IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

QOMPLX LLC,

Plaintiff,

v.

MICROSOFT CORPORATION,

Defendant.

Case No. 1:25-cv-01383

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

MICROSOFT CORP.
EXHIBIT 1020

1. Plaintiff QOMPLX LLC ("Plaintiff" or "Qomplx"), by its undersigned counsel, brings this Complaint against Defendant Microsoft Corporation ("Defendant" or "Microsoft"), for infringement of U.S. Patent Nos. 12,143,424 ("the '424 Patent," Ex. 1); 12,218,934 ("the '934 Patent," Ex. 2); 12,231,426 ("the '426 Patent," Ex. 3); 12,301,627 ("the '627 Patent," Ex. 4); 12,301,628 ("the '628 Patent," Ex. 5); and 11,539,663 ("the '663 Patent," Ex. 6) (collectively, the "Asserted Patents").

THE PARTIES

2. Plaintiff Qomplx is a corporation organized and existing under the laws of the State of New York, having a principal place of business at 1900 Reston Metro Plaza, Ste 600, Reston, Virginia 20190. Qomplx is a leading innovator with a focus on cybersecurity, analytics, simulation, and AI.

3. Qomplx is the assignee of and owns all right and title to the Asserted Patents.

4. On information and belief, Microsoft is a corporation organized under the laws of the State of Washington, with a place of business in this District located at 10900 Stonelake Boulevard, Suite 225, Austin, Texas 78759. Microsoft does business throughout the United States, including in this District.

JURISDICTION AND VENUE

5. The Court has subject matter jurisdiction under 28 U.S.C. § 1338, in that this action arises under federal statute, the patent laws of the United States (35 U.S.C. §§ 1, et seq.).

6. This Court has personal jurisdiction over Microsoft because it has substantial, systematic, and continuous contacts with this Judicial District. Microsoft has purposefully and voluntarily availed itself of the privileges of conducting business in the United States, the State of Texas, and this District by continuously and systematically placing goods into the

stream of commerce through an established distribution channel with the expectation that they will be purchased by consumers in this District. Microsoft directly and/or through intermediaries (including distributors, sales agents, and others), ships, distributes, sells, offers to sell, imports, advertises, makes, and/or uses its products (including but not limited to the products accused of infringement herein) in the United States, the State of Texas, and this District.

7. Upon information and belief, Microsoft conducts business within the State of Texas and in this District, and has designated Corporation Service Company d/b/a CSC − Lawyers Incorporating Service Company, 211 E. 7th Street, STE 620, Austin, Texas 78701-3218, as its agent for service of process in this District.

8. On information and belief, Microsoft has been registered to do business within the State of Texas under Texas Secretary of State File Number 0010404606 since about March 1987.

9. On information and belief, Microsoft employs one or more of its data centers in this District in furtherance of infringing acts in this District. For example, Microsoft maintains data centers in this District, located at: 5150 Rogers Road, San Antonio, Texas 78251, 5200 Rogers Road, San Antonio, Texas 78251, and 3823 Wiseman Boulevard, San Antonio, Texas 78251.

10. On information and belief, Microsoft has operated, and continues to operate, data centers supporting Microsoft products and services within the State of Texas, and within this District. Microsoft is building at least four additional data centers in this District, including data centers located at: 3545 Wiseman Boulevard, San Antonio, TX 78251; 3555 Westover Link, San Antonio, TX 78251; 2995 U.S. Highway 90 W, Castroville, TX 78009; and 15000 Block Lambda Drive, San Antonio, TX 78245. On information and belief,

Microsoft owns over $1 billion in property devoted to data centers in the Western District of Texas. Additionally, Microsoft has announced a $1.5 billion data center expansion in this District.[1]

11. Upon information and belief, Microsoft's data centers, including those in this District, include computer hardware (e.g., memory and processors) that store and execute software that performs some, or all, of the actions that infringe on the Asserted Patents. On information and belief, Microsoft has employed, is employing, and is offering to employ individuals in this District in furtherance of infringing acts in this District. On information and belief, these employees have direct personal knowledge about the accused products and Microsoft's infringing activities. These data centers that Microsoft operates constitute a regular and established physical presence in the District, including, but not limited to, ownership of or control over property, inventory, or infrastructure.

12. On information and belief, Microsoft has operated permanent office facilities within the State of Texas, and within this District, since at least 2000. The offices Microsoft maintains in this District include locations at 10900 Stonelake Boulevard, Suite 225, Austin, TX 78759 and Concord Park II, 401 East Sonterra Boulevard, Suite 300, San Antonio, TX 78258.

13. Microsoft operates offices in Austin, Texas for the purpose of selling, promoting, maintaining, and providing support for a suite of products, including the accused products.

14. On information and belief, Microsoft maintains a "Corporate Sales Office" in Austin, Texas at the following address: 10900 Stonelake Boulevard, Suite 225 Austin, TX,

[1] https://www.datacenters.com/news/microsoft-s-1-5-billion-data-center-expansion-in-san-antonio

78759; and Microsoft maintains a "Corporate Sales Office" in San Antonio, Texas at the following address: Concord Park II 401 East Sonterra Boulevard, Suite 300, San Antonio, TX, 78258.

15. On information and belief, one or more of the Accused Products are used, offered for sale, and sold in this District, including by Microsoft.

16. Qomplx's causes of action arise directly from Microsoft's business contacts and other activities in the State of Texas and this District.

17. Microsoft has also derived substantial revenues from its infringing acts within the State of Texas and this District.

18. In other recent actions, Microsoft has either admitted or not contested that this federal judicial district is a proper venue for patent infringement actions against it. *See, e.g.*, *Panther Innovations v. Microsoft Corp.*, No. 6:20-cv-01071, Dkt. No. 14; *Exafer Ltd. v. Microsoft Corp.*, No 1:20-cv- 00131, Dkt. No. 15; *Zeroclick, LLC v. Microsoft Corp.*, No. 1:20-cv-00272, Dkt. No. 14; and *California Institute of Technology v. Microsoft Corp*., No. 6:21-cv-00276, Dkt. No. 22.

19. Venue is proper in the Western District of Texas pursuant to 28 U.S.C. §§ 1391 and 1400(b) because Microsoft maintains regular and established places of business in this District and has committed acts of infringement within this District giving rise to this action.

20. Microsoft has committed acts of infringement in this District and does business in this District, including by committing acts of infringement in data centers in this District, and making sales and/or providing service and support for customers and/or end-users in this District. On information and belief, Microsoft purposefully and voluntarily sold one or more infringing products with the expectation they would be purchased in this District, and these infringing products have been and continue to be purchased in this District. Thus, Microsoft

has committed acts of infringement within the United States, the State of Texas, and this District.

21. Furthermore, Microsoft maintains corporate sales offices in this District, which, on information and belief, provide sales and support for the infringing products.

COUNT I

INFRINGEMENT OF U.S. PATENT NO. 12,143,424

22. Qomplx repeats and incorporates by reference each preceding paragraph as if fully set forth herein and further alleges:

23. The ’424 Patent, entitled “Rapid Predictive Analysis of Very Large Data Sets Using the Distributed Computational Graph,” was duly and lawfully issued on November 12, 2024 and assigned to QOMPLX LLC. A true and correct copy of the ’424 Patent is attached hereto as Exhibit 1.

24. The ’424 Patent names Jason Crabtree and Andrew Sellers as inventors.

25. The ’424 Patent claims priority to, among others, U.S. Application No. 14/925,974, filed October 28, 2015.

26. The ’424 Patent has been in full force and effect since its issuance. Qomplx owns all rights to the ’424 Patent that are necessary to bring this action.

27. The ’424 Patent, among other things, states that it “is in the field of analysis of very large data sets using distributed computational graph tools which allow for transformation of data through both linear and non-linear transformation pipelines.” **Ex. 1** at 3:32-35.

28. As the ’424 Patent explains, “[c]omputer database technology ha[d] allowed [a substantial amount of] information to be reliably stored for future retrieval and analysis,” but the prior art did “not have the tools to analyze all but a trickle into knowledge or informed

action." *Id.* at 3:61-4:20. Data retrieval and analysis methods had "proven to be too labor intensive and rigid to be of use in all but the more superficial and simple of campaigns." *Id.* at 4:11-4:18.

29. The '424 Patent further explains that "[d]ata pipelines, which are a progression of functions which each perform some action or transformation on a data stream, offer a mechanism to process quantities of data" in very large volumes. *Id.* at 4:31-34.

30. At the time of the invention, however, data pipelines were "extremely limited in what they d[id]." *Id.* at 4:34-35. For example, they could only perform limited tasks such as "move data from a web based merchant site to a distributed data store;" "extract all purchases and classify by product type and region;" and "store the result logs." *Id.* at 4:36-38. Or the data pipelines were "rigidly programmed and possibly required the uses of highly specific remote protocol calls to perform needed tasks." *Id.* at 4:39-40.

31. The '424 Patent overcomes the technological limitations of the prior art by proposing "a system for rapid predictive analysis of very large data sets using a distributed computational graph." *Id.* at 4:58-63.

32. Prior to the inventions claimed in the '424 Patent, data pipelines were extremely limited in the tasks that they could perform or were rigidly programmed. The distinct orchestration environment claimed in the '424 Patent provides a tangible improvement to computer technology by allowing substantially larger data sets to be analyzed in a timely and efficient manner. By enlisting additional computer systems on an as-needed basis, claimed embodiments can "introduce new transformation pipelines just as they are needed, creating only those that are ready to compute." *Id.* at 9:40-41. This architecture creates improvements to the functioning of computer systems utilizing the invention.

33. For example, claimed embodiments in the '424 Patent allow the system to "introduce new transformation pipelines just as they are needed, creating only those that are ready to compute" in order to compensate for the "exponential growth of resource consumption" that may arise in streaming data systems. *Id.* at 9:35-44. This enables the system to operate "in a timely and efficient manner" by providing "the ability to monitor for both operational issues within its components," "to learn and react to intermediate determinations of the analyses it runs," and to "self-modify to maintain optimal operation." *Id.* at 4:49-54.

34. The claims of the '424 Patent thus cover specific solutions that provide specific technological advantages to data processing systems. For example, they claim unconventional distributed architectures that enable smart scaling, which closely matches resource usage to resource needs even with unpredictable streaming data inputs, making the overall system more efficient. They allow the system to scale both up and down dynamically, an inherently difficult problem. For example, scaling down is a significant challenge due to the need to handle data processing that is "in flight" during the rescaling. The result is that these solutions can run in environments that periodically generate massive traffic, reducing average latency to acceptable limits. Additionally, they can scale down to smaller environments, reducing total compute utilization. In the cloud era, this reduced utilization leads to directly reduced cost. Moreover, and as the specification explains, such streaming input feeds can include a wide variety of sources, including for example "the internet," "arrays of physical sensors," "database servers," "electronic monitoring equipment," and "direct human interaction." *Id.* at 15:35-40. As the processing of such streaming data can be "limited" by "the resources of the system," *id.* at 9:64-67, the scaling capabilities recited in the claims provide discrete and tangible improvements to data processing architectures by allowing the architecture to

automatically enlist additional computing resources on an as-needed basis, distributing the increased compute load over multiple computer systems when necessary.

35. As another example, the claimed distributed computational graph combines knowledge of the network architecture and data flow in a unique way that enables the system to be more adaptable, by, for example, replicating certain processing steps to remove bottlenecks in overall processing time. As another example, the specification explains that the claimed distributed computational graph allows the system to "monitor for data searches or transformations that are processing slowly or may have hung and for results that are outside established data stability boundaries so that actions can be implemented to resolve the issue," enabling both "autonomous[ ]" action by the system itself and "status updates . . . made by administrators" or "direct changes to operational parameters by such." *Id.* at 16:53-61. This leads to improvements in the distribution of response latencies, both reducing average and worst-case latencies.

36. Accordingly, the claims of the '424 Patent provide discrete technological solutions to challenges unique to computer systems, such as the necessity of dynamic and efficient resource allocation in a data processing environment that requires analysis of streaming data from a variety of potentially disparate input sources.

37. The '424 Patent's improvements to very large data analysis systems are further described in the specification.

38. Microsoft is not currently licensed to practice the '424 Patent.

39. Qomplx is informed and believes, and thereon alleges, that Microsoft has infringed and continues to infringe one or more claims of the '424 Patent in violation of 35 U.S.C. § 271, either literally and/or under the doctrine of equivalents, by making, using, directing an entity to use, selling, and/or offering for sale in the United States, and/or

importing into the United States, without authorization, Microsoft products that practice one of more claims of the ’424 Patent, including without limitation Microsoft Fabric (“Fabric”) (collectively, the “Accused ’424 Products”).

40. For example and without limitation, the Accused ’424 Products embody every limitation of at least claim 1 of the ’424 Patent, both literally and under the doctrine of equivalents, as set forth below.

