

(51) International Patent Classification:
H04L 12/24 (2006.01)

(21) International Application Number:
PCT/US2020/062347

(22) International Filing Date:
25 November 2020 (25.11.2020)

(25) Filing Language:
English

(26) Publication Language:
English

(30) Priority Data:
16/699,424 29 November 2019 (29.11.2019) US
16/699,431 29 November 2019 (29.11.2019) US
16/699,440 29 November 2019 (29.11.2019) US
16/699,446 29 November 2019 (29.11.2019) US

(71) Applicant: **AMAZON TECHNOLOGIES, INC.**
[US/US]; PO Box 81226, Seattle, Washington 98108-1226 (US).

(72) Inventors: **QIAN, Baihu**; c/o Amazon Technologies, Inc., PO Box 81226, Seattle, Washington 98108-1226 (US). **CHERUKU, Nikhil Reddy**; c/o Amazon Technologies, Inc., P.O. Box 81226, Seattle, Washington 98108-1226 (US). **DEB, Bashuman**; c/o Amazon Technologies, Inc., P.O. Box 81226, Seattle, Washington 98108-1226 (US). **HASHMI, Omer**; c/o Amazon Technologies, Inc., P.O. Box 81226, Seattle, Washington 98108-1226 (US). **MISHRA, Alok**; c/o Amazon Technologies, Inc., P.O. Box 81226, Seattle, Washington 98108-1226 (US). **PENNEY, Alexander Justin**; c/o Amazon Technologies, Inc., P.O. Box 81226, Seattle, Washington 98108-1226 (US). **SPENDLEY, Thomas Nguyen**;

(54) Title: CONFIGURATION AND MANAGEMENT OF SCALABLE GLOBAL PRIVATE NETWORKS




FIG. 1

(57) **Abstract:** This disclosure describes techniques for configuring and managing scalable global private networks associated with a service provider. Different input mechanisms, such as an API, a UI, or a CLI may be utilized to configure, and manage a global private network that spans across the cloud in different geographic locations and connects to different stand-alone networks. The user may proactively use the input mechanisms to configure and query different network resources to reactively configure settings for reacting to one or more events. The input mechanisms may also be utilized to define the network resources to be modeled within the global private network as well as connections within the global network. A user may configure events/metrics to be monitored, tasks/workflows to be performed, and the like. In some configurations, a network management service (NMS) may perform health monitoring and reachability monitoring to identify possible issues in the global network.



c/o Amazon Technologies, Inc., P.O. Box 81226, Seattle,
Washington 98108-1226 (US).

(74) **Agent: SULLIVAN, Timothy P.** et al.; Lee & Hayes, P.C.,
601 W. Riverside Avenue, Suite 1400, Spokane, Washing-
ton 99201 (US).

(81) **Designated States** (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN,
KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD,
ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO,
NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW,
SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) **Designated States** (*unless otherwise indicated, for every
kind of regional protection available*): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

CONFIGURATION AND MANAGEMENT OF SCALABLE GLOBAL PRIVATE NETWORKS

RELATED APPLICATIONS

5 [0001] This international application is related to the following applications: U.S. Patent Application No. 16/669,424, filed on November 29, 2019, U.S. Patent Application No. 16/699,431, filed on November 29, 2019, U.S. Patent Application No. 16/699,440, filed on November 29, 2019, and U.S. Patent Application No. 16/699,466, filed on November 29, 2019, which are fully incorporated by reference herein.

10 BACKGROUND

[0002] Companies of all sizes are challenged with efficiently and reliably networking their branch offices. To connect a remotely located branch office, a company may utilize a private network that includes leased line circuits that are backhauled to a corporate data center and/or the use of lower-cost broadband Internet. As a company that utilizes leased
15 line circuits expands and add more branches, the company continues to add physical, leased line wide area networks (WANs). Scaling out these traditional WANs can be costly as companies have to provision and manage expensive new leased-lines and hardware. While some companies may attempt to utilize lower-cost broadband Internet to connect remote branches, the use of lower-cost broadband Internet to expand can also be difficult and time-
20 consuming. For example, a company may spend a significant amount of time and money developing custom solutions to utilize broadband Internet. These custom solutions may include custom software and may use software, and devices from various networking vendors. The company may also have to obtain Internet services from a variety of different broadband providers.

25 BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The detailed description is set forth below with reference to the accompanying figures. In the figures, the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The use of the same reference numbers in different figures indicates similar or identical items. The systems depicted in the
30 accompanying figures are not to scale and components within the figures may be depicted not to scale with each other.

[0004] FIG. 1 is a software and network architecture diagram showing aspects of configuration and management of a scalable global private network.

[0005] FIG. 2 is a software and network architecture diagram showing aspects of a network management system (NMS) interacting with different networks of a global network.

5 [0006] FIG. 3 is a software and network architecture diagram showing aspects of communication between a network and an NMS.

[0007] FIG. 4 illustrates an example system diagram in which networks may be linked using redundant pathways.

[0008] FIG. 5A is a diagram showing an exemplary graphical user interface for configuring and managing a scalable global private network.

10 [0009] FIG. 5B is a diagram showing an exemplary graphical user interface for registering a gateway to include within a global network.

[0010] FIG. 5C is a diagram showing an exemplary graphical user interface for viewing information about a gateway within a global network.

15 [0011] FIG. 5D is a diagram showing an exemplary graphical user interface for viewing a graph representing a portion of the global network and performing a search of the graph.

[0012] FIG. 5E is a diagram showing an exemplary graphical user interface for viewing a graph representing a global network.

[0013] FIG. 6 is a flow diagram showing an illustrative routine for configuration and management of a scalable global private network, according to some examples.

20 [0014] FIG. 7 is a flow diagram showing an illustrative routine for creating a gateway within the global network, according to some examples.

[0015] FIG. 8 is a flow diagram showing an illustrative routine for monitoring network endpoints for connectivity, according to some examples.

25 [0016] FIG. 9 is a flow diagram showing an illustrative routine 900 for monitoring the global network to generate metrics, according to some examples.

[0017] FIG. 10 is a system and network diagram that shows an illustrative operating environment including several data centers that can be configured to implement aspects of the functionality described herein.

30 [0018] FIG. 11 is a computing system diagram illustrating a configuration for a data center that can be utilized to implement aspects of the technologies disclosed herein.

[0019] FIG. 12 is a computer architecture diagram showing an illustrative computer hardware architecture for implementing a computing device that can be utilized to implement aspects of the various technologies presented herein.

DETAILED DESCRIPTION

[0020] This disclosure relates to techniques for configuring and managing scalable global private networks. As used herein, a “scalable global private network”, which may be referred to herein is a “global private network” or “global network”, is a network that
5 utilizes one or more service/cloud provider networks (sometimes referred to simply as a “cloud”) to span different geographic locations and connect different stand-alone networks. A “stand-alone network” may include networks that are external to the service provider network (e.g., at client-owned premises or client-managed data centers) and/or networks that include computing resources allocated within a service provider network on behalf of
10 a user (e.g., a virtual private cloud (VPC)). Using techniques described herein, a user of a service provider network may utilize a network management service (NMS) to configure, visualize, and manage a global private network that spans across the cloud in different geographic locations and connects to different stand-alone networks associated with the user.

15 [0021] According to some configurations, the NMS may expose different input mechanisms to interact with a global private network, such as a graphical user interface (GUI), a user interface (UI), a command line interface (CLI), an application programming interface (API), and the like. In some examples, the NMS may generate a connected graph using industry-standard graph description language to represent the global private network
20 and then display a visual representation of the network graph within a GUI. The network graph may include nodes that represent different aspects of the global network, such as computing resources, networking resources, and other indicators that show connections between different the resources, as well as data that may indicate information about the global private network (e.g., metrics, events, versions of the graph, changes to the
25 graph, ...). According to some configurations, the user, or some device, service, or component may also query the network graph. For instance, a user may provide a query to the NMS to determine whether a particular change has been made to the network, locate one or more resources within the network, view one or more metrics, and the like. These graph-based queries may be performed by a user to analyze the network at scale such that
30 the user does not have to generate a query for each different region and/or network that is part of the global network.

