



**EXHIBIT H**  
**U.S. Patent No. 12,301,628**

As used herein and with respect to the '628 Patent, the term "Accused '628 Fusion Products" means:

- (a) Microsoft products that incorporate, rely upon, interact with, or otherwise utilize Microsoft Fusion ("Fusion"), including at least Microsoft Sentinel ("Sentinel") and Microsoft Defender;
- (b) Any other systems, services, or 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 Fusion 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 Fusion Products.

**I. Claim 1**

<p>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:</p>	<p>The Accused '628 Fusion 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.</p> <p>For example, Microsoft Sentinel “is a cloud-native Security Information and Event Management (SIEM) and unified security platform . . . built on a modern data lake.”<sup>1</sup> The Sentinel data lake “is a fully managed, cloud-native data lake purpose-built for security operations. It unifies, retains, and analyzes security data at scale - providing the foundation for advanced analytics, AI-driven insights, and agentic defense.”<sup>2</sup> Microsoft’s documentation explains that Sentinel data lake is priced based on the amount of data processed, the amount of data analyzed, compute hours, and the amount of data stored.<sup>3</sup></p>
--	---

---


<sup>1</sup> Microsoft, *What is Microsoft Sentinel?*, available at <https://learn.microsoft.com/en-us/azure/sentinel/sentinel-overview> [hereinafter *What is Microsoft Sentinel?*].

<sup>2</sup> *Id.*

<sup>3</sup> Microsoft, *Microsoft Sentinel pricing*, available at <https://www.microsoft.com/en-us/security/pricing/microsoft-sentinel> [hereinafter *Sentinel Pricing*].

	<p><b>Sentinel</b></p> <p>Microsoft Sentinel data lake enables security teams to ingest, retain, and analyze massive volumes of security data cost-effectively. With separate compute and storage meters, the data lake allows defenders to flexibly run advanced data insights, machine learning, and forensic investigations from a single point.<sup>3</sup> <a href="#">Learn more.</a></p> <table border="1"> <thead> <tr> <th>SKU</th> <th>Meter type</th> <th>Price</th> </tr> </thead> <tbody> <tr> <td>Data lake ingestion</td> <td>Data Processed (GB)</td> <td>\$0.05 USD</td> </tr> <tr> <td>Data processing</td> <td>Data Processed (GB)</td> <td>\$0.1 USD</td> </tr> <tr> <td>Data lake query</td> <td>Data Analyzed (GB)</td> <td>\$0.005 USD</td> </tr> <tr> <td>Advanced Data Insights</td> <td>1 Compute Hour</td> <td>\$0.15 USD</td> </tr> <tr> <td>Data lake storage</td> <td>Data Stored (GB/Month)</td> <td>\$0.026 USD</td> </tr> </tbody> </table>	SKU	Meter type	Price	Data lake ingestion	Data Processed (GB)	\$0.05 USD	Data processing	Data Processed (GB)	\$0.1 USD	Data lake query	Data Analyzed (GB)	\$0.005 USD	Advanced Data Insights	1 Compute Hour	\$0.15 USD	Data lake storage	Data Stored (GB/Month)	\$0.026 USD
SKU	Meter type	Price																	
Data lake ingestion	Data Processed (GB)	\$0.05 USD																	
Data processing	Data Processed (GB)	\$0.1 USD																	
Data lake query	Data Analyzed (GB)	\$0.005 USD																	
Advanced Data Insights	1 Compute Hour	\$0.15 USD																	
Data lake storage	Data Stored (GB/Month)	\$0.026 USD																	
<p>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,</p>	<p>The software instructions of the Accused '628 Fusion 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.</p> <p>For example, Microsoft's documentation explains that Fusion "builds and continually updates a hyperconnected graph on large scale data sets" in which "nodes represent the entities and the activities, and the edges represent the relationships between the nodes. . . . The entities can be IP addresses, accounts, Cloud resources, virtual machines, etc."<sup>4</sup></p>																		

<sup>4</sup> Microsoft, *Behind the Scenes: The ML Approach for Detecting Advanced Multistage Attacks with Sentinel Fusion*, available at <https://techcommunity.microsoft.com/blog/microsoftsentinelblog/behind-the-scenes-the-ml-approach-for-detecting-advanced-multistage-attacks-with/3239236> [hereinafter *Fusion Behind the Scenes*].