41. Claim 1 of the ’424 Patent provides:

> A distributed computing cluster comprising:
> a first plurality of computer systems,
> wherein each respective computer system of the first plurality of computer systems comprises a memory that stores a respective first data,
> wherein the respective first data represents a respective portion of a distributed computational graph
> and wherein the distributed computational graph describes a flow of output data of a first transformation pipeline to an input of a second transformation pipeline,
> wherein a first computer system of the first plurality of computer systems is configured to:
> receive a first stream of input data from a first input feed,
> process the first stream of input data substantially in real time by executing software instructions that apply the first transformation pipeline to the first stream of input data to generate first pipeline output messages,
> process the respective first data stored in the memory of the first computer system to determine information about the second transformation pipeline,
> and transmit the first pipeline output messages to a second computer system of the first plurality of computer systems in accordance with the determined information,
> wherein the second computer system is configured to:
> receive the first pipeline output messages,
> and process the first pipeline output messages substantially in real time by executing software instructions that apply the second transformation pipeline to the first pipeline output messages to generate second pipeline output messages,
> wherein the first and second computer systems are distinct; and
> a second plurality of computer systems;
> wherein a third computer system of the first plurality of computer systems is configured to execute software instructions that cause a fourth computer system of the second plurality of computer systems to execute software instructions that apply at least one of the first transformation pipeline and the second transformation pipeline.

42. The Accused '424 Products meet every element of this claim. The further descriptions below, which are based on analysis of publicly available information, are preliminary examples and are non-limiting.

43. For example, and to the extent that the preamble is limiting, Fabric forms a distributed computing cluster.

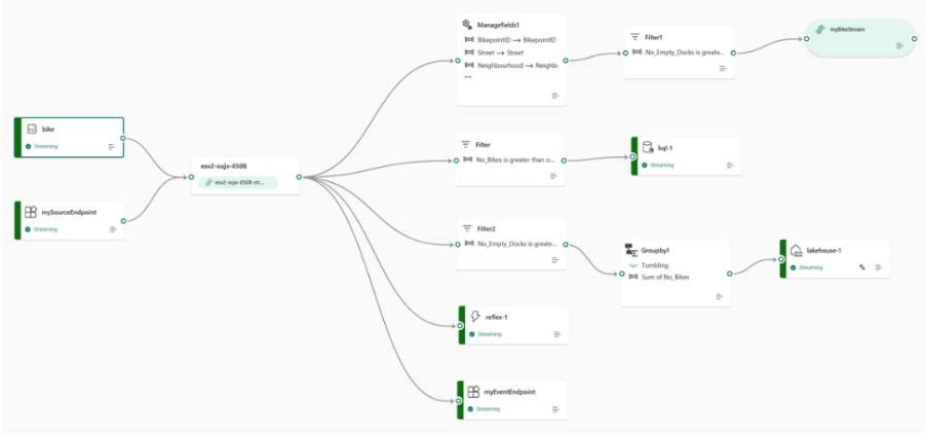
44. For example, Fabric is a unified data analytics platform that runs on a distributed computer cluster, and includes the Real-Time Intelligence tool, which includes Eventstream.[2] Eventstream runs across an auto-scaling plurality of computer systems.[3]

45. On information and belief, the Accused '424 Products comprise a first plurality of computer systems, and each respective computer system of a first plurality of computer systems of the Accused '424 Products comprises a memory that stores a respective first data, which represents a portion of a distributed computational graph. The distributed computational graph describes a flow of output data from a first transformation pipeline to an input of a second transformation. For example, Eventstream allows a user to define, and then execute, a distributed computational graph that describes the flow of data between transformation pipeline elements.[4]

[2] https://learn.microsoft.com/en-us/fabric/fundamentals/microsoft-fabric-overview

[3] https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-streams/overview?tabs=enhancedcapabilities

[4] https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-streams/overview?tabs=enhancedcapabilities



46. On information and belief, a computer system of the first plurality of computer systems of the Accused ’424 Products is configured to receive a stream of input data from an input feed. For example, Eventstream can receive various streams of input data from various input feeds.[5] Additionally, Eventstream can receive event streams from other computing steps.

Bring events into Fabric

The eventstreams feature provides you with various source connectors to fetch event data from the various sources. There are more sources available when you enable **Enhanced capabilities** at the time of creating an eventstream.

Enhanced capabilities | Standard capabilities

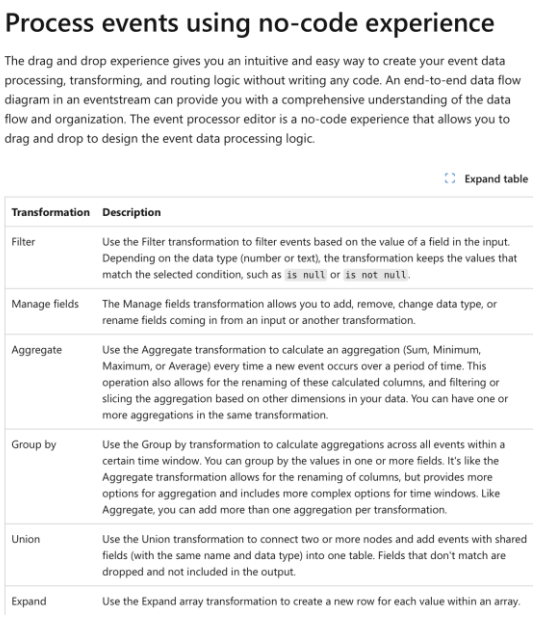
Expand table

Sources	Description
Azure Event Hubs	If you have an Azure event hub, you can ingest event hub data into Microsoft Fabric using Eventstream.
Azure IoT Hub	If you have an Azure IoT hub, you can ingest IoT data into Microsoft Fabric using Eventstream.
Sample data	You can choose **Bicycles**, **Yellow Taxi**, or **Stock Market events** as a sample data source to test the data ingestion while setting up an eventstream.
Custom App	The custom app feature allows your applications or Kafka clients to connect to Eventstream using a connection string, enabling the smooth ingestion of streaming data into Eventstream.

[5] https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-streams/overview?tabs=standardcapabilities

47. On information and belief, a computer system of the first plurality of computer systems of the Accused ’424 Products is configured to process the stream of input data substantially in real time by executing software instructions that apply the first transformation pipeline to the stream of input data to generate first pipeline output messages, process the respective first data to determine information about the second transformation pipeline, and transmit the first pipeline output messages to a second computer system of the first plurality of computer systems of the Accused ’424 Products in accordance with the determined information. On information and belief, a second computer system of the first plurality of computer systems of the Accused ’424 Products is configured to receive the first pipeline output messages, process those messages substantially in real time by applying the second transformation pipeline to the first pipeline output messages and generating second pipeline output messages. For example, Fabric is configured with various processing steps that receive, process substantially in real time, including by applying transformation pipelines, and then transmit streams of data.[6]

[6] https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-streams/overview?tabs=standardcapabilities



Process events using no-code experience

The drag and drop experience gives you an intuitive and easy way to create your event data processing, transforming, and routing logic without writing any code. An end-to-end data flow diagram in an eventstream can provide you with a comprehensive understanding of the data flow and organization. The event processor editor is a no-code experience that allows you to drag and drop to design the event data processing logic.

Expand table

Transformation	Description
Filter	Use the Filter transformation to filter events based on the value of a field in the input. Depending on the data type (number or text), the transformation keeps the values that match the selected condition, such as `is null` or `is not null`.
Manage fields	The Manage fields transformation allows you to add, remove, change data type, or rename fields coming in from an input or another transformation.
Aggregate	Use the Aggregate transformation to calculate an aggregation (Sum, Minimum, Maximum, or Average) every time a new event occurs over a period of time. This operation also allows for the renaming of these calculated columns, and filtering or slicing the aggregation based on other dimensions in your data. You can have one or more aggregations in the same transformation.
Group by	Use the Group by transformation to calculate aggregations across all events within a certain time window. You can group by the values in one or more fields. It's like the Aggregate transformation allows for the renaming of columns, but provides more options for aggregation and includes more complex options for time windows. Like Aggregate, you can add more than one aggregation per transformation.
Union	Use the Union transformation to connect two or more nodes and add events with shared fields (with the same name and data type) into one table. Fields that don't match are dropped and not included in the output.
Expand	Use the Expand array transformation to create a new row for each value within an array.

48. These transformation pipelines run across Fabric "Capacity Units," which map to multiple computer systems.[7]

[7] https://learn.microsoft.com/en-us/fabric/enterprise/licenses#capacity

SKU*	Capacity Units (CU)	Power BI SKU	Power BI v-cores
F2	2	-	0.25
F4	4	-	0.5
F8	8	EM/A1	1
F16	16	EM2/A2	2
F32	32	EM3/A3	4
F64	64	P1/A4	8
Trial	64	-	8
F128	128	P2/A5	16
F256	256	P3/A6	32
F512	512	P4/A7	64
F1024	1024	P5/A8	128
F2048	2048	-	256

49. On information and belief, a third computer system of the first plurality of computer systems of the Accused '424 Products is configured to cause a fourth computer system of the second plurality of computer systems of the Accused '424 Products to apply at least one of the transformation pipelines. For example, Fabric uses an autoscaler that—on an as-needed basis—automatically recruits additional computer systems to apply processing pipelines.[8]

[8] https://learn.microsoft.com/en-us/fabric/real-time-intelligence/event-streams/monitor-capacity-consumption#Note-2

- **Note 2: Eventstream Processor Per Hour**. The CU consumption rate of the Eventstream processor is correlated to the throughput of event traffic, the complexity of the event processing logic, and the partition count of input data:
  - With "Low" set in "Event throughput setting", the processor CU consumption rate starts at 1/3 base-rate (0.778 CU hour) and autoscale within 2/3 base-rate (1.555 CU hour), 1 base-rate (2.333 CU hour), 2 base-rates, and 4 base-rates.
  - With "Medium" set in "Event throughput setting", the processor CU consumption rate starts at 1 base-rate and autoscale within multiple possible base-rates.
  - With "High" set in "Event throughput setting", the processor CU consumption rate starts at 2 base-rates and autoscale within multiple possible base-rates.

50. Other Azure products such as Microsoft Service Fabric show that units of computation are often distributed "across a cluster of machines" "at massive scale."[9] On information and belief, Microsoft Fabric uses elements of architecture and orchestration systems related to Microsoft Service Fabric.

51. Accordingly, as illustrated above, the Accused '424 Products directly infringe one or more claims of the '424 Patent. Microsoft makes, uses, sells, offers for sale, and/or imports, in this District and/or elsewhere in the United States the Accused '424 Products and thus directly infringes the '424 Patent.

52. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the '424 Patent, as provided in 35 U.S.C. § 271(b), including at least by inducing others, such as Microsoft's customers and end-users, in this District and elsewhere in the United States, to use the Accused '424 Products in manners that infringe the '424 Patent. Microsoft induces such direct infringement through its affirmative acts of making, using, directing an entity to use, selling, offering to sell, and/or importing the Accused '424 Products, as well as by advertising the Accused '424 Products and providing instructions, documentation, and other information to its customers

[9] https://learn.microsoft.com/en-us/azure/service-fabric/service-fabric-overview#container-orchestration

and end-users to encourage and teach them how to use the infringing Accused '424 Products, including but not limited to by Microsoft providing in-store and online technical support, marketing materials, product manuals, advertisements, and other product documentation. At least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the '424 Patent and with the intent, or willful blindness, that the induced acts directly infringe the '424 Patent.

53. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the '424 Patent, as provided by 35 U.S.C. § 271(c), by contributing to direct infringement committed by others, such as Microsoft's customers and end-users, in this District and elsewhere in the United States. Microsoft's affirmative acts of selling and offering to sell the Accused '424 Products in this District and elsewhere in the United States, and causing the Accused '424 Products to be manufactured, used, sold, and offered for sale, contribute to Microsoft's customers and end-users using the Accused '424 Products, such that the '424 Patent is directly infringed. The accused components in the Accused '424 Products are material to the inventions claimed in the '424 Patent, are not staple articles or commodities of commerce, have no substantial non-infringing uses, and are known by Microsoft to be especially made or adapted for use in the infringement of the '424 Patent during at least the post-service period. Similarly, at least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the '424 Patent and with the intent, or willful blindness, that they cause direct infringement of the '424 Patent.

54. On information and belief, no licensed patent-practicing commercial products have been made, sold, offered for sale, or imported in a manner that would trigger the requirements of 35 U.S.C. § 287.

55. Microsoft's infringement of the '424 Patent has damaged and will continue to damage Qomplx.

56. Upon service of this complaint, Microsoft will have explicit written notice of its infringement of the '424 Patent. Nevertheless, on information and belief, Microsoft will, without authorization, knowingly, intentionally, purposefully, and deliberately continue to infringe the '424 Patent.

COUNT II

INFRINGEMENT OF U.S. PATENT NO. 12,218,934

57. Qomplx repeats and incorporates by reference each preceding paragraph as if fully set forth herein and further alleges:

58. The '934 Patent, entitled "Contextual and Risk-Based Multi-Factor Authentication," was duly and lawfully issued on February 4, 2025 and assigned to QOMPLX LLC. A true and correct copy of the '934 Patent is attached hereto as Exhibit 2.