[0022] According to some configurations, the NMS may associate metadata with the network graph. As an example, the NMS may annotate the network graph to indicate network capacity, metadata, state, and the like. The annotations may be associated with

nodes of the graph and/or edges of the graph that connect the nodes. The NMS may also generate subgraphs for portions of the global network, such that cloud, or stand-alone local networks can be imported/exported and merged into the global network. For instance, the user may request that a graph be generated of a stand-alone network that the user wants to add to the global network. After generating the network graph, the user may utilize the GUI to connect the stand-alone network to the global network.

[0023] According to some examples, the NMS exposes an API for the configuration and management of the global network. An API refers to an interface and/or communication protocol between a client and a server, such that if the client makes a request in a predefined format, the client should receive a response in a specific format or initiate a defined action. In the cloud provider network context, APIs provide a gateway for customers to access cloud infrastructure by allowing customers to obtain data from or cause actions within the cloud provider network, enabling the development of applications that interact with resources and services hosted in the cloud provider network. APIs can also enable different services of the cloud provider network to exchange data with one another. The user may utilize the API and/or some other input mechanism to configure and manage the global private network from one location. The user may proactively use the API to configure and query different network resources as well as use the API to reactively configure settings for reacting to one or more events. The events may indicate changes to the global network, such as but not limited to network additions, deletions, topology changes, and the like. The user may also use the API to configure alarms that may be triggered in response to a metric associated with the network changing (e.g., exceeding a value, dropping below a specified value, entering/leaving a particular range of values, ...).

[0024] Having the ability to configure and manage a global private network from a centralized location helps the user to build self-healing networks. For instance, the user may configure monitoring of different events and alarms and based on an occurrence of an alarm and/or an event perform an event-driven configuration (e.g., using AWS® Lambda) that automatically reacts to network changes quickly and provides the user with tools to automatically reconfigure their global private networks to mitigate impacts without human intervention. The API may also be utilized by the user to define the network resources to be modeled within the global private network, the connections between the network resources, the connections between a stand-alone network, the cloud network, and/or other networks.

[0025] In some examples, the user might also utilize the GUI, API, or CLI exposed by the NMS to configure metrics to be monitored, tasks/workflows to be performed (e.g., based on an occurrence of one or more events and/or alarms being triggered), and the like. In some configurations, the NMS may also perform operations on behalf of the user, such as health monitoring, reachability analysis (static and/or dynamic) and monitoring, and the like. As users run tasks across the cloud and stand-alone networks forming the global private network, the NMS may perform dynamic network reachability monitoring to assist the user in identifying possible problems with connectivity between different endpoints. For example, the NMS may utilize Two-Way Active Measurement Protocol (TWAMP), or some other dynamic analysis to identify connectivity between endpoints.

[0026] The Two-Way Active Measurement Protocol (TWAMP) is an open protocol that may be used to measure network performance between two network endpoints (e.g., devices in the network) that support the TWAMP framework. Generally, TWAMP is a framework that separates sessions based on the client/server architecture. The TWAMP client initiates a Transmission Control Protocol (TCP) connection and acts as a control-client and a session-sender, while the TWAMP server acknowledges the TCP connection and performs the roles of a server and a session-reflector. TWAMP-Control messages are exchanged between the endpoints and TWAMP-Test messages are exchanged between the session-sender and the session-reflector.

[0027] Managing reachability can be difficult in large, complex networks that span across both stand-alone networks and the cloud, as the monitoring and analysis of the reachability depends on the correct configuration of many resources and networking devices. Further, the global network may include firewalls, security groups, access control lists (ACLs)/network access control lists (NACLs) that are configured to block/allow traffic to flow between different network endpoints. In some cases, static network reachability analysis can be performed by the NMS before tasks are run across the global network. For instance, in some examples, the NMS may use formal methods, or some other verification technique, to determine whether the network is configured properly such that network endpoints are reachable. “Formal methods” refers to design techniques that use rigorously specified mathematical models to build software and hardware systems. Formal methods may utilize mathematical proof as a complement to dynamic testing in order to help ensure correct behavior. According to some configurations, the NMS is configured to perform static reachability analysis and/or dynamic reachability monitoring to determine network

reachability between endpoints and policy-based reachability to simplify the network monitoring and management.

[0028] According to some configurations, users may also define policies or intents on reachability, such as "VPC A can reach stand-alone network in CIDR range X" or "VPC D
5 cannot reach VPC E", and the NMS configures the network accordingly. This makes network management and configuration easier and less error-prone. In some examples, users may specify metrics and/or alarms for different endpoints within the network (e.g., monitor whether two points in a network, A and B, can or cannot reach each other).

[0029] While the techniques described herein are with reference to configuring and
10 managing global networks, the techniques are equally applicable to management of other networks. Additionally, while the examples herein discuss utilization of networks provided by a service provider, implementations are not so limited. For example, the techniques provided herein may be operating in other networks, across networks provide by different service providers, and the like.

15 [0030] Certain implementations and examples of the disclosure will now be described more fully below with reference to the accompanying figures, in which various aspects are shown. However, the various aspects may be implemented in many different forms and should not be construed as limited to the implementations set forth herein. The drawings herein are not drawn to scale. Like numerals represent like elements throughout the several
20 figures (which might be referred to herein as a "FIG." or "FIGS.").

[0031] FIG. 1 is a software and network architecture diagram 100 showing aspects of configuration and management of a scalable global network. It is to be appreciated that the environment 100 is merely illustrative and that the various configurations disclosed herein can be utilized in many different types of computing environments.

25 [0032] To provide functionality disclosed herein, the NMS 102 may include one or more computing resources 120. The computing resources 120 may be provided by a service provider that operates one or more service/cloud provider networks 104 (sometimes referred to simply as a "cloud"), which refers to a large pool of network-accessible computing resources (such as compute, storage, and networking resources, applications,
30 and services), which may be virtualized or bare-metal. The cloud can provide convenient, on-demand network access to a shared pool of configurable computing resources that can be programmatically provisioned and released in response to user commands. These resources can be dynamically provisioned and reconfigured to adjust to variable load. Cloud computing can thus be considered as both the applications delivered as services over

a publicly accessible network (e.g., the Internet, a cellular communication network, etc.) and the hardware and software in cloud provider data centers that provide those services.

[0033] A cloud provider network can be formed as a number of different regions, where a region is a separate geographical area in which the cloud provider clusters data centers.

5 Each region can include two or more availability zones connected to one another via a private high-speed network, for example a fiber communication connection. An availability zone (also known as an availability domain, or simply a “zone”) refers to an isolated failure domain including one or more data center facilities with separate power, separate networking, and separate cooling from those in another availability zone.

10 Preferably, availability zones within a region are positioned far enough away from one other that the same natural disaster (or other event) should not take more than one availability zone offline at the same time. Users can connect to availability zones of the cloud provider network via a publicly accessible network (e.g., the Internet, a cellular communication network) by way of a transit center (TC). TCs are the primary backbone

15 locations linking users to the cloud provider network and may be collocated at other network provider facilities (e.g., Internet service providers, telecommunications providers) and securely connected (e.g. via a VPN or direct connection) to the availability zones. Each region can operate two or more TCs for redundancy. The cloud provider network may deliver content from points of presence outside of, but networked with, these regions by

20 way of edge locations and regional edge cache servers.

[0034] The cloud provider network can provide on-demand, scalable computing platforms to users through a network, for example allowing users to have at their disposal scalable “virtual computing devices” via their use of the compute servers and block store servers. These virtual computing devices have attributes of a personal computing device

25 including hardware (various types of processors, local memory, random access memory (“RAM”), hard-disk and/or solid-state drive (“SSD”) storage), a choice of operating systems, networking capabilities, and pre-loaded application software. Each virtual computing device may also virtualize its console input and output (“I/O”) (e.g., keyboard, display, and mouse). This virtualization allows users to connect to their virtual computing

30 device using a computer application such as a browser, application programming interface, software development kit, or the like, in order to configure and use their virtual computing device just as they would a personal computing device. Unlike personal computing devices, which possess a fixed quantity of hardware resources available to the user, the hardware associated with the virtual computing devices can be scaled up or down depending upon

the resources the user requires. Users can choose to deploy their virtual computing systems to provide network-based services for their own use and/or for use by their users or clients. The computing resources 120 implemented by the NMS 102 and executed on behalf of one or more users of the service provider can be data processing resources, such as virtual machine (“VM”) instances, data storage resources, networking resources, data communication resources, network services, and other types of resources.