	<p><b>Graph forming:</b> Fusion builds and continually updates a hyperconnected graph on large scale data sets, typically millions of anomalous signals in a customer workspace. In the graph, the nodes represent the entities and the activities, and the edges represent the relationships between the nodes. The activities are the alerts and anomalies from different sources. The entities can be IP addresses, accounts, Cloud resources, virtual machines, etc.</p>  <p><i>Figure 1: Graph formed from a Microsoft Sentinel workspace</i></p> <p>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.</p>
wherein the first graph is a directed graph,	The first graph of the Accused '628 Fusion Products is a directed graph.  For example, Microsoft’s documentation shows directionality of edges in Fusion’s graph visualizations: <sup>5</sup>

---

<sup>5</sup> *Id.*




Figure 4: Expansion - probabilistic random walk



Figure 5: Expansion - aggregate weights and apply threshold

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

wherein the first plurality of entities comprises a plurality of accounts and a plurality of resources,

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;

The first plurality of entities of the Accused '628 Fusion Products comprises a plurality of accounts and a plurality of resources, and each edge of the first plurality of edges of the Accused Fusion Products corresponds to a respective relationship between a respective pair of entities of the first plurality of entities.

For example, as explained above, nodes in Fusion's graph representation "represent the entities and the activities, and the edges represent the relationships between the nodes. . . . The entities can be IP addresses, accounts, Cloud resources, virtual machines, etc.":<sup>6</sup>

**Graph forming:** Fusion builds and continually updates a hyperconnected graph on large scale data sets, typically millions of anomalous signals in a customer workspace. In the graph, the nodes represent the entities and the activities, and the edges represent the relationships between the nodes. The activities are the alerts and anomalies from different sources. The entities can be IP addresses, accounts, Cloud resources, virtual machines, etc.



Figure 1: Graph formed from a Microsoft Sentinel workspace

---

<sup>6</sup> *Id.*

<p>receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities,</p> <p>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;</p> <p>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;</p> <p>modify, in the hardware memory, the representation of the first graph to generate a modified representation of the first graph, wherein the</p>	<p>The software instructions of the Accused '628 Fusion Products receive streaming data comprising time-stamped data about events relating to one or more entities of the first plurality of entities, 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 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; and 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 plurality of nodes and the first edge is not of the first plurality of edges.</p> <p>For example, Microsoft's documentation explains that "Fusion correlates signals from multiple clouds, on-premise, and at the Edge for your entire enterprise, including anomalies, alerts from Microsoft products, as well as alerts from scheduled analytics rules . . . helping you to automatically detect sophisticated, multistage attacks":<sup>7</sup></p>
--	---


---

<sup>7</sup> *Id.*

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 plurality of nodes and the first edge is not of the first plurality of edges;

## From Anomalous Signals to High Fidelity Incidents: How Fusion Works End to End

Fusion operates a series of patented machine learning algorithms to look for advanced attacks from millions of anomalous signals. The process includes graph forming, attack pattern matching, expansion, scoring, and incident creation.



**Anomalous signals:** Fusion correlates signals from multiple clouds, on-premise, and at the Edge for your entire enterprise, including anomalies, alerts from Microsoft products, as well as alerts from scheduled analytics rules - both **built-in** and those **created by your security analysts** — helping you to automatically detect sophisticated, multistage attacks.

Microsoft’s documentation explains that Fusion uses this streaming data to “build[] and continually update[] a hyperconnected graph.”<sup>8</sup>

Microsoft’s documentation goes on to explain that Fusion collects data from multiple sources, including “[o]ut-of-the-box anomaly detections,” “[a]lerts from Microsoft services,” and “[a]lerts from scheduled analytics rules.”<sup>9</sup>

<sup>8</sup> *Id.*

<sup>9</sup> Microsoft, *Advanced multistage attack detection in Microsoft Sentinel*, available at <https://learn.microsoft.com/en-us/azure/sentinel/fusion> [hereinafter *Advanced Multistage Attack Detection*].