59. The '934 Patent names Jason Crabtree, Andrew Sellers, and Ian MacLeod as inventors.

60. The '934 Patent claims priority to, among others, U.S. Provisional No. 62/574,708, filed October 19, 2017.

61. The '934 Patent has been in full force and effect since its issuance. Qomplx owns all rights to the '934 Patent that are necessary to bring this action.

62. The '934 Patent, among other things, states that it "relates to the field of network security, particularly to multi-factor user authentication." **Ex. 3** at 2:15-16.

63. As the '934 Patent explains, multi-factor authentication ("MFA") systems "commonly used today include[] one-time use codes sent to a user's mobile device or email,

confirming through a uniquely generated link sent to the user, or using authenticator devices and apps that generate a code on-demand." *Id.* at 2:23-27.

64. The '934 Patent addresses a key "fault" in those systems: "over-reliance on a single method of delivery." *Id.* at 2:28-29. For example, "once a user's email is compromised, it may be trivial to gain access to that user's other accounts" because "password reset" requests "are usually conducted through the user's email" and "[o]ne time codes, such as those from banking websites, are also often sent to the user's email address." *Id*. at 2:29-34.

65. Departing from the prior art's limited and ineffective approach to authentication, the claims of the '934 Patent are directed to specific verification methods that improve computer security.

66. For example, Claim 1 identifies a specific technique by which a computer system may "determine whether an additional verification is required to grant access." Among other things, a computer system implementing the claimed authentication technology can, upon "receiv[ing] a request to authenticate a client" and "storing, in a multidimensional time-series database, information about the request," retrieve "historical information about previous access requests associated with the user account" from the multidimensional time-series database and use that information to determine whether the user account is associated with a "previous access request to a network resource" that "is anomalous relative to a baseline profile of access requests" and, if so, select from multiple verification methods to authenticate the client.

67. Accordingly, the authentication techniques claimed in the '934 Patent, including the particular sequence of steps recited in the claims to determine whether additional authentication is needed, "eliminate" MFA's "over-reliance on a single, and possibly compromised, method" by using "a combination of verification methods" and "dynamically

determin[ing] the varying amounts of verification needed, based on the contexts and risks associated with the connection," yielding higher levels of security and improving computer functionality. *Id.* at 2:35-40.

68. As the specification explains, the contextual and risk-based approach to authentication claimed in the '934 Patent allows a computer security system in certain embodiments to "intelligently integrate the large volume of data from a plurality of sources on an ongoing basis" and identify anomalous and potentially threatening network activity by "continuously poll[ing] the incoming traffic data for activities anomalous to [network] baseline." *Id*. at 8:23-35. These embodiments are able to react with less latency to attacks. They can also process a larger bandwidth of input data, resulting in the detection of a broader base of potential attack types. Both of these improvements to computer performance result in blocking accesses that prior art systems could not block. The specification provides numerous examples of such anomalous activity, including "a user attempting to gain access [to] several network resources such as workstations or servers in rapid succession, or a user attempting to gain access to a domain server of server with sensitive information using random userIDs or another user's userID and password, or attempts by any user to brute force crack a privileged user's password," among others. *Id.* at 8:35-46. Through the storage and retrieval of historical network access information, and the application of additional verification methods to incoming access requests based on that historical information as compared to the network's baseline profile, the claims of the '934 Patent improve on the prior art's approach to authentication technology, provide tangible improvements to network security, and mitigate the risk of cyberattack or data loss.

69. The '934 Patent's improvements to network security are further described in the specification.

70. Microsoft is not currently licensed to practice the ’934 Patent.

71. Qomplx is informed and believes, and thereon alleges, that Microsoft has infringed and continues to infringe one or more claims of the ’934 Patent in violation of 35 U.S.C. § 271, either literally and/or under the doctrine of equivalents, by making, using, directing an entity to use, selling, and/or offering for sale in the United States, and/or importing into the United States, without authorization, Microsoft products that practice one of more claims of the ’934 Patent, including without limitation products that incorporate, rely upon, interact with, or otherwise utilize Microsoft Entra ID (“Entra ID” or the “Accused ’934 Functionality”).

72. For example and without limitation, products that incorporate, rely upon, interact with, or otherwise utilize Entra ID, including at least the Microsoft Entra product family (collectively, the “Accused Entra Products”) embody every limitation of at least claim 1 of the ’934 Patent, both literally and under the doctrine of equivalents, as set forth below.

73. Claim 1 of the ’934 Patent provides:

> A computer system configured to execute software instructions stored on nontransitory machine-readable storage media, wherein the software instructions comprise instructions that:
>
> receive a request to authenticate a client, wherein the request comprises an identifier and a password,
>
> store, in a multidimensional time-series database, information about the request,
>
> determine whether the password corresponds to a user account identified by the identifier,
>
> determine whether an additional verification is required to grant access,
>
> wherein determining whether the additional verification is required to grant access comprises:
>
> retrieving, from the multidimensional time-series database, historical information about previous access requests associated with the user account, and
>
> determining, based at least on the historical information, whether the user account is associated with a previous access request to a network resource, wherein the previous access request to the network resource is anomalous relative to a baseline profile of access requests; and,
>
> based on the additional verification being required to grant access:
>
> select an additional verification method from a plurality of verification methods,

cause the client to be prompted to complete the additional verification method, and

determine whether the additional verification method has been completed correctly

74. By incorporating, relying upon, interacting with, or otherwise utilizing the Accused '934 Functionality, the Accused Entra Products meet every element of this claim. The further descriptions below, which are based on an analysis of publicly available information, are preliminary examples and non-limiting.

75. For example, and to the extent that the preamble is limiting, the Accused Entra Products are computer systems configured to execute software instructions stored on nontransitory machine-readable storage media. For example, Entra ID is a cloud-hosted login system that includes features to enforce and implement multifactor authentication such as Entra ID Protection.[10]

76. On information and belief, the Accused Entra Products execute software instructions that receive a request to authenticate a client, wherein the request comprises an identifier and a password. For example, Entra ID supports various forms of identifiers for users, including email addresses, bank account numbers, and IP addresses.[11] As another example, Entra ID also receives and manages passwords on behalf of users and applications.[12]

77. On information and belief, the Accused Entra Products execute software instructions that store, in a multidimensional time-series database, information about the request, and determine whether the password corresponds to a user account identified by the

[10] https://learn.microsoft.com/en-us/entra/fundamentals/whatis

[11] https://learn.microsoft.com/en-us/entra/fundamentals/identity-fundamental-concepts

[12] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-sspr-howitworks

identifier. For example, Entra ID stores historical data about correct and incorrect password use, among other data, to determine risk profiles for the future.[13]

78. Microsoft's current cloud-based systems rely on several platforms that store and manage multiple dimensions of time-series data. For example, modern versions of Microsoft Fabric, Azure Data Explorer, Azure Monitor, Microsoft Sentinel use the Kusto platform, described as "a powerful tool for exploring your data and discovering patterns, identifying anomalies and outliers, creating statistical modeling, and more."[14] The use of a multi-dimensional time-series database significantly improves the bandwidth and latency of both storing and querying past user behavior. This improves the bandwidth of events that can be stored and the latency on which events can be queried during a login. This can result in faster login times and the use of a larger volume of data than could be consulted in prior art systems. Microsoft's documentation explains that "KQL is used by many other Microsoft services."[15] On information and belief, the Accused Entra Products use a multi-dimensional time series database related to Kusto and KQL.

79. On information and belief, the Accused Entra Products execute software instructions that determine whether additional verification is required by retrieving historical information about previous access requests associated with the user account and using at least that information to determine whether the account is associated with a previous anomalous request to a network resource. For example, Entra ID uses data about the real-time and offline

[13] https://learn.microsoft.com/en-us/entra/id-protection/concept-identity-protection-risks#password-spray

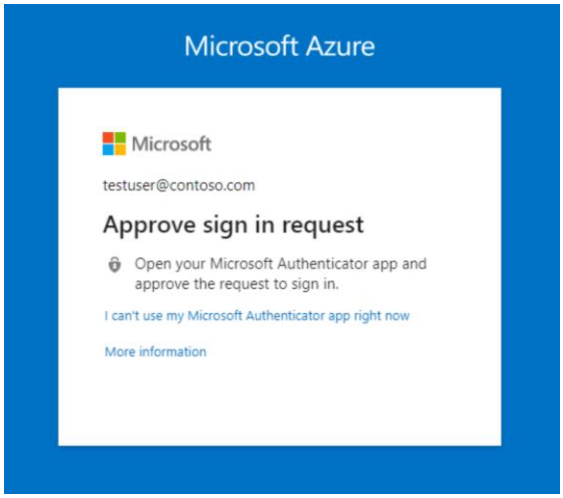
[14] https://learn.microsoft.com/en-us/kusto/query/

[15] *Id.*

behavior of users to determine whether an access request is anomalous and whether to trigger MFA.[16]

80. On information and belief, the Accused Entra Products execute software instructions that, when required, select an additional verification method from a plurality of verification methods. For example, Entra ID supports multiple additional verification methods, including Microsoft Authenticator, Windows Hello for Business, and voice calls.[17]

81. On information and belief, the Accused Entra Products execute software instructions that, when required, cause the client to be prompted to complete the additional verification method. For example, a Microsoft tutorial shows Entra ID prompting a client to complete an additional verification method, as shown.[18]



82. On information and belief, the Accused Entra Products execute software instructions that, when required, determine whether the additional verification method was

[16] https://learn.microsoft.com/en-us/entra/identity/authentication/tutorial-risk-based-sspr-mfa

[17] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-mfa-howitworks

[18] https://learn.microsoft.com/en-us/entra/identity/authentication/tutorial-enable-azure-mfa

completed correctly. For example, Microsoft documentation explains that the entire MFA process is implemented by Entra ID for natively supported methods, meaning that Entra ID determines whether the additional verification method has been completed correctly.[19]

83. Accordingly, as illustrated above, the Accused Entra Products directly infringe one or more claims of the '934 Patent. Microsoft makes, uses, sells, offers for sale, and/or imports, in this District and/or elsewhere in the United States the Accused Entra Products and thus directly infringes the '934 Patent.

84. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the '934 Patent, as provided in 35 U.S.C. § 271(b), including at least by inducing infringement by others, such as Microsoft's customers and end-users, in this District and elsewhere in the United States, to use the Accused Entra Products in manners that infringe the '934 Patent. Microsoft induces such direct infringement through its affirmative acts of making, using, directing an entity to use, selling, offering to sell, and/or importing the Accused Entra Products, as well as by advertising the Accused Entra Products and providing instructions, documentation, and other information to its customers and end-users to encourage and teach them how to use the infringing Accused Entra Products, including but not limited to by Microsoft providing in-store and online technical support, marketing materials, product manuals, advertisements, and other product documentation. At least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the '934 Patent and with the intent, or willful blindness, that the induced acts directly infringe the '934 Patent.

[19] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-mfa-howitworks

85. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the ’934 Patent, as provided by 35 U.S.C. § 271(c), by contributing to direct infringement committed by others, such as Microsoft’s customers and end-users, in this District and elsewhere in the United States. Microsoft’s affirmative acts of selling and offering to sell the Accused Entra Products in this District and elsewhere in the United States, and causing the Accused Entra Products to be manufactured, used, sold, and offered for sale, contribute to Microsoft’s customers and end-users using the Accused Entra Products, such that the ’934 Patent is directly infringed. The accused components in the Accused Entra Products are material to the inventions claimed in the ’934 Patent, are not staple articles or commodities of commerce, have no substantial non-infringing uses, and are known by Microsoft to be especially made or adapted for use in the infringement of the ’934 Patent during at least the post-service period. Similarly, at least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the ’934 Patent and with the intent, or willful blindness, that they cause direct infringement of the ’934 Patent.

86. On information and belief, no licensed patent-practicing commercial products have been made, sold, offered for sale, or imported in a manner that would trigger the requirements of 35 U.S.C. § 287.

87. Microsoft’s infringement of the ’934 Patent has damaged and will continue to damage Qomplx.

88. Upon service of this complaint, Microsoft will have explicit written notice of its infringement of the ’934 Patent. Nevertheless, on information and belief, Microsoft will, without authorization, knowingly, intentionally, purposefully, and deliberately continue to infringe the ’934 Patent.

COUNT III

INFRINGEMENT OF U.S. PATENT NO. 12,231,426

89. Qomplx repeats and incorporates by reference each preceding paragraph as if fully set forth herein and further alleges:

90. The ’426 Patent, entitled “Contextual and Risk-Based Multi-Factor Authentication,” was duly and lawfully issued on February 18, 2025 and assigned to QOMPLX LLC. A true and correct copy of the ’426 Patent is attached hereto as Exhibit 3.

91. The ’426 Patent names Jason Crabtree, Andrew Sellers, and Ian MacLeod as inventors.

92. The ’426 Patent claims priority to, among others, U.S. Provisional No. 62/574,708, filed October 19, 2017.

93. The ’426 Patent has been in full force and effect since its issuance. Qomplx owns all rights to the ’426 Patent that are necessary to bring this action.