[0035] The computing resources 120 utilized can be general-purpose or can be available in a number of specific configurations. For example, data processing resources can be available as physical computers or VM instances in a number of different configurations.

The VM instances can be configured to execute applications, including web servers, servers, media servers, database servers, some or all of the network services described above, and/or other types of programs. Data storage resources can include file storage devices, block storage devices, and the like. The NMS 102 can also include and utilize other types of computing resources not mentioned specifically herein.

[0036] According to some configurations, servers are utilized to provide at least a portion of the computing resources 120 and execute software components to provide functionality described herein, including functionality related to the configuration and management of global private networks. The software components can execute on a single server or in parallel across multiple servers in the NMS 102. In addition, a software component can consist of subcomponents executing on different servers or other computing devices in the NMS 102. Various components can be implemented as software, hardware, or any combination of the two. In this regard, it is to be appreciated that the NMS 102 shown in FIG. 1 has been simplified for discussion purposes and that many additional software and hardware components can be utilized.

[0037] A user 138 of the NMS 102 can utilize a computing device 134, or some other input device, to access the NMS 102 through a network 132. The user 138 may be a user of the service provider network 104 that provides computing resources within the service provider network 104. The computing device 134 is an input/output device configured to receive input associated with specifying parameters utilized by the network management service 122 to configure and manage global private networks. The computing device 134 may also present for display a user interface 136, that may be utilized by the user 138 to view a graphical representation of a global private network, interact with the graphical representation to set or view parameters associated with the global private network, and

the like. The user interface 136 may also be utilized by the user as a CLI to interact with the global private network.

[0038] The computing device 134 may be one or more devices, such as but not limited to a smart phone, a smart watch, a personal computer (“PC”), desktop workstation, laptop
5 computer, tablet computer, notebook computer, personal digital assistants (“PDA”), electronic-book reader, game console, set-top box, consumer electronics device, server computer, or any other type of computing device capable of connecting to the network 132 and communicating with the NMS 102.

[0039] As illustrated, the computing device 134 may couple with the NMS 102 over a
10 network 132. The network 132 may represent an array or wired networks, wireless networks (e.g., WiFi), or combinations thereof. The NMS 102 may provide a variety of different services (not shown) as a network-accessible platform that is implemented as a computing infrastructure of processors, storage, software, data access, and so forth that is maintained and accessible via the network 132, such as the Internet. These services may
15 not require end-user knowledge of the physical location and configuration of the system that delivers the services. Common expressions associated with these remote services, include “on-demand computing”, “software as a service (SaaS)”, “platform computing”, “network accessible platform”, and so forth.

[0040] The network 132 can be a local-area network (“LAN”), a wide-area network
20 (“WAN”), the Internet, or any other networking topology known in the art that connects the user devices to the NMS 102. The user 138 can use an application (not shown) executing on computing device 134 to access and utilize the functionality provided by NMS 102. In some examples, the application is a web browser application, such as the Amazon® Silk® web browser, or some other web browser. Generally, a web browser
25 application exchanges data with the computing devices in the NMS 102 using the hypertext transfer protocol (“HTTP”) over the network 132.

[0041] The application might also be a stand-alone client application configured for communicating with the NMS 102. The client application can also utilize any number of communication methods known in the art to communicate with the NMS 102 across the
30 network 132, including remote procedure calls, SOAP-based web services, remote file access, proprietary client-server architectures, and the like. According to some configurations, the application provides a user interface 136 that can be utilized by the user 138 for the configuration and management of one or more global private networks. The

user interface 136 may also be utilized to present data, and/or to interact with the NMS 102.

[0042] In some examples, web service users or, in general, clients may utilize or otherwise control a processing entity of the service provider to control, access, or otherwise manage other computing resources. As such, data associated with the processing entity and/or the computing resources of the service provider may be transmitted to or received from computing resources of a client's private network (or other local network) via one or more network connections. As used herein, a processing entity may be a computing resource of the service provider and may include one or more computing devices, such as instantiated virtual machine instances, configured to access data of the distributed computing system (e.g., provided by the distributed system and acting on behalf of a client or user of the system).

[0043] In some configurations, the service provider may also provide storage, access, and/or placement of one or more computing resources through a service such as, but not limited to, a web service, a cloud computing service, or other network-based data management service. For example, a user or processing entity acting on behalf of the user may access, via the service provider, data storage services and/or data management services such that access mechanisms may be implemented and/or provided by the service provider to the processing entity utilizing the computing resources. In some examples, computing resource services, such as those provided by the service provider, may include one or more computing resources accessible across one or more networks through user interfaces (UIs), application programming interfaces (APIs), and/or other interfaces where the one or more computing resources may be scalable and/or expandable as desired

[0044] As briefly discussed above, a user of a service provider network, and/or a component or device, may utilize the NMS 102 to configure, visualize, and manage a global private network that spans across the cloud in different geographic locations and connects to different stand-alone networks associated with the user. According to some configurations, the NMS 102 exposes different input mechanisms to interact with a global private network, such as a user interface (UI) 136, a command line interface (CLI), API(s) 140, and the like. In some examples, the NMS 102 may generate a connected graph using industry-standard graph description language to represent the global private network. The NMS 102, network management service 122, or some other component may display a visual representation of the network graph within a GUI, such as within UI 136 (e.g., See

FIGs. 5C, 5D, and 5E). The NMS 102 may store the graph as data 128 in the data store 126, or at some other location (e.g., in a data store of an external network 116).

[0045] The network graph may include nodes that represent different resources and/or connections between resources/networks, indicators that show connections between the nodes, as well as data that may indicate information about the global private network (e.g., metrics, events, versions of the global network, changes to the global network, ...). According to some configurations, the user 138 and/or service, device, or component may also query the network graph. For instance, a user or service (e.g., network management service 122 and/or other service(s) 124) may provide a query via a CLI, an API, using a search box within a GUI, to the NMS 102 to determine whether a particular change has been made to the network, locate one or more resources within the network, view one or more metrics, and the like. These graph-based queries may be performed by a user 138 to analyze the network at scale such that the user does not have to generate a query for each different geographic area of the global network and/or network that is part of the global network.

[0046] According to some configurations, the user 138 and/or the NMS 102 may associate metadata with the network graph. As an example, the NMS 102 may annotate the network graph to indicate network capacity, metadata, state, relationships between the connected nodes, and the like. The NMS 102 may also generate subgraphs for portions of the global network, such that cloud, or stand-alone local networks can be imported/exported and merged into the global network. For instance, the user 138 may request that a graph be generated of a stand-alone network (e.g., external network 116A), that the user wants to add to the global network. After generating the network graph, the user 138 may utilize the GUI, CLI, and/or the API to connect the stand-alone network to the global network.

[0047] As illustrated in FIG. 1, the user 138 has created a global network that spans three regions 118 including VPCs 112 and connects external networks 116. As briefly discussed above, the user 138 may utilize a UI 136, API(s) 140, or some other mechanism to configure and manage a global network. As a particular example, assume that user 138 has utilized UI 136 to configure and manage the global network as illustrated in FIG. 1. Using the UI 136, the user 138 may select network resources to include within the global network. For instance, the user 138 may utilize a GUI to define that network 118A associated with a first geographic location, which may be referred to herein as a “network”, includes VPC

112A, VPC 112B, and VPC 112C connected using a gateway 106A that are connected to resources 114A in external network 116A using external connector 108A.

[0048] As used herein, a “gateway” is a gateway service that enables users to connect Virtual Private Clouds (VPCs) 112 and stand-alone networks, such as external networks 116 across different geographic locations/areas and the cloud. As a company grows and the global network grows, a gateway makes it easier for managing point-to-point connectivity across many different VPCs 112 and external networks 116, with the ability to centrally manage connectivity policies. As illustrated, the user 138 may also attach a VPN 110A to a gateway 106A that is connected to the VPCs 112A – 112C, without requiring the user 138 to attach a VPN 110 to each VPC.