**Fusion for emerging threats** supports data collection and analysis from the following sources:

- Out-of-the-box anomaly detections
- Alerts from Microsoft services:
  - Microsoft Entra ID Protection
  - Microsoft Defender for Cloud
  - Microsoft Defender for IoT
  - Microsoft Defender XDR
  - Microsoft Defender for Cloud Apps
  - Microsoft Defender for Endpoint
  - Microsoft Defender for Identity
  - Microsoft Defender for Office 365
- Alerts from scheduled analytics rules. Analytics rules must contain kill-chain (tactics) and entity mapping information in order to be used by Fusion.

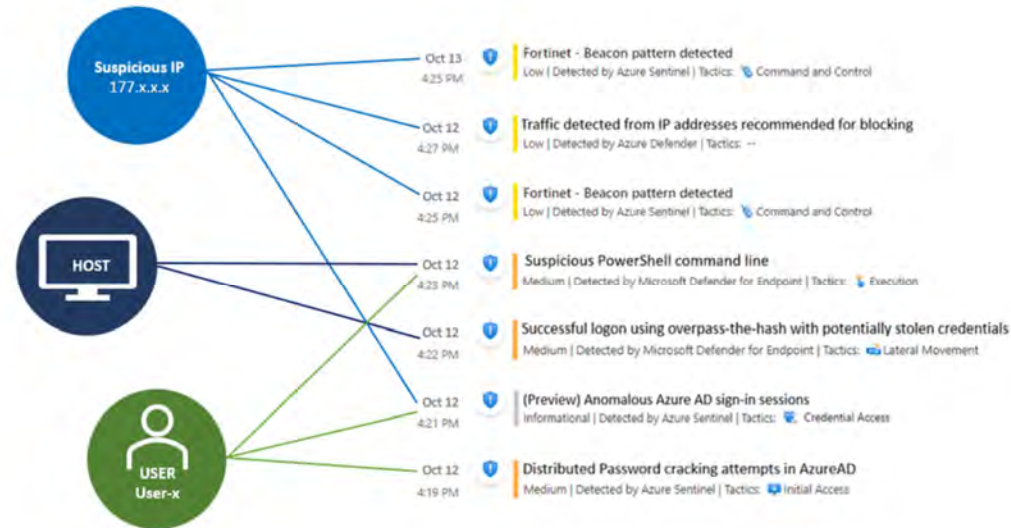
As another example, Microsoft's documentation provides the following example "possible attack" detected by Fusion that "started with initial access from the Cloud to end point execution, and then moved on to consistent beaconing from an internal IP address to a suspicious external one in roughly 24 hours":<sup>10</sup>

---

<sup>10</sup> Microsoft, *Detecting Emerging Threats with Microsoft Sentinel Fusion*, available at <https://techcommunity.microsoft.com/blog/microsoftsentinelblog/detecting-emerging-threats-with-microsoft-sentinel-fusion/2923835> [hereinafter *Detecting Emerging Threats*].

### An example

The example below shows a possible attack started with initial access from the Cloud to end point execution, and then moved on to consistent beaconing from an internal IP address to a suspicious external one in roughly 24 hours. The Fusion ML algorithms detected this attack by correlating anomaly (Anomalous Azure AD sign-in sessions), as well as alerts from custom scheduled rules, Azure Defender, and Microsoft Defender for Endpoint.



