,” including a set of correlated entities; the graph “shows the full scope of the attack,” including connections between “the different suspicious entities that are part of the attack with their related assets such as users, devices, and mailboxes.”¹⁰⁵ An example graph reproduced from Microsoft’s documentation is provided below:¹⁰⁶
--	---

¹⁰⁵ *Investigate Incidents*.

¹⁰⁶ *Id.*

	 Human-operated ransomware attack was launched from a compromised asset (attac... Alerts (43) Assets (9) Investigations (1) Evidence and Response (76) Summary Similar incidents (1) Incident details Assigned to: Minix Emails - AlpineDataHouse Incident ID: 5176 Classification: Not set Categories: Initial access, Execution, Privilege escalation, Defense evasion, Credential access, Discovery, Lateral movement, Ransomware, Suspicious activity First activity: Mar 4, 2025 2:55:54 PM Last activity: Mar 4, 2025 3:28:11 PM Incident description A combination of several suspicious remote desktop protocol (RDP) session activities have been detected on this device. Threat actors might be attempting to establish a foothold in the environment by using various reconnaissance and persistence methods, then evade detection by tampering with and turning off security features to complete malicious objectives. See logs 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.
for one or more of the first plurality of correlated nodes, perform a further correlation using the modified representation of the first graph to identify a second plurality of correlated nodes, wherein each of the second plurality of correlated nodes	For one or more of the first plurality of correlated nodes, the software instructions of the Accused '628 Defender Products perform a further correlation using the modified representation of the first graph to identify a second plurality of correlated nodes, wherein each of the second plurality of correlated nodes is connected through a respective edge of a third plurality of edges to the respective node of the first plurality of correlated nodes in the modified representation of the first graph.

is connected through a respective edge of a third plurality of edges to the respective node of the first plurality of correlated nodes in the modified representation of the first graph,	For example, as discussed above, the Accused '628 Defender Products identify incidents, which are “collections of related alerts [that] tell the full story of an attack,”¹⁰⁷ and monitor streaming data to “continue[] to monitor [incidents’] evolution, merging incidents together if situation warrants.”¹⁰⁸ As the streaming data is received and processed, the Accused '628 Defender Products merges incidents with commonalities, including based on common entities and events:¹⁰⁹ 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. In case of an incident merge, the entities and events of a “source incident” are migrated to a “target incident”:¹¹⁰
--	---

¹⁰⁷ *Incidents and Alerts.*

¹⁰⁸ *Alert Correlation.*

¹⁰⁹ *Incidents and Alerts.*

¹¹⁰ *Alert Correlation.*

	Details of the merge process When two or more incidents are merged, a new incident is <i>not</i> created to absorb them. Instead, the contents of one incident (the "source incident") are migrated into the other incident (the "target incident"), and the source incident is automatically closed. The source incident is no longer visible or available in the Defender portal, and any reference to it is redirected to the target incident. The source incident, though closed, remains accessible in Microsoft Sentinel in the Azure portal. Further, entities "follow the alerts they're linked to."¹¹¹ The Accused '628 Defender Products' correlation merge of incidents is an example of a further correlation. 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 representation of a second graph comprising representations of one or more of the first plurality of correlated nodes, representations of one or more of the second plurality of correlated nodes, representations of one or more of the second plurality	The software instructions of Accused '628 Defender Products generate a representation of a second graph comprising representations of one or more of the first plurality of correlated nodes, representations of one or more of the second plurality of correlated nodes, representations of one or more of the second plurality of edges, and representations of one or more of the third plurality of edges, wherein one or more of the second plurality of edges together with one or more of the third plurality of edges represents one or more event flows that could be involved in a cybersecurity attack, and generate a report comprising information associated with the one or more event flows.

¹¹¹ *Id.*

of edges, and representations of one or more of the third plurality of edges, wherein one or more of the second plurality of edges together with one or more of the third plurality of edges represents one or more event flows that could be involved in a cybersecurity attack, and generate a report comprising information associated with the one or more event flows,	For example, the blast radius graph “provides a unique unified view of both prebreach and post-breach information on the incident page”:¹¹² 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. Microsoft’s documentation explains that the blast radius graphs can be used to “[i]nstantly see the compromised component at the center of the graph and the paths to potentially compromised targets.”¹¹³ The blast radius is available from the incident graph page:¹¹⁴
--	--


¹¹² *Investigate Incidents.*

¹¹³ *Id.*

¹¹⁴ *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. 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.
--	--

VI. Claim 13

The computer system of claim 11,	See above for an analysis of Claim 11.
wherein the further correlation using the modified representation of the first graph further identifies a plurality of respective correlated edges that each involve the respective event or resource.	The further correlation using the modified representation of the first graph of Claim 11 further identifies a plurality of respective correlated edges that each involve the respective event or resource. For example, Defender identifies correlated edges between nodes “[i]f two nodes have more than one relationship.”¹¹⁵ In the example from Microsoft’s documentation reproduced below, a user account entity has multiple relationships to a resource group:¹¹⁶ Selecting an edge opens a side panel that provides more details about the connection properties. If two nodes have more than one relationship, a number appears on the edge, in place of an icon. You can find more information about these nodes’ relationships by hovering over the number or opening the side panel. 