[0049] Generally, a gateway 106 acts as a hub that controls how traffic is routed among the connected networks. According to some configurations, the gateway 106 utilizes a hub and spoke model, where the spokes connect the hub to the different VPCs 112 and/or VPNs 110. Any new VPC 112 added is simply connected to the gateway 106A and is then automatically available to every other network that is connected to the gateway 106A. In various examples, the routing/forwarding of network packets from one attached stand-alone network to another stand-alone network (e.g., VPC 112A to VPC 112B) is managed by one or more gateways 106 based on metadata and/or policies provided by the users. According to some examples, the gateways may be created using AWS® Transit Gateway.

[0050] In some configurations, the NMS 102 may expose functionality for incorporating an existing network into a global network. For instance, the user 138 may access a graph that represents network 118C and select an option within the UI 136 and/or programmatically through an API 140 that adds network 118C to the global network. According to some examples, the gateway 106D that connects the different networks 118A, 118B, and 118C may be added in response to a command received from a user and/or automatically in response to connecting a network being added to the global network. In the current example, the network 118C includes VPC 112H, VPC 112I, and VPC 112J connected using a gateway 106C that are connected to resources 114C in external network 116C using link 142 and external connector 108C. Gateway 106D connects network 118A to network 118C.

[0051] As another example, the user 138 may utilize a CLI to generate the representation of network 118B. In the current example, the network 118B includes VPC 112D, VPC 112E, VPC 112F, and VPC 112G connected using a gateway 106B that are connected to resources 114B in external network 116B using VPN 110B and external connector 108B.

Gateway 106D connects network 118B to network 118C, and network 118A. See FIGs. 5A – 5E and related description for example GUIs for configuring and managing a global network.

[0052] Networks 118 may have different resources 114 and connections. For example, network 118A may comprise a set of resources 114 at a data center or premise external to the service provider network's own data centers, which may be linked to the service provider network 104 using VPN 110 (virtual private network) tunnels or connections that utilize portions of the public Internet. Network 118C may also comprise resources 114C at premises outside the service provider network 104, connected to the service provider network 104 via dedicated physical links (which may be referred to as "direct connect" links), such as link 142, in the depicted example. The networks 118 may also include one or more virtual networks, such as VPCs 112, set up using resources located at the provider network's data centers. A virtual network may comprise a collection of networked resources (including, for example, virtual machines) allocated to a given client of the service provider network 104, which are logically isolated from (and by default, inaccessible from) resources allocated for other clients in other virtual networks. The client on whose behalf a virtual network is established may be granted substantial flexibility regarding network configuration for the resources of the virtual network (e.g., private IP addresses for virtual machines may be selected by the client without having to consider the possibility that other resources within other virtual networks may have been assigned the same IP addresses, subnets of the client's choice may be established within the virtual network, security rules may be set up by the client for incoming and outgoing traffic with respect to the virtual network, and so on). Similar flexibility may also apply to configuration settings at VPN-connected external networks such as external network 116A and external network 116B, and/or at external networks 140C connected via dedicated links, such as link 142, to the service provider network 104.

[0053] As briefly discussed above, the user 138, or some other device/component/service, may also utilize the UI 136, the API(s) 140, or some other input mechanism (e.g., speech) for the configuration and management of a global network. For instance, the user 138 may configure monitoring of different events and alarms and based on an occurrence of an alarm and/or an event perform an event-driven configuration (e.g., using AWS® Lambda) that automatically reacts to network changes quickly and provides the user with tools to automatically reconfigure their global private networks to mitigate impacts without human intervention.

[0054] In some examples, the user might also utilize the UI 136, API 140, and/or CLI exposed by the NMS 102 to configure metrics/events to be monitored, tasks/workflows to be performed (e.g., based on an occurrence of an event and/or a triggering of an alarm), and the like. In some configurations, the NMS 102 may also perform operations on behalf
5 of the user, such as health monitoring, reachability monitoring, and the like (See FIG. 2 and related discussion for more details). Additional details regarding the various components and processes described briefly above for configuring and managing global private networks will be provided below with regard to FIGS. 2-12.

[0055] FIG. 2 is a software and network architecture diagram showing aspects of a
10 network management system (NMS) 102 interacting with different networks 118 of a global network. FIG. 2 is similar to FIG. 1 in that it shows networks 118 that span different geographic locations but includes further details regarding performing reachability analysis and monitoring of resources within the networks 118.

[0056] In the example illustrated by FIG. 2, the global network includes four networks
15 118D – 118G. For example, network 118D may include one or more data centers in a country C1 (or territory, area, etc.), network 118E may include one or more premises in country C2, network 118F may include locations in state S1 of country C3, and network 118G may include resources in states S2 and S3 of country C3. In response to data received via the UI 136, the CLI, and/or the API(s) 140, the global network may be generated and
20 set up within each of the specified geographic areas/locations/regions. Furthermore, in the depicted example, the user may also have configured one or more external networks 116, such as 116D – 116G, at premises outside the service provider network's data centers to be part of the global network. As discussed above, such external networks 116 may be connected to the service provider network 104, for example, using VPNs 110 or dedicated
25 physical links 142 as discussed earlier.

[0057] In order to manage the routing of network traffic in a scalable manner between resources in various of the internal and external isolated networks, one or more gateways, such as gateway 106E, can be utilized. The gateway 106E, and gateways 106 included in the networks 118, provide network pathways or links that may be used to enable packets
30 to flow at desired levels of performance and availability. In many cases, there is more than one path between network endpoints. The paths of the global network within the service provider network 104 may be provisioned and managed (e.g., by adding/acquiring new fiber optic or other types of physical links, upgrading or otherwise changing the links and/or devices used) by the service provider network operator without obtaining input from

clients. According to some examples, the NMS 102 configures the pathways that connect the gateways 106 on behalf of the user 136. For example, the client may submit a programmatic request using the API(s) 140 indicating a set of stand-alone networks that are to be connected to one another using gateways 106, and the network management service 122 may configure the region-level gateways as well as one or more cross-region gateways, such as gateway 106E. In some examples, the user 138 may not even have to request the establishment of individual ones of the gateways. The user experience may thereby be made much smoother or easier than if the user 138 had to issue multiple programmatic requests.

10 [0058] As illustrated, FIG. 2 shows network 118D including external network 116D, network 118E including external network 116E, network 118F including external network 116E, and network 118G including external network 116G, coupled via gateway 106E. The NMS 102 includes network management service 122, monitoring service 204, reachability service 106, and other services 124. Each network 118 includes a network
15 manager 208 configured to perform operations relating to the configuration and management of the global private network. In some configurations, the network managers 208 are configured to interact with the different services exposed by the NMS 102.

[0059] As briefly discussed above, the reachability service 206 of the NMS 102 may perform static and/or dynamic network reachability monitoring to assist the user 138 in
20 identifying possible problems with connectivity between different endpoints within a global network. Managing network reachability can be difficult in large, complex networks that span across both stand-alone networks and the cloud, as the monitoring and analysis of the reachability depends on the correct configuration of many resources and networking devices. Further, the global network may include firewalls, security groups, access control
25 lists (ACLs)/network access control lists (NACLs) that are configured to block/allow traffic to flow between different network endpoints. According to some configurations, the reachability service 206 of the NMS 102 is configured to manage the monitoring of the different networks, the connectivity between the networks, and identify problems in network reachability between different network endpoints.

30 [0060] In some examples, the NMS 102 may receive information from the network managers 208 of the different networks 118 and provide the data to the reachability service 206 to determine reachability between different endpoints. For example, the NMS 102 may receive events data from a network 118 that a network change has been made (e.g., a device added/removed from the global network, setting(s) have been changed, ...) and/or metrics

data that is associated with performance of the network (e.g., performance of a gateway and/or some other node within the global network). The NMS 102 may also transmit instructions to the network managers 118 to perform monitoring of one or more computing resources and/or network resources within the associated region.

5 **[0061]** According to some configurations, a user 138 may define what endpoints to monitor. For instance, the user 138 may specify one or more policies that define what network endpoints to monitor for reachability. As an example, the user 138 may specify policies that are positively stated, such as "VPC A can reach stand-alone network in CIDR range X" or negatively stated, such as "VPC D cannot reach VPC E". The user 138 may
10 also specify what events to utilize in determining when to perform static/dynamic reachability analysis. For example, the reachability service 206 may perform static analysis in response to a change being made to the global network and perform dynamic analysis at some specified period (e.g., every minute, five minutes, ...). In some configurations, the reachability service 206 may monitor each of the different network endpoints defined
15 within the global network for reachability. This makes network management and configuration easier and less error-prone since the user does not have to individually access each network and analyze a flow of data between the networks.