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

For an anomalous event associated with a node in the modified representation of the first graph, the software instructions of the Accused '628 Fusion 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


<p>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;</p>	<p>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.</p> <p>For example, Microsoft’s documentation explains that Fusion is “a correlation engine based on scalable machine learning algorithms, to automatically detect multistage attacks by identifying combinations of anomalous behaviors and suspicious activities that are observed at various stages of the attack chain. Based on these discoveries, Microsoft Sentinel generates incidents that would otherwise be difficult to catch.”<sup>11</sup> Microsoft’s documentation goes on to explain that Fusion “can help you find the emerging and unknown threats in your environment by applying extended ML analysis and by correlating a broader scope of anomalous signals.”<sup>12</sup></p> <p>As another example, Microsoft’s documentation explains that Fusion performs “[a]ttack pattern matching” that correlates nodes “in the hyperconnected graph” based on an attack pattern “consist[ing] of activities (nodes), entities (nodes), and their relationships (edges)”.<sup>13</sup></p>
---	--

---

<sup>11</sup> Microsoft, *Configure multistage attack detection (Fusion) rules in Microsoft Sentinel*, available at <https://learn.microsoft.com/en-us/azure/sentinel/configure-fusion-rules> [hereinafter *Configure Fusion Rules*].

<sup>12</sup> *Advanced Multistage Attack Detection*

<sup>13</sup> *Fusion Behind the Scenes*.

	<p><b>Attack pattern matching:</b> Fusion keeps a large set of attack patterns in a knowledge pool, including known attack patterns and ML generated emerging attack patterns. The known attack patterns are derived from past true positive incidents and security research. We will deep dive into how ML generates the emerging attack patterns in the next section of the blog.</p> <p>An attack pattern consists of activities (nodes), entities (nodes), and their relationships (edges). In this step, Fusion constantly takes attack patterns from the knowledge pool and identifies matches in the hyperconnected graph. Those identified matches are called subgraphs. This step reduces the millions of anomalous signals to a smaller set of subgraphs representing possible attacks. In the example below, three attack patterns are matched in the graph. There are 4 nodes and 3 edges in the top subgraph.</p>  <p><i>Figure 2: Simplified graph shows nodes and edges from attack pattern matching</i></p> <p>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.</p>
<p>for one or more of the first plurality of correlated nodes, perform a further correlation using the modified</p>	<p>For one or more of the first plurality of correlated nodes, the software instructions of the Accused '628 Fusion 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</p>

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;

plurality of correlated nodes in the modified representation of the first graph.

For example, Microsoft's documentation explains that Fusion "expands the matched attack patterns to discover additional activities and entities that are relevant".<sup>14</sup>

**Expansion:** During the expansion phase, Fusion expands the matched attack patterns to discover additional activities and entities that are relevant.




Figure 3: Simplified graph shows nodes and edges from attack pattern matching and expansion

Microsoft's documentation explains that "in the graph, the edges represent the relationship between the nodes. Fusion first uses an ML algorithm to assign a weight to each edge in the full graph to determine the relevance of the nodes by taking information including time range, kill chain intent, severity, entity

<sup>14</sup> *Id.*

	<p>type into consideration.”<sup>15</sup></p> <p>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.</p>
<p>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 second 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</p>	<p>The software instructions of the Accused '628 Fusion 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.</p> <p>For example, Microsoft’s documentation explains that Fusion “calculate[s] the killchain reachability of an attack and identif[ies] the nodes that have highest relevance in a real attack. In the example below, all the colored nodes (orange, yellow, green) are relevant to an attack. After the scoring round, Fusion only surfaces the nodes that have the highest relevance (orange and yellow colored nodes) in an incident. This way the security analysts only need to investigate a focused set of the most relevant activities and entities to quickly understand an attack”:<sup>16</sup></p>

---

<sup>15</sup> *Id.*

<sup>16</sup> *Id.*

could be involved in a cybersecurity attack; and generate a report comprising information associated with the one or more event flows.

**Scoring and incident creation:** Once the subgraphs representing possible attacks are identified, Fusion applies a round of scoring and triggers incidents that includes the most relevant alerts, anomalies, and entities to further reduce alert volume and speedup investigation.

In this step, Fusion uses k-nearest neighbors (KNN) to calculate the killchain reachability of an attack and identify the nodes that have highest relevance in a real attack. In the example below, all the colored nodes (orange, yellow, green) are relevant to an attack. After the scoring round, Fusion only surfaces the nodes that have the highest relevance (orange and yellow colored nodes) in an incident. This way the security analysts only need to investigate a focused set of the most relevant activities and entities to quickly understand an attack.