94. The ’426 Patent, among other things, states that it “relates to the field of network security, particularly to multi-factor user authentication.” **Ex. 3** at 2:15-16.

95. As the ’426 Patent explains, MFA systems “commonly used today include[] one-time use codes sent to a user’s mobile device or email, confirming through a uniquely generated link sent to the user, or using authenticator devices and apps that generate a code on-demand.” *Id.* at 2:23-27.

96. The ’426 Patent addresses a key “fault” in those systems: “over-reliance on a single method of delivery.” *Id.* at 2:28-29. For example, “once a user’s email is compromised, it may be trivial to gain access to that user’s other accounts” because “password reset” requests “are usually conducted through the user’s email” and “[o]ne time codes, such as those from banking websites, are also often sent to the user’s email address.” *Id.* at 2:29-34.

97. Departing from the prior art's limited and ineffective approach to authentication, the claims of the '426 Patent are directed to specific verification methods that improve computer security.

98. For example, Claim 1 identifies a specific technique by which a computer system may "determine whether an additional verification is required to grant access." Among other things, a computer system implementing the claimed authentication technology can, upon "receiv[ing] a request to authenticate a client" that "comprises a first identifier and a password" corresponding to a first user account and "storing, in a multidimensional time-series database, information about the request," retrieve "historical information about previous access requests associated with the first user account" from the multidimensional time-series database and use that information to determine whether the first user account is associated with a "previous access request to authenticate" that "comprised a second identifier not associated with the first user account" and, if so, select from multiple verification methods to authenticate the client.

99. Accordingly, the authentication techniques claimed in the '426 Patent, including the particular sequence of steps recited in the claims to determine whether additional authentication is needed, "eliminate[s]" MFA's "over-reliance on a single, and possibly compromised, method" by using "a combination of verification methods" and "dynamically determin[ing] the varying amounts of verification needed, based on the contexts and risks associated with the connection," yielding higher levels of security and improving computer functionality. *Id.* at 2:35-40.

100. As the specification explains, the contextual and risk-based approach to authentication claimed in the '426 Patent allows a computer security system in certain embodiments to "intelligently integrate the large volume of data from a plurality of sources

on an ongoing basis" and identify anomalous and potentially threatening network activity by "continuously poll[ing] the incoming traffic data for activities anomalous to [network] baseline." *Id.* at 8:23-35. These embodiments are able to react with less latency to attacks. They can also process a larger bandwidth of input data, resulting in the detection of a broader base of potential attack types. Both of these improvements to computer performance result in blocking accesses that prior art systems could not block. The specification provides numerous examples of such anomalous activity, including "a user a user attempting to gain access several network resources such as workstations or servers in rapid succession, or a user attempting to gain access to a domain server of server with sensitive information using random userIDs or another user's userID and password, or attempts by any user to brute force crack a privileged user's password," among others. *Id.* at 8:35-46. Through the storage and retrieval of historical access information about the first user account, and the application of additional verification methods to an incoming authentication request based on whether the first user account is associated with a prior authentication request, the claims of the '426 Patent improve on the prior art's approach to authentication technology, provide tangible improvements to network security, and mitigate the risk of cyberattack or data loss.

101. The '426 Patent's improvements to network security are further described in the specification.

102. Microsoft is not currently licensed to practice the '426 Patent.

103. Qomplx is informed and believes, and thereon alleges, that Microsoft has infringed and continues to infringe one or more claims of the '426 Patent in violation of 35 U.S.C. § 271, either literally and/or under the doctrine of equivalents, by making, using, directing an entity to use, selling, and/or offering for sale in the United States, and/or importing into the United States, without authorization, Microsoft products that practice one

of more claims of the ’426 Patent, including without limitation products that incorporate, rely upon, interact with, or otherwise utilize Microsoft Entra ID (“Entra ID” or the “Accused ’426 Functionality”).

104. For example and without limitation, products that incorporate, rely upon, interact with, or otherwise utilize Entra ID, including at least the Microsoft Entra product family (collectively, the “Accused Entra Products”) embody every limitation of at least claim 1 of the ’426 Patent, both literally and under the doctrine of equivalents, as set forth below.

105. Claim 1 of the ’426 Patent provides:

> A computer system configured to execute software instructions stored on nontransitory machine-readable storage media, wherein the software instructions comprise instructions that:
> receive a request to authenticate a client, wherein the request comprises a first identifier and a password,
> store, in a multidimensional time-series database, information about the request,
> determine whether the password corresponds to a first user account identified by the first identifier,
> determine whether an additional verification is required to grant access,
> wherein determining whether the additional verification is required to grant access comprises:
> retrieving, from the multidimensional time-series database, historical information about previous access requests associated with the first user account, and
> determining, based at least on the historical information, whether the first user account is associated with a previous request to authenticate, wherein the previous request to authenticate comprised a second identifier not associated with the first user account; and,
> based on the additional verification being required to grant access:
> select an additional verification method from a plurality of verification methods,
> cause the client to be prompted to complete the additional verification method, and
> determine whether the additional verification method has been completed correctly

106. By incorporating, relying upon, interacting with, or otherwise utilizing the Accused ’426 Functionality, the Accused Entra Products meet every element of this claim.

The further descriptions below, which are based on an analysis of publicly available information, are preliminary examples and non-limiting.

107. For example, and to the extent that the preamble is limiting, the Accused Entra Products are computer systems configured to execute software instructions stored on nontransitory machine-readable storage media. For example, Entra ID is a cloud-hosted login system that includes features to enforce and implement multifactor authentication such as Entra ID Protection.[20]

108. On information and belief, the Accused Entra Products execute software instructions that receive a request to authenticate a client, wherein the request comprises a first identifier and a password. For example, Entra ID supports various forms of identifiers for users, including email addresses, bank account numbers, and IP addresses.[21] As another example, Entra ID also receives and manages passwords on behalf of users and applications.[22]

109. On information and belief, the Accused Entra Products execute software instructions that store, in a multidimensional time-series database, information about the request, and determine whether the password corresponds to a first user account identified by the first identifier. For example, Entra ID stores historical data about correct and incorrect password use, among other data, to determine risk profiles for the future.[23]

110. Microsoft's current cloud-based systems rely on several platforms that store and manage multiple dimensions of time-series data. For example, modern versions of

[20] https://learn.microsoft.com/en-us/entra/fundamentals/whatis

[21] https://learn.microsoft.com/en-us/entra/fundamentals/identity-fundamental-concepts

[22] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-sspr-howitworks

[23] https://learn.microsoft.com/en-us/entra/id-protection/concept-identity-protection-risks#password-spray

Microsoft Fabric, Azure Data Explorer, Azure Monitor, Microsoft Sentinel use the Kusto platform, described as "a powerful tool for exploring your data and discovering patterns, identifying anomalies and outliers, creating statistical modeling, and more.[24] "The use of a multi-dimensional time-series database significantly improve the bandwidth and latency of both storing and querying past user behavior. This improves the bandwidth of events that can be stored and the latency on which events can be queried during a login. This can result in faster login times and the use of a larger volume of data than could be consulted in prior art systems. Microsoft's documentation explains that "KQL is used by many other Microsoft services."[25] On information and belief, the Accused Entra Products use a multi-dimensional time series database related to Kusto and KQL.

111. On information and belief, the Accused Entra Products execute software instructions that determine whether additional verification is required by retrieving historical information about previous access requests associated with the first user account and using at least that information to determine whether the first user account is associated with a previous request to authenticate, which comprised a second identifier not associated with the first user account. For example, Entra ID uses data about the real-time and offline behavior of users to determine whether an access request is suspicious and thus whether to trigger MFA.[26]

112. On information and belief, the Accused Entra Products execute software instructions that, when required, select an additional verification method. For example, Entra

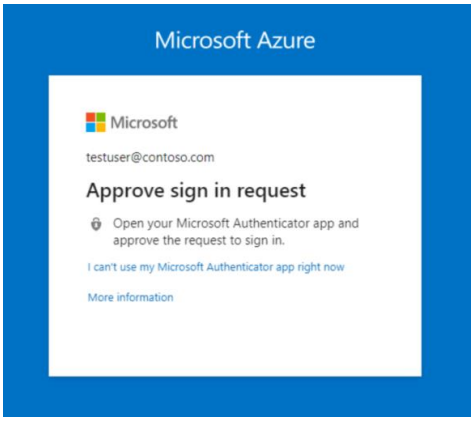
[24] https://learn.microsoft.com/en-us/kusto/query/

[25] *Id.*

[26] https://learn.microsoft.com/en-us/entra/identity/authentication/tutorial-risk-based-sspr-mfa

ID supports multiple additional verification methods, including for example, Microsoft Authenticator, Windows Hello for Business, and voice calls.[27]

113. On information and belief, the Accused Entra Products execute software instructions that, when required, cause the client to be prompted to complete the additional verification method. For example, many of the additional verification methods supported by Entra ID cause the client to be prompted to complete verification, as shown.[28]



114. On information and belief, the Accused Entra Products execute software instructions that, when required, determine whether the additional verification method was completed correctly. For example, Microsoft documentation explains that Entra ID implements the entire MFA process for natively supported methods, meaning that Entra ID determines whether the additional verification method has been completed correctly.[29]

115. Accordingly, as illustrated above, the Accused Entra Products directly infringe one or more claims of the ’426 Patent. Microsoft makes, uses, sells, offers for sale, and/or

[27] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-mfa-howitworks

[28] https://learn.microsoft.com/en-us/entra/identity/authentication/tutorial-enable-azure-mfa

[29] https://learn.microsoft.com/en-us/entra/identity/authentication/concept-mfa-howitworks

imports, in this District and/or elsewhere in the United States the Accused Entra Products and thus directly infringes the ’426 Patent.

116. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the ’426 Patent, as provided in 35 U.S.C. § 271(b), including at least by inducing infringement by others, such as Microsoft’s customers and end-users, in this District and elsewhere in the United States, to use the Accused Entra Products in manners that infringe the ’426 Patent. Microsoft induces such direct infringement through its affirmative acts of making, using, directing an entity to use, selling, offering to sell, and/or importing the Accused Entra Products, as well as by advertising the Accused Entra Products and providing instructions, documentation, and other information to its customers and end-users to encourage and teach them how to use the infringing Accused Entra Products, including but not limited to by Microsoft providing in-store and online technical support, marketing materials, product manuals, advertisements, and other product documentation. At least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the ’426 Patent and with the intent, or willful blindness, that the induced acts directly infringe the ’426 Patent.

117. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the ’426 Patent, as provided by 35 U.S.C. § 271(c), by contributing to direct infringement committed by others, such as Microsoft’s customers and end-users, in this District and elsewhere in the United States. Microsoft’s affirmative acts of selling and offering to sell the Accused Entra Products in this District and elsewhere in the United States, and causing the Accused Entra Products to be manufactured, used, sold, and offered for sale, contribute to Microsoft’s customers and end-users using the Accused Entra Products, such that the ’426 Patent is directly infringed. The

accused components in the Accused Entra Products are material to the inventions claimed in the ’426 Patent, are not staple articles or commodities of commerce, have no substantial non-infringing uses, and are known by Microsoft to be especially made or adapted for use in the infringement of the ’426 Patent during at least the post-service period. Similarly, at least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the ’426 Patent and with the intent, or willful blindness, that they cause direct infringement of the ’426 Patent.

118. On information and belief, no licensed patent-practicing commercial products have been made, sold, offered for sale, or imported in a manner that would trigger the requirements of 35 U.S.C. § 287.

119. Microsoft’s infringement of the ’426 Patent has damaged and will continue to damage Qomplx.

120. Upon service of this complaint, Microsoft will have explicit written notice of its infringement of the ’426 Patent. Nevertheless, on information and belief, Microsoft will, without authorization, knowingly, intentionally, purposefully, and deliberately continue to infringe the ’426 Patent.

COUNT IV

INFRINGEMENT OF U.S. PATENT NO. 12,301,627

121. Qomplx repeats and incorporates by reference each preceding paragraph as if fully set forth herein and further alleges:

122. The ’627 Patent, entitled “Correlating Network Event Anomalies Using Active and Passive External Reconnaissance to Identify Attack Information,” was duly and lawfully issued on May 13, 2025 and assigned to QOMPLX LLC. A true and correct copy of the ’627 Patent is attached hereto as Exhibit 4.

123. The ’627 Patent names Jason Crabtree, Andrew Sellers, and Richard Kelley as inventors.

124. The ’627 Patent claims priority to, among others, U.S. Application No. 17/237,346, filed April 22, 2021.

125. The ’627 Patent has been in full force and effect since its issuance. Qomplx owns all rights to the ’627 Patent that are necessary to bring this action.

126. The ’627 Patent, among other things, states that it “relates to the field of computer management, and more particularly to the field of cybersecurity and threat analytics.” **Ex. 4** at 2:28-30.

127. As the ’627 Patent explains, “[u]nderstanding the cybersecurity profile of an organization is a complex endeavor, and” the complexity “increases exponentially with the size of the organization” because “each component of the organization’s network connects to multiple other components.” *Id.* at 2:34-39.