[0062] In some examples, users may specify to determine network connectivity for different endpoints within the network (e.g., monitor whether two points in a network, A
20 and B, can or cannot reach each other). One or more alarms may be configured to trigger in response to a value of metric being out of range, exceeding a specified value, below a specified value, and the like. According to some examples, the monitoring service 204, and the reachability service 206 may monitor various network resources/parameters.

[0063] In some configurations, the monitoring service 204 is configured to obtain metrics
25 data and/or event data from the networks 118. Generally, metrics data includes metrics that identify a performance of a computing resource and/or a network element. For example, the monitoring service 204 may instruct a network manager 208 to monitor, collect and store metrics data from various network resources, applications, and services operating in the network 118. In some configurations, the monitoring service 204 collects specified
30 and/or default metrics relating to use of resources. For example, each region may collect metrics data relating to CPU utilization, data transfer, disk usage, memory usage, bandwidth utilized, latency, and the like.

[0064] In some examples, the monitoring service 204 collects metrics associated with gateways 106. For example, the metrics may include the number of bytes received by the

gateway, The number of bytes sent from the gateway 106, the number of packets received by the gateway 106, the number of packets sent by the gateway 106, the number of packets dropped by a gateway 106, the number of packets dropped by a gateway 106 because they did not match a route, and the like. According some configurations, the metrics may also include the number of bytes sent to each connection of the gateway 106, the number of packets received by each connection of the gateway 106, the number of packets sent by the gateway 106 to each connection of the gateway, and the like. As such, the user 138 may monitor metrics associated with an overall performance of a gateway 106, as well as performance of a gateway 106 with each connection of the gateway 106.

10 [0065] In some configurations, the monitoring service 204 may collect metrics for VPNs 110, such as a state of the tunnel, a number of bytes received through a VPN tunnel, bytes sent through the VPN tunnel, and the like. The VPN metrics may be aggregated per VPN tunnel and per VPN connection. According to some examples, the monitoring service 204 may also collect metrics for VPNs 110 such as a number of down tunnels terminated on a device, a number of bytes received through a device or link, a number of bytes sent through a device or link, and the like. The metrics may be aggregated by device and/or by link.

[0066] According to some examples, different events may be published by one or more services, such as by network management service 122 and/or other services 124. For instance, in some configuration, the following events may be published a network topology change, a routing update, a network status change, a gateway attachment created, a gateway 106 attachment deleted, a gateway 106 added, a gateway 106 deleted, a route/path created in a gateway 106 route table, a route deleted in gateway 106 route table, a route replaced in gateway 106 route table, a VPN 110 connection created, a VPN 110 connection deleted, VPN 110 connection's gateway 106 changed, a VPN 110 tunnel's IPsec session went down, a VPN 110 tunnel's IPsec session is now up, a VPN 110 tunnel's session went down, a VPN 110 tunnel's session is now up, a VPN 110 tunnel's endpoint instance replaced, a route added for VPN 110 connection, route removed for VPN 110 connection, and the like.

[0067] The monitoring service 204 may obtain the data from the different networks 118D and utilize this data to monitor operational performance, troubleshoot issues, and spot trends within each of the different networks 118 forming the global network.

[0068] For instance, each network manager 208 may be configured to collect data for the gateways 106. For instance, the bandwidth usage between the VPCs 112 and a VPN 110 connection, packet flow count, packet drop count, and the like may be monitored. In some

examples, information on the IP traffic routed through a gateway may also be monitored. According to some configurations, the network management service 122, the monitoring service 204, and/or the reachability service 206 collects this data in form of logs and metrics. This allows a user to access the metrics data and reachability data from a single
5 location rather than having to access each network in order to obtain the metrics data and reachability data for the entire global network. Instead of monitoring individual systems and applications in silos (server, network, database, etc.), the different components of the global network may be monitored as a complete stack (e.g., applications, infrastructure, and services). This data may be then be used the network management service 122, the
10 monitoring service 204, the reachability service 206, and/or some other service 124 to trigger alarms, create logs, and generate events that may be used to perform automated tasks (e.g., take a corrective action, provide warnings to a user 138 via the UI 136, ...). Being able to access the data from a single location may reduce the Mean Time to Resolution (MTTR) to address a problem that is detected or identified within the network.

15 **[0069]** A number of different types of metrics may be utilized, including for example latency metrics associated with individual ones of the network pathways between networks, bandwidth metrics associated with individual ones of network pathways, packet loss metrics associated with individual ones of the network pathways, or flow count metrics associated with individual ones of network pathways. In some examples, a user 138 may
20 be provided indications of pathways that are available for inter-region traffic between stand-alone networks (e.g., the user may be informed that some paths pass through country C1, others pass through countries C2 and C3, and the like). In some configurations, the user 138 may be provided metrics for inter-region traffic (e.g., total number of packets transmitted between networks 118, latencies for packets sent between different endpoints,
25 and the like.

[0070] As discussed above, the NMS may be configured to perform static reachability analysis and/or dynamic reachability monitoring to determine network reachability between endpoints and policy-based reachability to simplify the network monitoring and management. In some examples, the reachability service 206 may utilize formal methods,
30 or some other verification technique, to determine whether the network is configured properly such that network endpoints are reachable. For instance, the user 138, or some other user associated with the service provider network 104, or some other authorized user may generate mathematical models that model a global network. The reachability service 206 may apply these formal methods in response to changes being made to the global

network and/or at a request of the user 138 and/or by some other service, device, or component. In this way, if the user requests a change to the global network that would result in a loss of network connectivity, the NMS 102 may provide this information and/or recommendations to correct this detected loss of connectivity.

5 [0071] The monitoring of the global network allows a user 138 of the service provider network 104 to gain actionable insights that help the user 138 optimize application performance, manage resource utilization, and understand system-wide operational health of the global network. In some configurations, the NMS 102 may utilize monitoring service 204 for collecting, aggregating, and summarizing compute utilization information like
10 TWAMP data, CPU, memory, disk, and network data, as well as diagnostic information network reachability between various endpoints, to help the user 138 isolate network issues and resolve the issues quickly. In some examples, the network management service 102 may perform a self-correcting action (e.g., by executing one or more workflows) in response to an alarm triggered by metrics data or monitoring data.

15 [0072] FIG. 3 is a software and network architecture diagram 300 showing aspects of communication between a network 118 and an NMS 102. It is to be appreciated that the environment 300 is merely illustrative and that the various configurations disclosed herein can be utilized in many different types of computing environments.

[0073] In the example illustrated in FIG. 3, the NMS 102 includes an API 302, a global
20 workflow service 304, a global data store 306, an event publisher 308 to publish events 312, and a metrics publisher 310 to publish metrics 314. The NMS 102 communicates with a network manager 208 of the network 118 via the communication channel 316. As illustrated, the network manager 208 includes a regional workflow service 318, a regional data store 320, an event service 322, and a metrics service 324. The NMS 102 and the
25 network manager 208 may utilize a different number of components according to other configurations.

[0074] According to some examples, the API 302 may be configured to expose functionality for interacting with the NMS 102 for configuring and managing global networks that span across different geographic locations/areas/regions. As discussed
30 above, the API 302 may be utilized by the user 138 to specify parameters associated with the configuration and management of the global network. The API 302 may also be configured to interact with the network management service 122, and other services 124 of the service provider network 104, such as a metrics service 324, and an event service 322.

[0075] According to some configurations, communication takes place between the NMS 102 in the service provider network 104 and each region. In these configurations, a network manager 208 in one location does not directly communicate with another network manager 208 that is located in a different location. In other configurations, network managers 208 may communicate with each other.

[0076] In some examples, the global data store 306 is configured to store data associated with the global networks associated with different users. The regional data store 320 is configured to store data associated with the network resources that are located within a particular geographic area. The data stores may include one or more databases for storing different types of data, such as a SQL database, a not only SQL (NoSQL) database, a graph database, and the like.