Figure 6: Simplified graph shows nodes and edges from attack pattern matching, expansion and scoring

Microsoft’s documentation goes on to provide an example “Possible multistage attack activities detected” report, reproduced below, which includes information associated with one or more event flows that could be involved in a cybersecurity attack:<sup>17</sup>

The example in *Figure 7* shows a possible attack that started with initial access from the Cloud to endpoint execution, and then moved on to consistent beaconing from an internal IP address to a suspicious external IP address, and possible Command and Control in roughly 24 hours. The Fusion ML algorithms detected this attack by correlating an anomaly (Anomalous Azure AD sign-in sessions), as well as alerts from custom scheduled rules, Azure Defender, and Microsoft Defender for Endpoint.

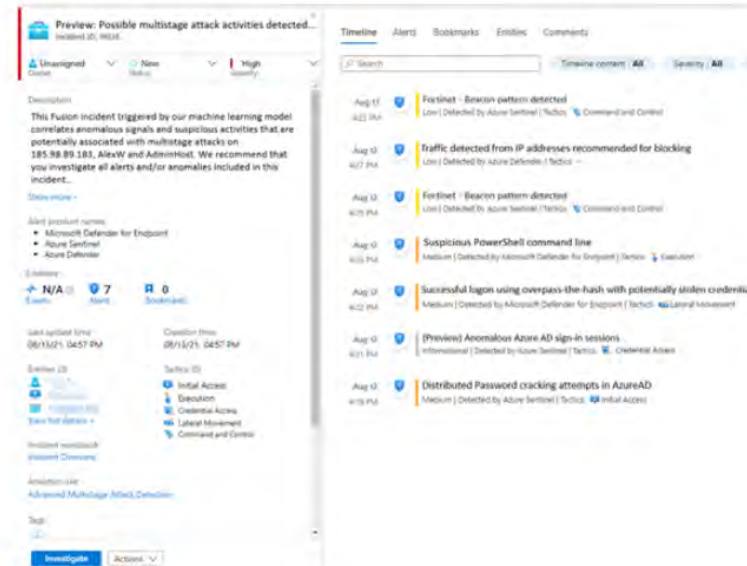


Figure 7: Fusion incident in Microsoft Sentinel workspace

Based on Qomplx’s analysis to date, this element appears to be practiced literally, including as charted

<sup>17</sup> *Id.*

	<p>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.</p>
--	---

**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, Microsoft's documentation explains that Fusion generates "Fusion incidents" when it detects "advanced multistage attacks." <sup>18</sup> The example reproduced from Microsoft's documentation below shows a Fusion incident for possible multistage attack activities: <sup>19</sup>

---

<sup>18</sup> Microsoft, *Scenarios detected by the Microsoft Sentinel Fusion engine*, available at <https://learn.microsoft.com/en-us/azure/sentinel/fusion-scenario-reference> [hereinafter *Fusion Scenarios*].

<sup>19</sup> *Fusion Behind the Scenes*

**Preview: Possible multistage attack activities detected...**  
Incident ID: 19834

Unassigned Owner | New Status | High Severity

**Description**  
This Fusion incident triggered by our machine learning model correlates anomalous signals and suspicious activities that are potentially associated with multistage attacks on 185.98.89.183, AlexW and AdminHost. We recommend that you investigate all alerts and/or anomalies included in this incident...