128. The ’627 Patent further explains that prior art “cybersecurity rating methods (for example, the Common Vulnerability Scoring Systems, or CVSS) fail to adequately profile and rate the cybersecurity profiles of organizations because they do not incorporate sufficient information.” *Id.* at 2:40-44.

129. The ’627 Patent overcomes the technological limitations of the prior art by providing “a system and method for correlating network event anomalies . . . that can identify attack patterns and points of origin based on observed anomalies and their relationships to other observed network events.” *Id.* at 2:52-57.

130. For example, Claim 1 of the ’627 Patent describes a computer system that, among other things, represents a plurality of network entities and relationships as a directed graph, analyzes streaming data relating to one or more of those entities to identify an entity

and a relationship that do not correspond to the plurality, updates the graph to account for that entity and relationship, identifies a potential attack path involving an entity by identifying additional entities that can be reached by using the first entity and generates a report identifying the first entity and an entity that can be reached by using the first entity. As the specification explains, this approach "describes the interconnectedness between nodes and the various event pathways that can occur, including potential (or actual) attac[k] paths that an intruder may take advantage of when attempting to compromise the network or any particular nodes," allowing analysis of "what dependencies must be satisfied for any given event or node behavior, what the effects of the event or behavior are, what nodes are in turn affected by those outcomes, and other such directed flow information." *Id.* at 22:3-12. This, in turn, enables the system "to identify any possible event flows and interactions between the affected nodes, indicating causative effect pathways that are triggered by events and node behaviors, which in the case of an anomaly allows insight into what effects the anomaly may have or what may have led up to the event in question." *Id.* at 22:12-17.

131. Accordingly, the systems claimed by the '627 Patent, including as described in Claim 1, represent discrete improvements in network security that "enable[] more effective mitigation of attacks and resolution of flaws in a network." *Id.* at 22:35-36. Unlike prior methods, which "focus on current symptoms of an attack," the architectures claimed by the '627 Patent allow computing systems to "directly identify and address the root cause of the issue." *Id.* at 22:40-42.

132. For example, by representing on a directed graph an attack path involving an entity that could be involved in an attack on a network, the claims of the '627 Patent make it "possible to determine a point of origin . . . and the starting conditions that were in place," which "enables more effective mitigation of attacks and resolution of flaws in a network, by

identifying the precise causes of vulnerabilities and deficiencies so they may be addressed directly," improving on the prior art's "traditional approaches that are inherently limited to mitigating the outcomes of an unknown attack" and "focus on current symptoms of an attack." *Id.* at 22:33-42. As another example, by identifying the entities, such as accounts and resources, that can be reached by using the entity that may be involved in an attack, the systems claimed in the '627 Patent can "analyze a user account and identify its access capabilities" such as "what files, directories, devices or domains an account may have access to," which "may be used to produce a 'blast radius' calculation" that "identif[ies] exactly what resources are at risk as a result of [a potential] intrusion and where security personnel should focus their attention." *Id.* at 23:17-27.

133. Thus, by producing a "meaningful and contextualized visualization of a security infrastructure that reflects the current state of the internal relationships present in the infrastructure," the claims of the '627 Patent enable more efficient network defense and more efficient ways of using computers. *Id.* at 27:67-28:3. The '627 Patent claims a specific mechanism by which the computer system can store this contextualized representation of the security infrastructure inside an organization. This mechanism allows for a computer system to represent and continuously update a representation of a larger organization with reduced latency than prior art systems. The '627 Patent also claims a specific mechanism for simulating attacks within this representation. This mechanism allows for faster graph traversal, also allowing for the management of larger graphs with reduced latency. Each of these improvements to computer systems improve computer security by providing a more complete understanding of potential attack paths, their reach, and the consequences of any such attack, using a dynamic system map that is regularly updated rather than a static (and potentially outdated) understanding of the system.

134. The ’627 Patent’s improvements to computer network defense systems are further described in the specification.

135. Microsoft is not currently licensed to practice the ’627 Patent.

136. Qomplx is informed and believes, and thereon alleges, that Microsoft has infringed and continues to infringe one or more claims of the ’627 Patent in violation of 35 U.S.C. § 271, either literally and/or under the doctrine of equivalents, by making, using, directing an entity to use, selling, and/or offering for sale in the United States, and/or importing into the United States, without authorization, Microsoft products that practice one of more claims of the ’627 Patent, including without limitation products that incorporate, rely upon, interact with, or otherwise utilize Microsoft Security Exposure Management (“MSEM” or the “Accused ’627 Functionality”).

137. For example and without limitation, products that incorporate, rely upon, interact with, or otherwise utilize MSEM (collectively, the “Accused ’627 Products”) embody every limitation of at least claim 1 of the ’627 Patent, both literally and under the doctrine of equivalents, as set forth below.

138. Claim 1 of the ’627 Patent provides:

> 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:
> store in the hardware memory a representation of a first graph, wherein the representation of the first graph comprises representations of a first plurality of nodes corresponding to a first plurality of entities and further comprises representations of a first plurality of edges, wherein the first graph is a directed graph,
> wherein the first plurality of entities comprises a plurality of accounts and a plurality of resources, and
> wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities;
> 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,

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,

identify, based on the modified representation of the first graph, an attack path that could be involved in an attack involving the first entity, wherein identifying the attack path comprises:

identifying a second entity that can be reached using the first entity, wherein the second entity corresponds to a second node, and the second node is related by one or more edges to the first node corresponding to the first entity in the modified representation of the first graph; and,

identifying a third entity that can be reached using the second entity, wherein the third entity corresponds to a third node, and the third node is related by one or more edges to the second node in the modified representation of the first graph; and

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

139. By incorporating, relying upon, interacting with, or otherwise utilizing the Accused ’627 Functionality, the Accused ’627 Products meet every element of this claim. The further descriptions below, which are based on an analysis of publicly available information, are preliminary examples and non-limiting.

140. For example, and to the extent that the preamble is limiting, the Accused ’627 Products are, on information and belief, computer systems with a hardware memory configured to execute software instructions stored on nontransitory machine-readable storage

media. For example, MSEM is a cloud software platform that executes instructions stored on nontransitory machine-readable storage.[30]

141. On information and belief, the Accused '627 Products execute software instructions that store in the hardware memory a representation of a first graph, comprising representations of a first plurality of nodes corresponding to a first plurality of entities, and representations of a first plurality of edges. For example, MSEM stores graphs, including the "exposure graph" consisting of nodes and edges stored in a pair of tables.[31]

Schema tables

The exposure graph relies on the following tables:

- *ExposureGraphNodes*
- *ExposureGraphEdges*

142. On information and belief, the Accused '627 Products execute software instructions that store the first graph as a directed graph. For example, Microsoft's website explains that MSEM stores a directed graph with edges that connect sources to targets.[32]

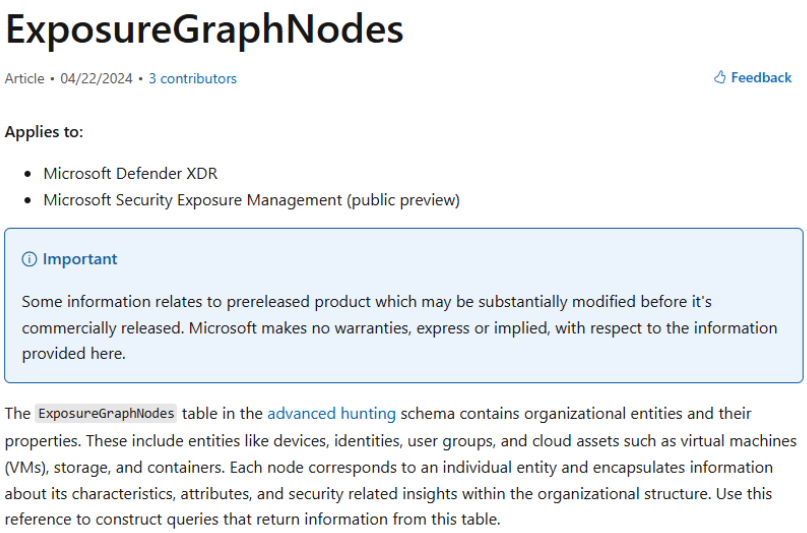
- `SourceNodeId` (`string`) - Node ID of the edge's source. Example: "12346aa0-10a5-587e-52f4-280bfc014a08"
- `SourceNodeName` (`string`) - The source node display name. Example: "mdvmaas-win-123"
- `SourceNodeLabel` (`string`) - The source node label. Example: "microsoft.compute/virtualmachines"
- `SourceNodeCategories` (`Dynamic` (json)) - The categories list of the source node.
- `TargetNodeId` (`string`) - The node ID of the edge's target. Example: "45676aa0-10a5-587e-52f4-280bfc014a08"
- `TargetNodeName` (`string`) - Display name of the target node. Example: gke-test-cluster-1
- `TargetNodeLabel` (`string`) - The target node label. Example: "compute.instances"
- `TargetNodeCategories` (`Dynamic` (json)) - The categories list of the target node.

[30] https://learn.microsoft.com/en-us/security-exposure-management/microsoft-security-exposure-management

[31] https://learn.microsoft.com/en-us/security-exposure-management/schemas-operators

[32] https://learn.microsoft.com/en-us/security-exposure-management/schemas-operators

143. On information and belief, the Accused '627 Products execute software instructions wherein a first plurality of entities comprises a plurality of accounts and a plurality of resources. For example, Microsoft's website explains that MSEM stores graphs in which the nodes include identities and assets.[33]



ExposureGraphNodes

Article • 04/22/2024 • 3 contributors Feedback

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 `ExposureGraphNodes` 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.

144. On information and belief, the Accused '627 Products execute software instructions wherein each edge of the first plurality of edges corresponds to a respective relationship between a respective pair of entities. For example, Microsoft's website explains that MSEM stores graphs wherein edges in the graph are labeled, describing the relationship by which the entities are connected.[34]

[33] https://learn.microsoft.com/en-us/defender-xdr/advanced-hunting-exposuregraphnodes-table

[34] https://learn.microsoft.com/en-us/defender-xdr/advanced-hunting-exposuregraphedges-table

Column name	Data type	Description
EdgeId	string	Unique identifier for the relationship/edge
EdgeLabel	string	The edge label like "routes traffic to"
SourceNodeId	string	Node ID of the edge's source
SourceNodeName	string	Source node display name
SourceNodeLabel	string	Source node label
SourceNodeCategories	dynamic	Categories list of the source node in JSON format
TargetNodeId	string	Node ID of the edge's target
TargetNodeName	string	Display name of the target node
TargetNodeLabel	string	Target node label
TargetNodeCategories	dynamic	The categories list of the target node in JSON format
EdgeProperties	dynamic	Optional data relevant for the relationship between the nodes in JSON format

145. On information and belief, Microsoft saves this graph in a system based on Kusto and the Kusto Query Language. The infringing use of such a graph database significantly improves the bandwidth of information that can be stored and the latency in which events can be queried.[35]

Graph Kusto Query Language (KQL) operators

Microsoft Security Exposure Management relies on exposure graph tables and unique exposure graph operators to enable operations over graph structures. The graph is built from tabular data using the `make-graph` operator, and then queried using graph operators.