[0077] The global workflow service 304 and the regional workflow service 318 provide functionality associated with performing actions relating to workflows. The global workflow service 304 and the regional workflow service 318 provide functionality for performing different tasks and managing intertask dependencies, scheduling, and concurrency in accordance with the defined logical flow. For instance, the global workflow service 304 and/or the regional workflow service 318 may be implemented using AWS® Lambda and AWS® Step Functions, AWS® Simple Workflow service, and the like. Generally, the workflow services may execute code and access different computing resources, such as computing resources 120 in the service provider network 104 and/or resources 114 included in one or more stand-alone networks 116, or other networks.

[0078] The event service 322 is configured to receive and identify different events. For example, the event service 322 may be configured to identify changes in a network, such as changes in a gateway, changes in VPN, and the like. As discussed above, the events may include events such as but not limited to network topology changed, routing updates, network status changed, gateway 106 updates (e.g., attachment created/deleted, gateway 106 added/ deleted, route/path altered in a gateway 106 route table, ...), VPN 110 updates (e.g., VPN 110 connection created/deleted/changed, VPN 110 tunnel changes, ...), and the like. According to some configurations, the event service 322 identifies the changes based on data generated by the metrics service 324, the reachability service 206, and/or some other service, device or component. In some examples, the event service 322 is a service that runs code without provisioning or managing servers (e.g., AWS® Lambda).

[0079] The metrics service 324 is configured to generate metrics that may be utilized by the NMS 102 and/or the user 138. In some configurations, the metrics service 324 may

utilize one or more services provided by the service provider network 104. The metrics service 324 collects metrics from network resources, and applications to monitor operational performance, troubleshoot issues, and spot trends within the global network.

[0080] FIG. 4 illustrates an example system environment 400 in which networks 118 may be linked using redundant pathways. As illustrated in FIG. 4, networks 118 are redundantly connected to more than one instance of NMS 122. For example, each of the networks 118A, 118B, and 118C may be connected to both NMS 122A and NMS 122B within the global network. NMS 122A and NMS 122B may be located in different geographic regions and/or in the same geographic region. Further, the networks 118 may be connected to more than two instances of NMS 122. According to some examples, data between the NMS 122A and NMS 122B are replicated. In this way, in case an NMS 102 goes down (or is experiencing issues that affect performance), the redundant NMS may provide services to the networks 118.

[0081] FIGs. 5A – 5E are diagrams showing exemplary graphical user interfaces for configuring and management of scalable global private networks. In some examples, the user configures and manages a global private network using a graphical user interface (GUI). In other examples, the user may use a command line interface, or utilize an Application Programming Interface (API). FIGs. 5A, 5B, 5C, 5D, and 5E illustrates example graphical user interfaces 500, 520, 540, 570, and 592 where the user 138 can configure and manage scalable global private networks. In some instances, the NMS 102 may provide data for displaying a GUI to a display associated with the user computing device 138.

[0082] In the example illustrated in FIG. 5A, GUI 500 shows user interface (UI) elements for selecting a global network. More or fewer UI elements may be included within GUI 500. As illustrated, the GUI 500 includes a search global network UI element 504 to enter a search term to locate a defined global network. An edit global network UI element 506 may also be provided that, when selected, is used to edit a global network. For instance, the user 138 may select “Global Network 1” as indicated by indicator 514.

[0083] GUI 500 also includes a view global network UI element 508 to view a global network. For example, selecting the view global network UI element 508 may cause a graph of the global network to be displayed (See FIG. 5D and 5E for example graphs of a global network). Create global network UI element 512 may be used to delete a previously created global network. Delete global network UI element 510 may be used to delete a

previously created global network. For example, selecting the delete global network UI element 510 may cause the selected global network to be deleted.

[0084] FIG. 5B shows GUI 520 that includes user interface (UI) elements for registering a gateway 106 to include within a global network. More or fewer UI elements may be included within GUI 520. As illustrated, the GUI 520 includes a gateways UI element 522, a devices UI element 524, a regions UI element 526, a connections UI element 528, and a configuration UI element 530 that may be selected to create, configure and/or edit a gateway 106, devices, regions 118, connections (e.g., VPNs 110, links 142), and/or other configurations relating to a global network.

[0085] GUI 520 also includes a search gateway UI element 532 for a user 138 to enter a search term to locate a gateway 106. In the current example, the user 138 has identified three gateways to include within the global network as indicated by indicator 534. GUI 530 also includes a cancel network UI element 536 to cancel registration of a gateway 106 and a register gateway UI element 538 to include one or more gateways 106 as part of the global network.

[0086] FIG. 5C shows GUI 540 that includes user interface (UI) elements for viewing information about a gateway 106 within a global network. More or fewer UI elements may be included within GUI 540. As illustrated, the GUI 540 includes a home UI element 542 to go to a home display, an events UI element 544 to configure alarms, actions, and tasks to perform, and a policies UI element 546 to configure policies for the global network.

[0087] The GUI 540 also includes a view of a map 548 that displays a world view that includes gateway UI elements 554A – 554C that show different gateways 106 that are part of the global map. UI element 556 provides an indication to the user 138 to select one of the gateway UI elements 554 to obtain details about the gateway 106. In the current example, the user has selected gateway UI element 554A. In response to selection of the gateway UI element 554A, graphical window 552 is displayed. Graphical window 552 includes a graph view that shows the connections of gateway 106F to VPCs 112K – VPC 112M, link 560, connection 558, and gateway 106A. More or less information may be shown within graphical window 552.

[0088] GUI 540 also shows a create gateway UI element 562, an import network UI element 564, and a remove gateway UI element 566. Indicator 568 shows additional details relating to gateway 106A and gateway 106F.

[0089] FIG. 5D shows GUI 570 that includes user interface (UI) elements for viewing a graph representing a portion of the global network and performing a search of the graph.

More or fewer UI elements may be included within GUI 570. As illustrated, the GUI 570 includes a gateways UI element 572 selectable to access information about gateways 106 of the global network, a devices UI element 574 selectable to access information about devices of the global network, a regions UI element 576 selectable to access information about networks of the global network, a connections UI element 578 to selectable to access information about connections of the global network, and a configuration UI element 580 selectable to access configuration information about the global network.

[0090] The GUI 570 also includes a graph display area 586 that display a connected graph representing a portion of the global network. In the current example, the graph display area 586 shows a portion of a global network connected to gateway 106G. As can be seen, the gateway 106G is connected to a VPC 114N and a VPN 110L. The VPC 114N is connected to VPCs 114O – 1140R, which are connected to resources 114M – 114S. VPN 110L is connected to VPC 110S, link 142L, and direct connection 108A. The user 138 may identify what to display in the graph display area 586 using selection UI elements 584. For instance, the user 138 has selected to show devices, links, sites. In other examples, the selection UI elements 584 may include other options to show more or fewer details. Similarly, the user 138 may select elements to collapse within the graph. For instance, if the user selected to collapse the VPN connections, then the graph display area may not show the connections to VPN 110L.

[0091] GUI 570 also includes a graph search UI element 588 for a user 138 to enter a search term to locate a portion or resource within the graph. In the current example, the user 138 has entered the search term “EAST” in the graph search UI element 590. In response to the search, the network management service 122 has returned two results “GATEWAY 2” and “GATEWAY 3”, and the user has selected “GATEWAY 2” to be illustrated within graph display area 586 as indicated by indicator 594.

[0092] FIG. 5E shows GUI 592 that includes user interface (UI) elements for viewing a graph representing a global network. More or fewer UI elements may be included within GUI 592. As illustrated, the GUI 592 includes UI elements selectable to access information about the global network. The GUI 5920 displays a connected graph representing the global network as illustrated in FIG. 1.

[0093] As can be seen, the gateway 106G is connected to network 118A, network 118B, and network 118C. Network 118A includes gateway 106A connected to gateway 106G, VPCs 112A – 112C, and VPN 110A. VPN 110A is connected to connector 108A that is connected to resources 114R1 that includes resources (114A – 114C) of an external

network. Network 118B includes gateway 106B connected to gateway 106G, VPCs 112E – 112G, and VPN 110B. VPN 110B is connected to connector 108B that is connected to resources 114R3 that includes resources (114F – 114H) of an external network. Network 118C includes gateway 106C connected to gateway 106G, VPCs 112H – 112J, and link
5 142 coupled to connector 108C that is connected to resources 114R2 that includes resources (114D and 114E) of an external network.