**Alert product names**

- Microsoft Defender for Endpoint
- Azure Sentinel
- Azure Defender

**Evidence**  
N/A Events | 7 Alerts | 0 Bookmarks

Last update time: 08/13/21, 04:57 PM | Creation time: 08/13/21, 04:57 PM

**Entities (3)**  
Initial Access, Execution, Credential Access, Lateral Movement, Command and Control

**Tactics (5)**  
Initial Access, Execution, Credential Access, Lateral Movement, Command and Control

**Timeline** | Alerts | Bookmarks | Entities | Comments

- Aug 13 4:25 PM: Fortinet - Beacon pattern detected (Low) | Detected by Azure Sentinel | Tactics: Command and Control
- Aug 12 4:27 PM: Traffic detected from IP addresses recommended for blocking (Low) | Detected by Azure Defender | Tactics: --
- Aug 12 4:25 PM: Fortinet - Beacon pattern detected (Low) | Detected by Azure Sentinel | Tactics: Command and Control
- Aug 12 4:23 PM: Suspicious PowerShell command line (Medium) | Detected by Microsoft Defender for Endpoint | Tactics: Execution
- Aug 12 4:22 PM: Successful logon using overpass-the-hash with potentially stolen credentials (Medium) | Detected by Microsoft Defender for Endpoint | Tactics: Lateral Movement
- Aug 12 4:21 PM: (Preview) Anomalous Azure AD sign-in sessions (Informational) | Detected by Azure Sentinel | Tactics: Credential Access
- Aug 12 4:19 PM: Distributed Password cracking attempts in AzureAD (Medium) | Detected by Azure Sentinel | Tactics: Initial Access

Investigate | Actions

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 5**

<p>The computer system of claim 4,</p>	<p>See above for an analysis of Claim 4.</p>
<p>wherein the anomalous 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.</p>	<p>The anomalous event of Claim 4 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.</p> <p>For example, as explained above, Fusion “supports data collection and analysis” from a variety of sources, including Sentinel’s “[o]ut-of-the-box anomaly detections.”<sup>20</sup> Such out-of-the-box anomaly detections include, for example, “UEBA anomalies” which “detect[] anomalies based on each entity's baseline historical behavior across various environments. Each entity's baseline behavior is set according to its own historical activities, those of its peers, and those of the organization as a whole. Anomalies can be triggered by the correlation of different attributes such as action type, geo-location, device, resource, ISP, and more”:<sup>21</sup></p> <p><b>UEBA anomalies</b></p> <p>Some of Microsoft Sentinel's anomaly detections come from its <a href="#">User and Entity Behavior Analytics (UEBA) engine</a>, which detects anomalies based on each entity's baseline historical behavior across various environments. Each entity's baseline behavior is set according to its own historical activities, those of its peers, and those of the organization as a whole. Anomalies can be triggered by the correlation of different attributes such as action type, geo-location, device, resource, ISP, and more.</p>

<sup>20</sup> Microsoft, *Advanced multistage attack detection in Microsoft Sentinel*, available at <https://learn.microsoft.com/en-us/azure/sentinel/fusion> [hereinafter *Advanced Multistage Attack Detection*].

<sup>21</sup> Microsoft, *Use customizable anomalies to detect threats in Microsoft Sentinel*, available at <https://learn.microsoft.com/en-us/azure/sentinel/soc-ml-anomalies> [hereinafter *Customizable Anomalies*].

	<p>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.</p>
--	--

**IV. Claim 6**

<p>The computer system of claim 1,</p>	<p>See above for an analysis of Claim 1.</p>
<p>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.</p>	<p>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.</p> <p>For example, Microsoft’s documentation explains that “Fusion keeps a large set of attack patterns in a knowledge pool, including known attack patterns and ML generated emerging attack patterns. The known attack patterns are derived from past true positive incidents and security research”:<sup>22</sup></p>

---

<sup>22</sup> *Fusion Behind the Scenes.*

**Attack pattern matching:** Fusion keeps a large set of attack patterns in a knowledge pool, including known attack patterns and ML generated emerging attack patterns. The known attack patterns are derived from past true positive incidents and security research. We will deep dive into how ML generates the emerging attack patterns in the next section of the blog.

An attack pattern consists of activities (nodes), entities (nodes), and their relationships (edges). In this step, Fusion constantly takes attack patterns from the knowledge pool and identifies matches in the hyperconnected graph. Those identified matches are called subgraphs. This step reduces the millions of anomalous signals to a smaller set of subgraphs representing possible attacks. In the example below, three attack patterns are matched in the graph. There are 4 nodes and 3 edges in the top subgraph.



Figure 2: Simplified graph shows nodes and edges from attack pattern matching

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.