146. On information and belief, the Accused '627 Products execute software instructions that receive streaming data comprising updates about events relating to entities of the first plurality. For example, Microsoft's website explains that MSEM collects data from a large set of sources, including Microsoft Defender and Microsoft Entra ID, among others.[36]

147. On information and belief, the Accused '627 Products execute software instructions that use non-identical portions of the streaming data to identify a first relationship

[35] https://learn.microsoft.com/en-us/security-exposure-management/schemas-operators#graph-kusto-query-language-kql-operators

[36] https://learn.microsoft.com/en-us/security-exposure-management/overview-data-connectors

between a pair of entities of the first plurality, not corresponding to any of the first plurality of edges, and modify the representation of the first graph accordingly. For example, Microsoft's website explains that MSEM is continuously updated into a "unified and up-to-date view of your inventory and attack surface."[37]

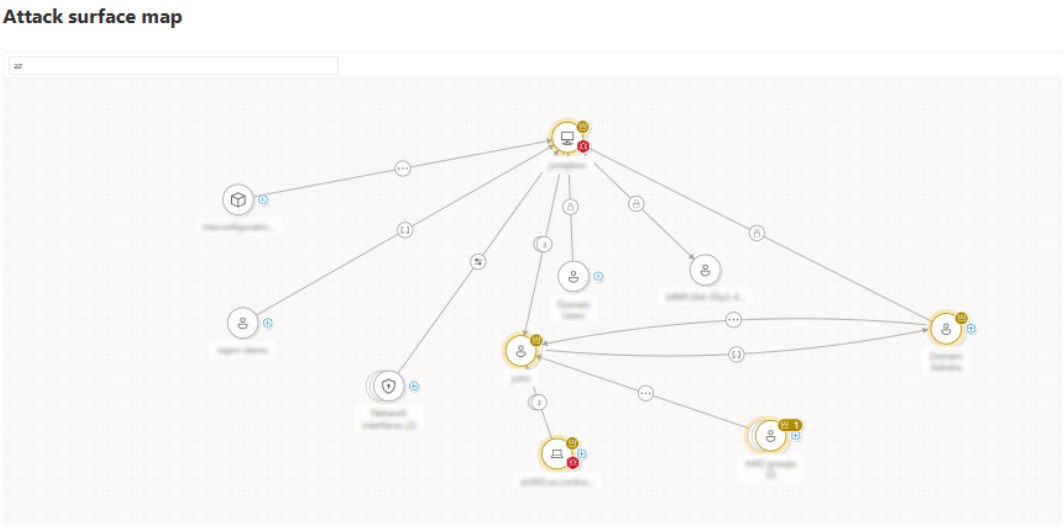
148. On information and belief, the Accused '627 Products execute software instructions that identify, based on the first graph's modified representation, an attack path that could be involved in an attack involving the first entity. For example, Microsoft's website explains that MSEM uses the updated graph to compute attack paths.[38]

149. On information and belief, the Accused '627 Products execute software instructions that identify second and third entities, corresponding to second and third nodes, which are related by at least one edge to the first and second nodes/entities, respectively. For example, Microsoft's website shows that MSEM identifies attack paths that can be many levels deep. This identifies and displays how different nodes may be reached along a path.[39]

[37] https://learn.microsoft.com/en-us/security-exposure-management/microsoft-security-exposure-management#what-can-i-do-with-security-exposure-management

[38] https://learn.microsoft.com/en-us/security-exposure-management/critical-asset-management

[39] *Id.*



150. On information and belief, the Accused ’627 Products execute software instructions that generate reports identifying the first entity and at least one of the second and third entities. For example, Microsoft’s website shows that MSEM generates reports on these attack paths.[40]

151. Accordingly, as illustrated above, the Accused ’627 Products directly infringe one or more claims of the ’627 Patent. Microsoft makes, uses, sells, offers for sale, and/or imports, in this District and/or elsewhere in the United States the Accused ’627 Products and thus directly infringes the ’627 Patent.

152. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the ’627 Patent, as provided in 35 U.S.C. § 271(b), including at least by inducing infringement by others, such as Microsoft’s customers and end-users, in this District and elsewhere in the United States, to use the Accused ’627 Products in manners that infringe the ’627 Patent. Microsoft induces such direct infringement through its affirmative acts of making, using, directing an entity to use, selling, offering to sell, and/or importing the Accused ’627 Products, as well as by advertising

[40] *Id.*

the Accused ’627 Products and providing instructions, documentation, and other information to its customers and end-users to encourage and teach them how to use the infringing Accused ’627 Products, including but not limited to by Microsoft providing in-store and online technical support, marketing materials, product manuals, advertisements, and other product documentation. At least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the ’627 Patent and with the intent, or willful blindness, that the induced acts directly infringe the ’627 Patent.

153. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the ’627 Patent, as provided by 35 U.S.C. § 271(c), by contributing to direct infringement committed by others, such as Microsoft’s customers and end-users, in this District and elsewhere in the United States. Microsoft’s affirmative acts of selling and offering to sell the Accused ’627 Products in this District and elsewhere in the United States, and causing the Accused ’627 Products to be manufactured, used, sold, and offered for sale, contribute to Microsoft’s customers and end-users using the Accused ’627 Products, such that the ’627 Patent is directly infringed. The accused components in the Accused ’627 Products are material to the inventions claimed in the ’627 Patent, are not staple articles or commodities of commerce, have no substantial non-infringing uses, and are known by Microsoft to be especially made or adapted for use in the infringement of the ’627 Patent during at least the post-service period. Similarly, at least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the ’627 Patent and with the intent, or willful blindness, that they cause direct infringement of the ’627 Patent.

154. On information and belief, no licensed patent-practicing commercial products have been made, sold, offered for sale, or imported in a manner that would trigger the requirements of 35 U.S.C. § 287.

155. Microsoft's infringement of the '627 Patent has damaged and will continue to damage Qomplx.

156. Upon service of this complaint, Microsoft will have explicit written notice of its infringement of the '627 Patent. Nevertheless, on information and belief, Microsoft will, without authorization, knowingly, intentionally, purposefully, and deliberately continue to infringe the '627 Patent.

COUNT V

INFRINGEMENT OF U.S. PATENT NO. 12,301,628

157. Qomplx repeats and incorporates by reference each preceding paragraph as if fully set forth herein and further alleges:

158. The '628 Patent, entitled "Correlating Network Event Anomalies Using Active and Passive External Reconnaissance to Identify Attack Information," was duly and lawfully issued on May 13, 2025 and assigned to QOMPLX LLC. A true and correct copy of the '628 Patent is attached hereto as Exhibit 5.

159. The '628 Patent names Jason Crabtree, Andrew Sellers, and Richard Kelley as inventors.

160. The '628 Patent claims priority to, among others, U.S. Application No. 17/237,346, filed April 22, 2021.

161. The '628 Patent has been in full force and effect since its issuance. Qomplx owns all rights to the '628 Patent that are necessary to bring this action.

162. The ’628 Patent, among other things, states that it “relates to the field of computer management, and more particularly to the field of cybersecurity and threat analytics.” **Ex. 5** at 2:27-29.

163. As the ’628 Patent explains, “[u]nderstanding the cybersecurity profile of an organization is a complex endeavor, and” the complexity “increases exponentially with the size of the organization” because “each component of the organization’s network connects to multiple other components.” *Id.* at 2:33-38

164. The ’628 Patent further explains that prior art “cybersecurity rating methods (for example, the Common Vulnerability Scoring Systems, or CVSS) fail to adequately profile and rate the cybersecurity profiles of organizations because they do not incorporate sufficient information.” *Id.* at 2:39-43.

165. The ’628 Patent overcomes the technological limitations of the prior art by providing “a system and method for correlating network event anomalies . . . that can identify attack patterns and points of origin based on observed anomalies and their relationships to other observed network events.” *Id.* at 2:51-56.

166. For example, Claim 1 of the ’628 Patent describes a computer system that, among other things, represents a plurality of network entities and relationships as a directed graph, analyzes streaming data about events relating to one or more of those entities to identify an entity and a relationship that do not correspond to the plurality, modifies the graph to account for that entity and relationship, performs a series of correlations and generates a representation of a second graph to identify nodes and edges that represent one or more event flows that could be involved in a cybersecurity attack, and generates a report comprising information associated with the one or more event flows. As the specification explains, this approach to network anomaly detection and response “describes the interconnectedness

between nodes and the various event pathways that can occur, including potential (or actual) attac[k] paths that an intruder may take advantage of when attempting to compromise the network or any particular nodes," allowing analysis of "what dependencies must be satisfied for any given event or node behavior, what the effects of the event or behavior are, what nodes are in turn affected by those outcomes, and other such directed flow information." *Id.* at 22:6-15. This, in turn, enables the system "to identify any possible event flows and interactions between the affected nodes, indicating causative effect pathways that are triggered by events and node behaviors, which in the case of an anomaly allows insight into what effects the anomaly may have or what may have led up to the event in question." *Id.* at 22:15-20.

167. Accordingly, the systems claimed by the '628 Patent, including as described in Claim 1, represent discrete improvements in network security that "enable[] more effective mitigation of attacks and resolution of flaws in a network." *Id.* at 22:38-39. The systems claimed by the '628 Patent, including as described in Claim 1, similarly represent a specific mechanism by which the computer system can store this contextualized representation of the security infrastructure inside an organization. This mechanism allows for a computer system to represent and continuously update a representation of a larger organization with reduced latency than prior art systems. The '628 Patent also claims a specific mechanism for simulating attacks within this representation. This mechanism allows for faster graph traversal, also allowing for the management of larger graphs with reduced latency. Each of these improvements to computer systems, unlike prior methods, which "focus on current symptoms of an attack," allow computing systems to "directly identify and address the root cause of the issue." *Id.* at 22:43-45.

168. For example, by "identify[ing] possible event flows between affected nodes," the claims of the '628 Patent provide "further insight into [a network] anomaly and any

possible effects or outcomes it may produce." *Id.* at 21:64-67. By representing "an identified event anomaly" on a directed graph, the systems claimed in the '628 Patent "may be used to determine event dependencies based on the identified behavioral patterns and causative relationships in the graph," including by identifying "causative events, behaviors, and conditions that led up to [a network] anomaly." *Id.* at 22:24-30. This enables "following the path of event flows back in time . . . to determine a point of origin for the anomaly and the starting conditions that were in place," which in turn "enables more effective mitigation of attacks and resolution of flaws in a network, by identifying the precise causes of vulnerabilities and deficiencies so they may be addressed directly." *Id.* at 22:33-41. As the specification explains, this approach is an improvement over the prior art's "traditional approaches that are inherently limited to mitigating the outcomes of an unknown attack" and "focus on current symptoms of an attack." *Id.* at 22:41-45.

169. Thus, by producing a "meaningful and contextualized visualization of a security infrastructure that reflects the current state of the internal relationships present in the infrastructure," the claims of the '628 Patent enable more efficient network defense and more efficient ways of using computers. *Id.* at 28:1-28:4. The '628 Patent's claims also improve computer security by providing a more complete understanding of potential event flows, their reach, and the consequences of any such an attack, using a dynamic system map that is regularly updated rather than a static (and potentially outdated) understanding of the system.

170. The '628 Patent's improvements to computer network defense systems are further described in the specification.

171. Microsoft is not currently licensed to practice the '628 Patent.

172. Qomplx is informed and believes, and thereon alleges, that Microsoft has infringed and continues to infringe one or more claims of the '628 Patent in violation of 35

U.S.C. § 271, either literally and/or under the doctrine of equivalents, by making, using, directing an entity to use, selling, and/or offering for sale in the United States, and/or importing into the United States, without authorization, Microsoft products that practice one of more claims of the '628 Patent, including without limitation products that incorporate, rely upon, interact with, or otherwise utilize Microsoft Fusion ("Fusion" or the "Accused '628 Functionality").

173. For example and without limitation, products that incorporate, rely upon, interact with, or otherwise utilize Fusion, including at least Microsoft Sentinel ("Sentinel") and Microsoft Defender (collectively, the "Accused Fusion Products") embody every limitation of at least claim 1 of the '628 Patent, both literally and under the doctrine of equivalents, as set forth below.

174. Claim 1 of the '628 Patent provides:

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

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 edge;

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 modified representation of the first graph;

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 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;

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 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.

175. By incorporating, relying upon, interacting with, or otherwise utilizing the Accused ’628 Functionality, the Accused Fusion Products meet every element of this claim. The further descriptions below, which are based on analysis of publicly available information, are preliminary examples and are non-limiting.

176. For example, and to the extent that the preamble is limiting, the Accused Fusion Products are, on information and belief, computer systems with hardware memory configured to execute software instructions stored on nontransitory machine-readable storage

media. For example, Fusion is a cloud-based software platform that executes instructions stored on nontransitory storage.[41]

177. On information and belief, the Accused Fusion Products execute software instructions that store in the hardware memory a representation of a first graph, comprising representations of a first plurality of nodes corresponding to a first plurality of entities, and representations of a first plurality of edges. For example, Microsoft's website explains that Fusion builds a graph containing nodes and edges.[42]

> **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.

178. On information and belief, the Accused Fusion Products execute software instructions that store the first graph as a directed graph. For example, Microsoft's website explains that Fusion shows directions on the graph in visualizations of expansion.[43] This, for example, indicates that the graph is a directed graph.

[41] https://learn.microsoft.com/en-us/azure/sentinel/fusion

[42] https://techcommunity.microsoft.com/blog/microsoftsentinelblog/behind-the-scenes-the-ml-approach-for-detecting-advanced-multistage-attacks-with/3239236

[43] *Id.*

- ***Run probabilistic random walk***: 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.