[0094] As illustrated with regard to FIG. 5D, the user 138 may identify what to display in the graph display area using selection UI elements 584. In the current example of FIG. 5E, the user 138 has selected to show the entire global network. In other examples, the
10 selection UI elements 584 may include other options to show more or fewer details. Similarly, the user 138 may select elements to collapse within the graph.

[0095] FIGs. 6-9 are flow diagrams showing illustrative routines 600, 700, 800, and 900 for configuration and management of global scalable networks, according to examples disclosed herein. It should be appreciated that the logical operations described herein with
15 respect to FIG. 6, FIG. 6, FIG. 8, FIG. 9, and the other FIGS., can be implemented (1) as a sequence of computer implemented acts or program modules running on a computing system and/or (2) as interconnected machine logic circuits or circuit modules within the computing system.

[0096] The implementation of the various components described herein is a matter of
20 choice dependent on the performance and other requirements of the computing system. Accordingly, the logical operations described herein are referred to variously as operations, structural devices, acts, or modules. These operations, structural devices, acts, and modules can be implemented in software, in firmware, in special purpose digital logic, and any combination thereof. It should also be appreciated that more or fewer operations can be
25 performed than shown in the FIGS. and described herein. These operations can also be performed in parallel, or in a different order than those described herein. Some or all of these operations can also be performed by components other than those specifically identified.

[0097] FIG. 6 is a flow diagram showing an illustrative routine 600 for configuration and
30 management of a scalable global private network, according to some examples. The routine 600 may be performed by computing resources 120 associated with the NMS 102 and/or other computing resources, such as computing resources associated with some other network or system.

[0098] At 610, one or more interfaces are provided for configuration and management of a scalable global network. As discussed above, the NMS 102 may expose one or more API(s) 140, a CLI, and/or a UI 136, such as a graphical user interface. A user, such as user 138, may utilize the interfaces to configure, view, and manage global networks.

5 [0099] At 620, a request is received to perform an operation for configuration/management of a scalable global private network. For example, the request may be to add a stand-alone network to the current global network, change a policy (e.g., security policy) associated with the network, add a resource to one or more of the networks forming the global private network, configure one or more events and/or actions, configure
10 monitoring, or perform some other action or operation. As discussed above, the request may be received from a user, such as user 138 of a service provider network 104. In some examples, the network management service 122 within NMS 102 receives the request.

[0100] At 630, the action to perform is identified. As discussed above, NMS 102 may receive the request and identify the service, component, or device to perform the requested
15 operation. In some examples, the NMS102 may cause one or more workflows to be performed in response to the request. For instance, the NMS may identify that the request alters the global network, such as an update a security policy, add a resource to a particular region, monitor one or metrics, determine reachability between specified nodes within the global network, and the like.

20 [0101] At 640, static analysis of the global network may be performed. As discussed above, the network reachability service 206 may perform a static analysis using formal methods, or some other verification technique, in response to a request from a user 138 and/or in response to some other event (e.g., a change of network topology). In some configurations, the static analysis may be performed before the network change is
25 implemented. In this way, if a network connectivity issues is detected by the analysis, the user 138 may be provided with data indicating the issue and/or recommendations to correct the issue with network connectivity.

[0102] At 650, the request to perform the action is transmitted to one or more of the locations of the global private network. For example, if the request is to change a security
30 policy, the network management service 122 may provide the security policy to the network managers 208 to implement the change within each of the geographic locations and provide the change to the external networks 116 when determined.

[0103] At 660, the metrics/events to monitor within the locations of the networks forming the global network are configured. As discussed above, the network management service

122, the event service 322, the metrics service 324, or the reachability service 206 may transmit instructions to the network managers 208, or some other service, device or component, to monitor one or more network resources within the different locations.

[0104] At 670, the scalable global private network is monitored. As discussed above, each of the networks at the different locations may monitor events and metrics and provide metrics data and events data back to the NMS 102 for further analysis and/or actions. In some examples, the reachability service 206 utilizes monitoring information from the different locations to identify when there may be a connectivity issue between specified endpoints in the global network. See FIG. 8 and FIG. 9 for additional details.

[0105] At 680, network data is provided to a user and/or some other device or component. As discussed above, the NMS 102 may provide data to the user 138 via a UI 136, such as a GUI, and/or use the data to trigger one or more alarms and/or perform one or more workflows.

[0106] FIG. 7 is a flow diagram showing an illustrative routine 700 for creating a gateway 106 within the global network. The routine 700 may be performed by computing resources 120 associated with the NMS 102 and/or other computing resources associated with the global network. While routine 700 illustrates a request to create a gateway, a similar routine may be performed to perform other actions within the global network.

[0107] At 710, a request is received to create a gateway 106 for a particular network location. As discussed above, the user 138 may utilize an API 140, a CLI, or a UI 136 to request to create a gateway 106. In some configurations, the network management service 122 receives the request to create a gateway 106.

[0108] At 720, the request to create the gateway is transmitted to the location in which the gateway 106 is to be created. As discussed above, the network management service 122 may transmit the request to a network manager 208 to create the gateway in the network 118.

[0109] At 730, a response is received from the region indicating whether the gateway was created and/or whether any problems occurred during the creation of the gateway 106. In some configurations, the network manager 208 transmits a message to the network management service 122 indicating whether the creation of the gateway 106 was successful.

[0110] At 740, a decision is made as to whether the creation of the gateway 106 was successful. When the creation of the gateway 106 was successful, the routine moves to 760. When the creation of the gateway 106 was not successful, the routine moves to 750.

[0111] At 750, a notification of a problem creating the gateway 106 is provided. According to some examples, the notification may be provided to the user 138 via a UI 136. In other examples, the notification may be provided to a service, component and/or device.

5 [0112] At 760, the gateway may be connected to other regions when determined. For example, the network management service 122 may determine from the request that the gateway 106 is to be connected to one or more other nodes located within one or more other regions.

[0113] FIG. 8 is a flow diagram showing an illustrative routine 800 for monitoring
10 network endpoints for connectivity, according to some examples. The routine 800 may be performed by computing resources 120 associated with the NMS 102 and/or other computing resources associated with the global network.

[0114] At 810, the network endpoints to monitor for connectivity are identified. As discussed above, the reachability service 206 may identify network endpoints to check for
15 connectivity based on input specified by the user 138 as well as identify other network endpoints to monitor for connectivity, such as gateways 106, VPCs 112, VPNs 110, connectors 108, links 142, and the like.

[0115] At 820, static analysis of the global network may be performed. As discussed above, the network reachability service 206 may perform a static analysis using formal
20 methods, or some other verification technique, in response to a request from a user 138 and/or in response to some other event (e.g., a change of network topology).

[0116] At 830, the network is dynamically monitored. As discussed above, the reachability service 206, may receive monitoring data, such as events data indicating an occurrence of one or more events and metrics data, from the network managers 208 that
25 may be used to determine connectivity between network endpoints within a region and/or network endpoints that span more than one region. In some examples, the reachability service 206 may perform dynamic network connectivity checks between network endpoints specified by the user 138 as well as other network endpoints, at predetermined times (e.g., every minute, five minutes, ...).

30 [0117] At 840, a determination is made as to whether there is connectivity between network endpoints. For example, the reachability service 206 may identify that one or more network endpoints are not reachable and/or that specified network endpoints are reachable.

[0118] At 850, a decision is made as to whether the network endpoints have network connectivity. As discussed above, the reachability service 206 may decide whether the

network endpoints have network connectivity. When the network endpoints have network connectivity, the routine 800 returns to 820. When the network endpoints do not have network connectivity, routine 800 flows to 860.

5 [0119] At 650, an action is caused to be performed. As discussed above, the reachability service 206 may provide connectivity data to the user 138 via the UI 136 indicating the connectivity issue. The reachability service 206 might also provide the data to one or more other services, devices, or components that in turn execute a workflow to address the network connectivity issue. The routine 800 may end or return to 820.

10 [0120] FIG. 9 is a flow diagram showing an illustrative routine 900 for monitoring the global network to generate metrics, according to some examples. The routine 900 may be performed by computing resources 120 associated with the NMS 102 and/or other computing resources associated with the global network.

15 [0121] At 910, the network resources to monitor are identified. As discussed above, the monitoring service 204 may identify network resources to monitor based on input specified by the user 138 and/or identify network resources to monitor based on other specifications. For example, the reachability service 206 may instruct the monitoring service 204 to monitor specified metrics and/or events. In other examples, the monitoring service 204 may monitor computing resources for default metrics.