**V. 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 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, as explained above, nodes in Fusion “represent . . . entities” such as “IP addresses, accounts, Cloud resources, virtual machines, etc.” <sup>23</sup>  As another example, Microsoft’s documentation provides a list of the types of entities identified by Sentinel: <sup>24</sup>

---

<sup>23</sup> *Fusion Behind the Scenes*.

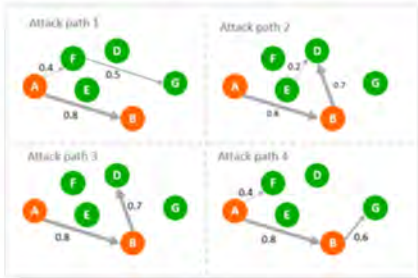
<sup>24</sup> Microsoft, *Entities in Microsoft Sentinel*, available at <https://learn.microsoft.com/en-us/azure/sentinel/entities>.

The following types of entities are currently identified in Microsoft Sentinel:

- Account
- Host
- IP address
- URL
- Azure resource
- Cloud application
- DNS resolution
- File
- File hash
- Malware
- Process
- Registry key
- Registry value
- Security group
- Mailbox
- Mail cluster
- Mail message
- Submission mail

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 10**

<p>The computer system of claim 1,</p>	<p>See above for an analysis of Claim 1.</p>
<p>wherein at least some of the edges of the second plurality of edges or of the third plurality of edges in the representation of the second graph are assigned a numerical weight that indicates a likelihood of a successful cybersecurity attack gaining access from one node to another.</p>	<p>At least some of the edges of the second plurality of edges or of the third plurality of edges in the representation of the second graph of Claim 1 are assigned a numerical weight that indicates a likelihood of a successful cybersecurity attack gaining access from one node to another.</p> <p>For example, Microsoft’s documentation explains that Fusion applies a “probabilistic kill chain model” “to determine viable attack paths in the graph from the matched patterns. The model runs multiple times to simulate different attack paths.”<sup>25</sup></p> <ul style="list-style-type: none"> <li>• <b>Run probabilistic random walk:</b> a probabilistic kill chain model is then applied to determine viable attack paths in the graph from the matched patterns. The model runs multiple times to simulate different attack paths. In the example below, A and B represent the nodes in a matched attack pattern and D, E, F, G represent the relevant activities and entities. In the real world, the subgraphs and attack paths are much more complicated and can be time consuming for security analysts to manually complete the process.</li> </ul>  <p>Figure 4: Expansion - probabilistic random walk</p> <p>“[A]fter simulating the different attack paths, Fusion aggregates weights across multiple runs.”<sup>26</sup></p>

<sup>25</sup> *Fusion Behind the Scenes.*

<sup>26</sup> *Id.*

- **Aggregate weights and apply threshold:** after simulating the different attack paths, Fusion aggregates weights across multiple runs. The algorithm first applies a threshold at the subgraph level to drop the subgraphs that represent unlikely attack paths. It then applies a threshold for each edge to determine how far to expand the graph.



Figure 5: Expansion - aggregate weights and apply threshold

As another example, Microsoft’s documentation explains that Fusion “applies a round of scoring” by “calculat[ing] the killchain reachability of an attack and identify[ing] the nodes that have highest relevance in a real attack.”<sup>27</sup>

---

<sup>27</sup> *Id.*

**Scoring and incident creation:** Once the subgraphs representing possible attacks are identified, Fusion applies a round of scoring and triggers incidents that includes the most relevant alerts, anomalies, and entities to further reduce alert volume and speedup investigation.

In this step, Fusion uses k-nearest neighbors (KNN) to calculate the killchain reachability of an attack and identify the nodes that have highest relevance in a real attack. In the example below, all the colored nodes (orange, yellow, green) are relevant to an attack. After the scoring round, Fusion only surfaces the nodes that have the highest relevance (orange and yellow colored nodes) in an incident. This way the security analysts only need to investigate a focused set of the most relevant activities and entities to quickly understand an attack.




Figure 6: Simplified graph shows nodes and edges from attack pattern matching, expansion and scoring

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.