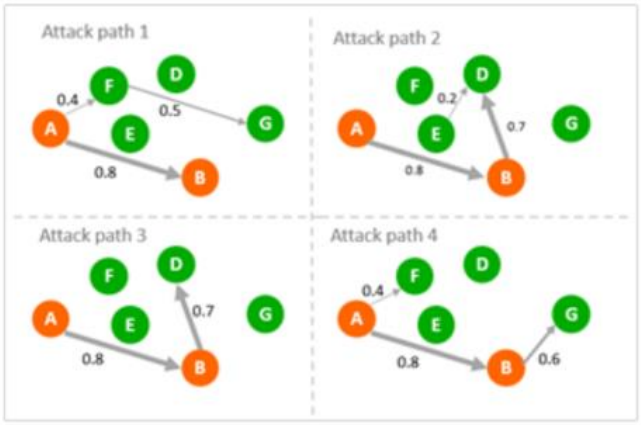


Figure 4: Expansion - probabilistic random walk

179. On information and belief, the Accused Fusion Products execute software instructions wherein a first plurality of entities comprises a plurality of accounts and a plurality of resources, and 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. For example, Microsoft's website explains that Fusion stores graphs in which nodes represent entities, including at least accounts and resources, while edges represent relationships between these nodes and entities.[44]

180. On information and belief, the Accused Fusion Products execute software instructions that receive streaming data comprising updates about events relating to entities of the first plurality, and use non-identical portions of the streaming data to identify a first relationship between a pair of entities of the first plurality, not corresponding to any of the first plurality of edges, and modify the representation of the first graph accordingly. For example, Microsoft's website explains that Fusion receives streaming data about the entities

[44] https://techcommunity.microsoft.com/blog/microsoftsentinelblog/behind-the-scenes-the-ml-approach-for-detecting-advanced-multistage-attacks-with/3239236

in the graph, identifying nodes and edges that are not in the graph.[45] This information is used to update the graph in such a way that the graph is continually updated.[46]

181. On information and belief, for an anomalous event associated with a node in the modified representation of the first graph, the Accused Fusion products execute software instructions that identify a first plurality of correlated nodes, in which each node corresponds to a respective event/resource associated with the anomalous event, and is connected by an 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, Microsoft's website explains that starting with events in the continually updated graph, Fusion correlates a group of connected nodes with an attack pattern.[47]

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.

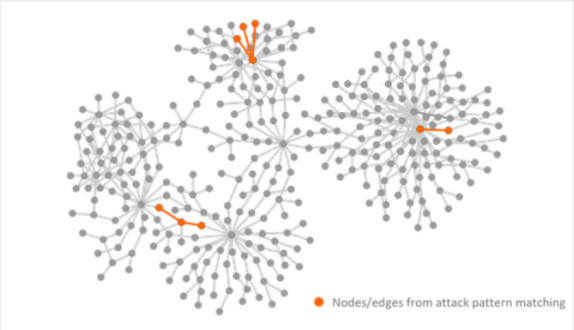


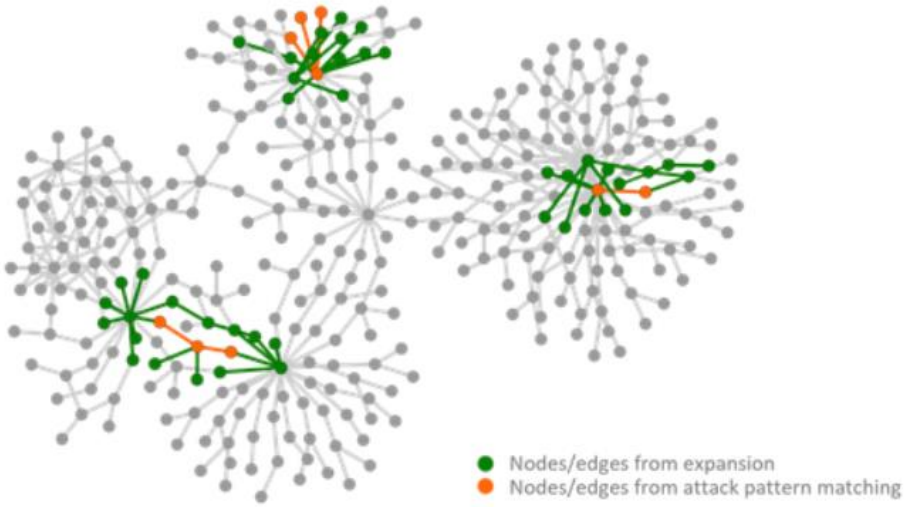
Figure 2: Simplified graph shows nodes and edges from attack pattern matching

[45] https://learn.microsoft.com/en-us/azure/sentinel/soc-ml-anomalies

[46] https://techcommunity.microsoft.com/blog/microsoftsentinelblog/behind-the-scenes-the-ml-approach-for-detecting-advanced-multistage-attacks-with/3239236

[47] https://techcommunity.microsoft.com/blog/microsoftsentinelblog/behind-the-scenes-the-ml-approach-for-detecting-advanced-multistage-attacks-with/3239236

182. On information and belief, the Accused Fusion Products execute software instructions that identify, based on the first graph's modified representation, an attack path that could be involved in an attack involving the first entity. For example, Microsoft's website explains that Fusion performs a second, further correlation step to identify further correlated nodes.[48]



183. On information and belief, the Accused Fusion Products execute software instructions 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, Microsoft's website shows that in Fusion, further correlated nodes are identified through their connection to the other nodes.[49]

184. On information and belief, the Accused Fusion Products execute software instructions that generate a representation of a second graph comprising representations of one or more of both the first and second plurality of correlated nodes, and representations of

[48] *Id.*

[49] *Id.*

one or more of both the second and third plurality of edges, wherein certain 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, Microsoft's website shows that Fusion generates a report including a graph comprising nodes from the identified nodes, connected by edges, showing the possible event flow of a cybersecurity attack.[50]

185. Accordingly, as illustrated above, the Accused Fusion Products directly infringe one or more claims of the '628 Patent. Microsoft makes, uses, sells, offers for sale, and/or imports, in this District and/or elsewhere in the United States the Accused Fusion Products and thus directly infringes the '628 Patent.

186. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the '628 Patent, as provided in 35 U.S.C. § 271(b), including at least by inducing infringement by others, such as Microsoft's customers and end-users, in this District and elsewhere in the United States, to use the Accused Fusion Products in manners that infringe the '628 Patent. Microsoft induces such direct infringement through its affirmative acts of making, using, directing an entity to use, selling, offering to sell, and/or importing the Accused Fusion Products, as well as by advertising the Accused Fusion Products and providing instructions, documentation, and other information to its customers and end-users to encourage and teach them how to use the infringing Accused Fusion Products, including but not limited to by Microsoft providing in-store and online technical support, marketing materials, product manuals, advertisements, and other product documentation. At least as of service of this Complaint, Microsoft performs

[50] https://techcommunity.microsoft.com/blog/microsoftsentinelblog/behind-the-scenes-the-ml-approach-for-detecting-advanced-multistage-attacks-with/3239236

these affirmative acts with knowledge of the ’628 Patent and with the intent, or willful blindness, that the induced acts directly infringe the ’628 Patent.

187. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the ’628 Patent, as provided by 35 U.S.C. § 271(c), by contributing to direct infringement committed by others, such as Microsoft’s customers and end-users, in this District and elsewhere in the United States. Microsoft’s affirmative acts of selling and offering to sell the Accused Fusion Products in this District and elsewhere in the United States, and causing the Accused Fusion Products to be manufactured, used, sold, and offered for sale, contribute to Microsoft’s customers and end-users using the Accused Fusion Products, such that the ’628 Patent is directly infringed. The accused components in the Accused Fusion Products are material to the inventions claimed in the ’628 Patent, are not staple articles or commodities of commerce, have no substantial non-infringing uses, and are known by Microsoft to be especially made or adapted for use in the infringement of the ’628 Patent during at least the post-service period. Similarly, at least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the ’628 Patent and with the intent, or willful blindness, that they cause direct infringement of the ’628 Patent.

188. On information and belief, no licensed patent-practicing commercial products have been made, sold, offered for sale, or imported in a manner that would trigger the requirements of 35 U.S.C. § 287.

189. Microsoft’s infringement of the ’628 Patent has damaged and will continue to damage Qomplx.

190. Upon service of this complaint, Microsoft will have explicit written notice of its infringement of the ’628 Patent. Nevertheless, on information and belief, Microsoft will,

without authorization, knowingly, intentionally, purposefully, and deliberately continue to infringe the ’628 Patent.

COUNT VI

INFRINGEMENT OF U.S. PATENT NO. 11,539,663

191. Qomplx repeats and incorporates by reference each preceding paragraph as if fully set forth herein and further alleges:

192. The ’663 Patent, entitled “System and Method for Midserver Facilitation of Long-Haul Transport of Telemetry for Cloud-Based Services,” was duly and lawfully issued on December 27, 2022 and assigned to QOMPLX, Inc.. A true and correct copy of the ’663 Patent is attached hereto as Exhibit 6.

193. The ’663 Patent names Mika Chasman, Jeffrey Chung, Jason Crabtree, Luka Jurukovski, Richard Kelley, Artem Panasenkov, and Andrew Sellers as inventors.

194. The ’663 Patent claims priority to, among others, U.S. Application No. 15/141,752, filed April 28, 2016, and U.S. Provisional No. 62/568,291, filed October 4, 2017.

195. The ’663 Patent has been in full force and effect since its issuance. Qomplx owns all rights to the ’663 Patent that are necessary to bring this action.

196. The ’663 Patent, among other things, states that it “relates to the field of computer technology, more specifically to the field of computer architectures for enterprise data collection, analysis, and transmission to cloud-based services.” **Ex. 6** at 7:5-8.

197. As the ’663 Patent explains, there were “numerous problems” with “heterogeneous data transfer between cloud services and large offices or campuses for organizations,” including a “lack of reliable data collection methods, poor standardized support for connection-oriented protocols by network appliances, security concerns with unfiltered or poorly filtered data, and bandwidth concerns with constantly streaming data

which may result in network slowdown due to unprioritized data transfer." *Id.* at 7:23-31. For "larger business enterprises" with "thousands of computing devices sending data to a cloud-based service on separate connections," each connection represented "an additional security risk." *Id.* at 7:31-34.

198. At the time of the invention, "data collection and models [did] not scale well for adding new data sources and flexible adhoc queries." *Id.* at 7:34-36. This resulted in "too much data being passed, no context for data and data sources oftentimes, unqueryable data and data sources, [an] inability to flexibly and quickly add new data sources such as new devices or user accounts which generate new data for analysis, and log management and data storage becom[ing] expensive and disorganized." *Id.* at 7:36-42.

199. The '663 Patent overcomes the technological problems in the prior art by proposing "a system and method that uses midservers integrated with the business enterprise computer infrastructure and the cloud-based infrastructure to collect, aggregate, analyze, transform, and securely transmit data from a multitude of computing devices and peripherals at an external network to a cloud-based service." *Id*. at 7:61-66.

200. The inventions of the '663 Patent improve computer technology by "aggregating data at midservers," allowing "multiple connections [to] be presented over the network as a single secure connection to enterprise cloud-based systems" *Id*. at 9:49-51. The specific configurations claimed by the '663 Patent enable "[t]housands of connections" from a large enterprise to be "reduced to a single connection or a small number of connections." *Id*. at 9:52-54.

201. "Midserver architecture also solves the problem that not all devices support secure data transport." *Id*. at 9:60-61. "In order to support system log traffic the data must be wrapped in a secure protocol before leaving the network." *Id*. at 9:63-65. The architecture

taught by the '663 Patent "can provide this type of capability by collecting and wrapping the data before it leaves the network." *Id.* at 9:65-67.

202. For example, Claim 1 of the '663 Patent describes a system for ingesting data into a cloud service from an external network, wherein among other things a midserver automatically installs a virtual appliance software application configured to load stored configurations on the midserver, establishes a secure network connection to an external network, receives data over a local network, transforms at least a portion of the received data and retransmits the received data over the secure network connection to the external network as a single data stream. The specification explains that the claimed midserver architecture, by receiving data over a local network and transforming at least a portion of that received data at the midserver, "allows for large-scale, reliable ingestion (i.e., one or more of collection, aggregating, analysis (pre-processing), transformation (pre-processing), and secure transmission) of data into a cloud-based service from an external network," which in turn "improves data consistency, reliability efficiency of bandwidth usage, and security." *Id.* at 11:22-28. For example, the claimed systems have direct reductions on the peak bandwidth that will be consumed by a network using such a system, which directly reduces costs.

203. As another example, by establishing a secure network connection and transmitting data as a single data stream over that connection, the system described in Claim 1 of the '663 Patent uses "the midserver as a gateway to the cloud-based service," which "dramatically reduces the number of connections at the business enterprise's network edge, greatly reducing the number of avenues of attack and improving network security." *Id.* at 11:28-32. It also allows a network administrator to ensure that the number of *unencrypted* flows to the cloud-based service is zero. Many network attacks are based on finding a single unsecured endpoint inside an enterprise to exploit. The claimed systems prevent these attacks

and improve computer security by isolating the entire class of third-party event-generating endpoints behind the single, secure midserver.

204. The ’663 Patent’s improvements to computer architectures for enterprise data collection, analysis, and transmission to cloud-based services are further described in the specification.

205. Microsoft is not currently licensed to practice the ’663 Patent.

206. Qomplx is informed and believes, and thereon alleges, that Microsoft has infringed and continues to infringe one or more claims of the ’663 Patent in violation of 35 U.S.C. § 271, either literally and/or under the doctrine of equivalents, by making, using, directing an entity to use, selling, and/or offering for sale in the United States, and/or importing into the United States, without authorization, Microsoft products that practice one of more claims of the ’663 Patent, including without limitation products that incorporate, rely upon, interact with, or otherwise utilize Azure Monitor pipeline at edge (“Edge Pipeline” or the “Accused ’663 Functionality”).

207. For example and without limitation, products that incorporate, rely upon, interact with, or otherwise utilize Edge Pipeline, including at least Azure Monitor (collectively, the “Accused Edge Pipeline Products”) embody every limitation of at least claim 1 of the ’663 Patent, both literally and under the doctrine of equivalents, as set forth below.