¹¹⁵ *Understand Graph Icons.*

¹¹⁶ *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.
--	---

VII. Claim 14

The computer system of claim 11,	See above for an analysis of Claim 11.
wherein the first event is an anomalous event.	The first event of Claim 11 is an anomalous event. For example, Microsoft’s documentation explains that Defender collects signals such as alerts, which “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.”¹¹⁷ 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.

¹¹⁷ *Investigate Alerts.*

VIII. Claim 15


The computer system of claim 13,	See above for an analysis of Claim 13.
wherein the first event is identified by comparing one or more of the events relating to one or more entities of the first plurality of entities to a pattern of normal event behavior.	The first event of Claim 13 is identified by comparing one or more of the events relating to one or more entities of the first plurality of entities to a pattern of normal event behavior. For example, Defender “detects anomalous behavior such as impossible-travel, credential access, and unusual downloading, file sharing, or mail forwarding activity.”¹¹⁸ Microsoft’s documentation explains that Defender “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.”¹¹⁹

¹¹⁸ *Deploy Defender XDR*.

¹¹⁹ 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*].

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.



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.

As another example, Defender uses “out-of-the-box user and entity behavioral analytics (UEBA) and machine learning (ML)” to “run advanced threat detection across your cloud environment” and “detect[] and collat[e] results, targeting numerous behavioral anomalies across your users and the machines and

	devices connected to your network.”¹²⁰ For example, Defender detects “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.”¹²¹ 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.
--	---

¹²⁰ 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*].

¹²¹ *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.
--	---

IX. Claim 17

The computer system of claim 11,	See above for an analysis of Claim 11.
wherein the computer system is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that identify a portion of the modified representation of the first graph that matches a known attack pattern, wherein the portion of the modified representation of the first graph comprises the node associated with the first event.	The computer system of Claim 11 is further configured to execute software instructions stored on nontransitory machine-readable storage media comprising software instructions that identify a portion of the modified representation of the first graph that matches a known attack pattern, wherein the portion of the modified representation of the first graph comprises the node associated with the first event. For example, as explained above, Defender XDR utilizes predictive analytics based on, among other things, “posture, activity, and scenario context to identify potential attack paths and targets,” using “threat intelligence, attacker behavior, [and] past incidents.”¹²² As another example, Defender XDR uses, among other things, “known attacker tools and tactics.”¹²³ Defender XDR continually updates the “organization’s exposure graph,” to identifying “potential attack paths,” “the blast radius,” and “paths attackers are most likely to take, factoring in past behaviors, asset characteristics, and environmental vulnerabilities.”¹²⁴ For example, in the screenshot from Microsoft’s documentation shown below, the user interface displays, along with an incident graph for a “[h]ands-on keyboard attack,” a list of various MITRE tactics and techniques.”¹²⁵


¹²² *Predictive Shielding*.

¹²³ *Id.*

¹²⁴ *Id.*

¹²⁵ *Manage Predictive Shielding*.

Exhibit G ('628 Patent)
November 24, 2025



	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.
--	---

X. Claim 20

The computer system of claim 11,	See above for an analysis of Claim 11.
wherein at least one entity of the first plurality of entities is at least one of a user, a place, a device, a resource, an activity, an event, a group, or a service.	At least one of entity of the first plurality of entities of Claim 11 is at least one of a user, a place, a device, a resource, an activity, an event, a group, or a service. For example, as explained above, Microsoft's documentation explains that graph nodes pertain to a wide variety of entity types in an environment, including, for example, entities like "[a]pp service plan[s]," "[d]evice[s], virtual machine[s], Microsoft Azure Logic App[s]," "public IP address[es], network security group[s]," "SQL data store[s], Azure Monitor Log Analytics workspace[s], storage account[s], Azure Event Hubs," "Kubernetes cluster[s]," "[k]ey vault[s]," "Azure DevOps repositories," "[u]ser account[s]," and others.¹²⁶ 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.

¹²⁶ *Understand Graph Icons.*

XI. Claim 22

The computer system of claim 11,	See above for an analysis of Claim 11.
wherein the node associated with the first event corresponds to a third security vulnerability.	The node associated with the first event of Claim 11 corresponds to a third security vulnerability. For example, as noted above, Defender XDR uses “threat intelligence, breach likelihood predictions, business contexts, and device assessments to quickly prioritize the biggest vulnerabilities in your organization,” using “critical details including related CVEs,” and “vulnerabilities being exploited in the wild.”¹²⁷ Microsoft’s documentation further explains that Defender “uses endpoint sensors to scan and detect for [CVEs] and other vulnerabilities in an organization.”¹²⁸ Microsoft’s documentation explains that nodes associated with vulnerabilities are denoted by a vulnerability indicator that “[i]ndicates that at least one vulnerability was detected on the entity.”¹²⁹ 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 . 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.

¹²⁷ *Vulnerability Management.*

¹²⁸ *TVM Weaknesses.*

¹²⁹ *Understanding Graph Icons.*