208. Claim 1 of the ’663 Patent provides:

> A system for ingestion of data into a cloud-based service from an external network, comprising:
> a midserver comprising at least a processor, a memory, and a plurality of programming instructions stored in the memory and operating on the processor, wherein the plurality of programming instructions, when operating on the processor, cause the processor to:
> automatically install a virtual appliance software application, the virtual appliance software application configured to automatically load a plurality of stored configurations on the midserver;
> establish a secure network connection to an external network;

receive data over a local network from a plurality of computing devices;
apply a plurality of transformations to at least a portion of the received data; and
retransmit the received data over the secure connection as a single data stream.

209. By incorporating, relying upon, interacting with, or otherwise utilizing the Accused '663 Functionality, the Accused Edge Pipeline Products meet every element of this claim. The further descriptions below, which are based on analysis of publicly available information, are preliminary examples and are non-limiting.

210. For example, and to the extent that the preamble is limiting, Edge Pipeline is a system for ingestion of data into a cloud-based service from an external network.[51]

211. On information and belief, the Accused Edge Pipeline Products comprise midservers comprising at least a processor, a memory, and a plurality of programming instructions stored in the memory and operating on the processor. For example, Microsoft's website explains that Edge Pipeline runs on Kubernetes clusters, including midservers on customer networks. These servers contain processors, memories, and programming instructions.[52]

212. On information and belief, the Accused Edge Pipeline Products contain programming instructions that operate on the processor to automatically install a virtual appliance software application. For example, when the Azure portal is used, "all required components are created," including on the Azure Arc cluster.[53] Microsoft documentation

[51] https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/edge-pipeline-configure?tabs=Portal

[52] https://learn.microsoft.com/en-us/azure/azure-arc/kubernetes/overview

[53] https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/edge-pipeline-configure?tabs=Portal

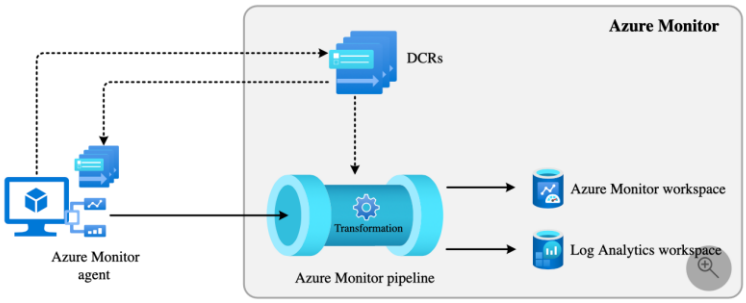
shows that services are created automatically.[54] These are collections of "pods", each containing one or more containers, which are virtual machines.[55]

213. On information and belief, the virtual appliance software application is configured to automatically load a plurality of stored configurations on the Midserver. For example, Azure Monitor distributes data collection rules to entities within the system, including to pipelines, to configure the flow of data.[56]

Data collection rule associations (DCRA)

Data collection rule associations (DCRAs) are used to associate a DCR with a monitored resource. This is a many-to-many relationship, where a single DCR can be associated with multiple resources, and a single resource can be associated with multiple DCRs. This allows you to develop a strategy for maintaining your monitoring across sets of resources with different requirements.

For example, the following diagram illustrates data collection for Azure Monitor agent (AMA) running on a virtual machine. When the agent is installed, it connects to Azure Monitor to retrieve any DCRs that are associated with it. In this scenario, the DCRs specify events and performance data to collect, which the agent uses to determine what data to collect from the machine and send to Azure Monitor. Once the data is delivered, the cloud pipeline runs any transformation specified in the DCR to filter and modify the data and then sends the data to the specified workspace and table.



214. Documentation shows that "edge pipeline configuration" is automatically configured.[57]

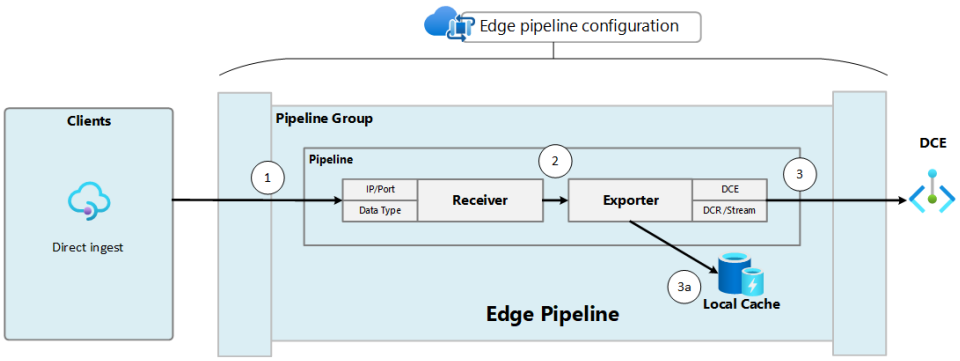
---

[54] *Id.*

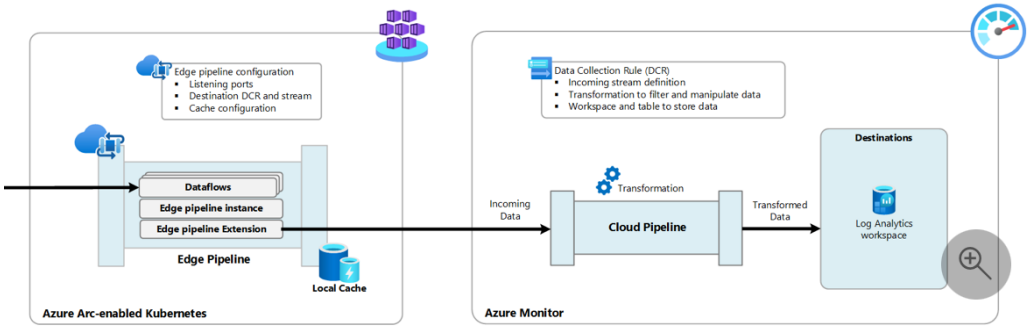
[55] https://kubernetes.io/docs/concepts/services-networking/service/

[56] https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/data-collection-rule-overview

[57] https://learn.microsoft.com/en-us/azure/azure-monitor/data-collection/edge-pipeline-configure?tabs=Portal



215. On information and belief, the Accused Edge Pipeline Products contain programming instructions that operate on the processor to establish a secure network connection to an external network and receive data over a local network from a plurality of computing devices. For example, Edge Pipeline sends data from the various dataflows it is collecting to the cloud.[58] Microsoft' documentations shows that all site-to-cloud communication is performed over a secure network connection.[59]



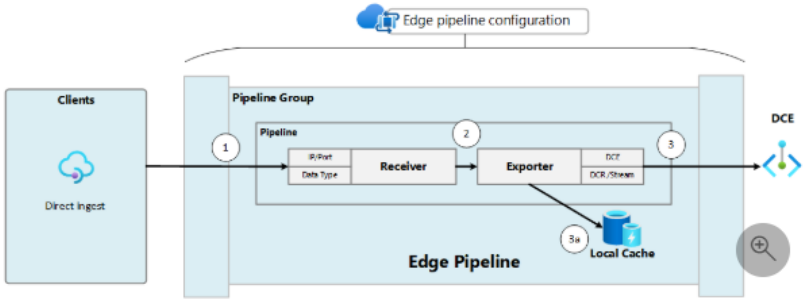
216. On information and belief, the Accused Edge Pipeline Products contain programming instructions that operate on the processor to apply a plurality of transformations to at least a portion of the received data. For example, the data collection rules define a

[58] https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/edge-pipeline-configure?tabs=Portal

[59] https://learn.microsoft.com/en-us/azure/security/fundamentals/double-encryption#data-in-transit

transformation for data in flight to be applied to at least a portion of the received data.[60] As a further example, Edge Pipeline can store data, performing a "storage transformation."[61]

3a.	Exporter stores data in the local cache if it can't connect to the DCE.	Persistent volume for the cache and configuration of the local cache is enabled in the pipeline configuration.



217. On information and belief, the Accused Edge Pipeline Products contain programming instructions that operate on the processor to retransmit the received data over the secure connection as a single data stream. For example, Microsoft's documentation explains that the cloud pipeline's exporter retransmits all data for a given data collection rule to a "Data Collection Endpoint" at a single URL.[62]

218. Accordingly, as illustrated above, the Accused Edge Pipeline Products directly infringe one or more claims of the '663 Patent. Microsoft makes, uses, sells, offers for sale, and/or imports, in this District and/or elsewhere in the United States the Accused Edge Pipeline Products and thus directly infringes the '663 Patent.

219. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the '663 Patent, as provided

[60] https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/edge-pipeline-configure?tabs=Portal

[61] *Id.*

[62] https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/edge-pipeline-configure?tabs=Portal

in 35 U.S.C. § 271(b), including at least by inducing infringement by others, such as Microsoft's customers and end-users, in this District and elsewhere in the United States, to use the Accused Edge Pipeline Products in manners that infringe the '663 Patent. Microsoft induces such direct infringement through its affirmative acts of making, using, directing an entity to use, selling, offering to sell, and/or importing the Accused Edge Pipeline Products, as well as by advertising the Accused Edge Pipeline Products and providing instructions, documentation, and other information to its customers and end-users to encourage and teach them how to use the infringing Accused Edge Pipeline Products, including but not limited to by Microsoft providing in-store and online technical support, marketing materials, product manuals, advertisements, and other product documentation. At least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the '663 Patent and with the intent, or willful blindness, that the induced acts directly infringe the '663 Patent.

220. Additionally, and/or in the alternative, on information and belief, Microsoft has also indirectly infringed and continues to indirectly infringe the '663 Patent, as provided by 35 U.S.C. § 271(c), by contributing to direct infringement committed by others, such as Microsoft's customers and end-users, in this District and elsewhere in the United States. Microsoft's affirmative acts of selling and offering to sell the Accused Edge Pipeline Products in this District and elsewhere in the United States, and causing the Accused Edge Pipeline Products to be manufactured, used, sold, and offered for sale, contribute to Microsoft's customers and end-users using the Accused Edge Pipeline Products, such that the '663 Patent is directly infringed. The accused components in the Accused Edge Pipeline Products are material to the inventions claimed in the '663 Patent, are not staple articles or commodities of commerce, have no substantial non-infringing uses, and are known by Microsoft to be especially made or adapted for use in the infringement of the '663 Patent during at least the

post-service period. Similarly, at least as of service of this Complaint, Microsoft performs these affirmative acts with knowledge of the '663 Patent and with the intent, or willful blindness, that they cause direct infringement of the '663 Patent.

221. Microsoft's infringement of the '663 Patent has damaged and will continue to damage Qomplx.

222. Upon service of this complaint, Microsoft will have explicit written notice of its infringement of the '663 Patent. Nevertheless, on information and belief, Microsoft will, without authorization, knowingly, intentionally, purposefully, and deliberately continue to infringe the '663 Patent.

PRAYER FOR RELIEF

WHEREFORE, Qomplx respectfully requests relief against Defendant Microsoft as follows:

A. A judgment declaring that Microsoft has infringed, either literally and/or under the doctrine of equivalents, one or more claims of each of the Asserted Patents;

B. An order and judgment enjoining Microsoft and its officers, directors, agents, employees, affiliates, attorneys, and all others acting in privity or in concert with it, and their subsidiaries, divisions, successors and assigns, from further acts of infringement of the Asserted Patents, or in the alternative, an ongoing royalty for all further acts of infringement;

C. A judgment awarding Qomplx all damages adequate to compensate for Microsoft's infringement of the Asserted Patents, including, but not limited to, lost profits, and in no event less than a reasonable royalty for Microsoft's acts of infringement, including all pre-judgment and post-judgment interest at the maximum rate permitted by law;

D. An order awarding Qomplx supplemental damages, including interest, with an accounting, as needed;

E. Costs of suit and reasonable attorneys' fees; and

F. Any other remedy to which Qomplx may be entitled, including under any other law, that this Court may deem just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Federal Rule of Civil Procedure 38(b), Qomplx hereby demands a jury trial on all issues so triable.

Respectfully submitted,

/s/ Paige Arnette Amstutz
Paige Arnette Amstutz
Texas Bar No. 00796136
pamstutz@scottdoug.com
Robert Pierce "Robby" Earle
Texas Bar No. 24124566
rearle@scottdoug.com
SCOTT DOUGLASS MCCONNICO LLP
303 Colorado Street, Suite 2400
Austin, TX 78701
Tel: (512) 495-6300
Fax: (512) 495-6399

Alan J. Heinrich (*pending pro hac vice)*
aheinrich@irell.com
Ian R. Washburn (*pending pro hac vice)*
iwashburn@irell.com
Amy E. Proctor (*pending pro hac vice)*
aproctor@irell.com
IRELL & MANELLA LLP
1800 Avenue of the Stars, Suite 900
Los Angeles, CA 90067
Tel: (310) 277-1010
Fax: (310) 203-7199