20 [0122] At 920, the networks are configured to monitor specified computing resources. As discussed above, the monitoring service 204, may instruct the network managers 208 to monitor specified computing resources for one or more metrics and/or events.

25 [0123] At 930, monitoring data is received from the different network locations. As discussed above, the monitoring service 204 may receive monitoring data including the metrics from each of the different networks 118. In this way, the user 138 may access the metrics for the global network without having to access each location separately.

30 [0124] At 940, an action is caused to be performed. As discussed above, the monitoring service 204 may provide events data, and/or metric data to the user 138 via the UI 136. The monitoring service 203 might also provide the data to one or more other services, devices, or components that in turn execute a workflow. For example, the metric data may be utilized to trigger one or more alarms, and/or cause one or more other task flows to be performed (e.g., automatically choose a different data store when the data store is experiencing a high volume of requests). Similarly, an occurrence of an event may cause one or more flows to be performed, such as by global workflow service 304 and/or regional

workflow service 318. In some examples, the workflow may be performed by one or more computing resources located in an external network that is part of the global network.

[0125] FIG. 10 is a system and network diagram that shows one illustrative operating environment for the configurations disclosed herein that includes an NMS 102 that can be configured to provide the functionality described above. As discussed above, the NMS 102 can execute network services that provide computing resources for implementing the functionality disclosed herein. The computing resources implemented by the NMS 102 can be data processing resources, such as virtual machine (“VM”) instances, data storage resources, networking resources, data communication resources, network services, and other types of resources.

[0126] The computing resources utilized can be general-purpose or can be available in a number of specific configurations. For example, data processing resources can be available as physical computers or VM instances in a number of different configurations. The VM instances can be configured to execute applications, including web servers, servers, media servers, database servers, some or all of the network services described above, and/or other types of programs. Data storage resources can include file storage devices, block storage devices, and the like. The NMS 102 can also include and utilize other types of computing resources not mentioned specifically herein.

[0127] As also discussed above, the computing resources provided by the NMS 102 are enabled in one implementation by one or more data centers 1004A-1004D (which might be referred to herein singularly as “a data center 1004” or collectively as “the data centers 1004”). The data centers 1004 are facilities utilized to house and operate computer systems and associated components. The data centers 1004 typically include redundant and backup power, communications, cooling, and security systems. The data centers 1004 can also be located in geographically disparate locations. One illustrative configuration for a data center 1004 that can be utilized to implement the technologies disclosed herein will be described below with regard to FIG. 11.

[0128] The users can access the services provided by the NMS 102 over a network 1002, which can be a wide area communication network (“WAN”), such as the Internet, an intranet or an Internet service provider (“ISP”) network or a combination of such networks. For example, and without limitation, a computing device 1000 operated by a user or other user of the NMS 102, such as the computing device 134, can be utilized to access the NMS 102 by way of the network 1002. It should be appreciated that a local-area network (“LAN”), the Internet, or any other networking topology known in the art that connects the

data centers 1004 to remote users and other users can be utilized. It should also be appreciated that combinations of such networks can also be utilized.

[0129] FIG. 11 is a computing system diagram that illustrates examples for a data center 1004 that can be utilized to configure and manage a global network, and the other
5 functionality disclosed herein. The example data center 1004 shown in FIG. 11 includes several server computers 1102A-1102F (which might be referred to herein singularly as “a server computer 1102” or in the plural as “the server computers 1102”).

[0130] The server computers 1102 can be standard tower, rack-mount, or blade server computers configured appropriately for providing various types of computing resources
10 1110 for implementing the functionality disclosed herein. As mentioned above, the computing resources 1110 provided by the data center 1004 can be data processing resources such as VM instances or hardware computing systems, data storage resources, database resources, networking resources, and others. Some of the servers 1102 can also be configured to execute network services 1112A-1112E, respectively, capable of
15 instantiating, providing and/or managing the computing resources 1110A-1110E.

[0131] The data center 1004 shown in FIG. 11 also includes a server computer 1102F that can execute some or all of the software components described above. For example, and without limitation, the server computer 1102F can be configured to execute functionality described herein, such as network management system functionality 1120 and
20 other available services 1122. The server computer 1102F can also be configured to execute other components and/or to store data for providing some or all of the functionality described herein. In this regard, it should be appreciated that components or different instances of the services can execute on many other physical or virtual servers in the data centers 1004 in various configurations.

[0132] In the example data center 1004 shown in FIG. 11, an appropriate LAN 1108 is also utilized to interconnect the server computers 1102A-1102F. The LAN 1108 is also connected to the network 1002 illustrated in FIG. 10. It should be appreciated that the configuration of the network topology described herein has been greatly simplified and that
25 many more computing systems, software components, networks, and networking devices can be utilized to interconnect the various computing systems disclosed herein and to provide the functionality described above. Appropriate load balancing devices or other types of network infrastructure components can also be utilized for balancing a load between each of the data centers 1004A-1004D, between each of the server computers
30 1102A-1102F in each data center 1004, and, potentially, between computing resources

1110 in each of the data centers 1004. It should be appreciated that the configuration of the data center 1004 described with reference to FIG. 11 is merely illustrative and that other implementations can be utilized.

[0133] FIG. 12 shows an example computer architecture for a computer 1200 capable of
5 executing program components for implementing the functionality described above. The computer architecture shown in FIG. 12 illustrates a conventional server computer, workstation, desktop computer, laptop, tablet, network appliance, e-reader, smartphone, or other computing device, and can be utilized to execute any of the software components presented herein.

10 [0134] The computer 1200 includes a baseboard 1202, or “motherboard,” which is a printed circuit board to which a multitude of components or devices can be connected by way of a system bus or other electrical communication paths. In one illustrative configuration, one or more central processing units (“CPUs”) 1204 operate in conjunction with a chipset 1206. The CPUs 1204 can be standard programmable processors that
15 perform arithmetic and logical operations necessary for the operation of the computer 1200.

[0135] The CPUs 1204 perform operations by transitioning from one discrete, physical state to the next through the manipulation of switching elements that differentiate between and change these states. Switching elements can generally include electronic circuits that maintain one of two binary states, such as flip-flops, and electronic circuits that provide an
20 output state based on the logical combination of the states of one or more other switching elements, such as logic gates. These basic switching elements can be combined to create more complex logic circuits, including registers, adders-subtractors, arithmetic logic units, floating-point units, and the like.

[0136] The chipset 1206 provides an interface between the CPUs 1204 and the remainder
25 of the components and devices on the baseboard 1202. The chipset 1206 can provide an interface to a RAM 1208, used as the main memory in the computer 1200. The chipset 1206 can further provide an interface to a computer-readable storage medium such as a read-only memory (“ROM”) 1210 or non-volatile RAM (“NVRAM”) for storing basic process that help to startup the computer 1200 and to transfer information between the
30 various components and devices. The ROM 1210 or NVRAM can also store other software components necessary for the operation of the computer 1200 in accordance with the configurations described herein.

[0137] The computer 1200 can operate in a networked environment using logical connections to remote computing devices and computer systems through a network, such

as the network 1208. The chipset 1206 can include functionality for providing network connectivity through a NIC 1212, such as a gigabit Ethernet adapter. The NIC 1212 is capable of connecting the computer 1200 to other computing devices over the network 1208. It should be appreciated that multiple NICs 1212 can be present in the computer 1200, connecting the computer to other types of networks and remote computer systems.

[0138] The computer 1200 can be connected to a mass storage device 1218 that provides non-volatile storage for the computer. The mass storage device 1218 can store an operating system 1220, programs 1222, workload control user interface 1224, and data, which have been described in greater detail herein. The mass storage device 1218 can be connected to the computer 1200 through a storage controller 1214 connected to the chipset 1206. The mass storage device 1218 can consist of one or more physical storage units. The storage controller 1214 can interface with the physical storage units through a serial attached SCSI (“SAS”) interface, a serial advanced technology attachment (“SATA”) interface, a fiber channel (“FC”) interface, or other type of interface for physically connecting and transferring data between computers and physical storage units.

[0139] The computer 1200 can store data on the mass storage device 1218 by transforming the physical state of the physical storage units to reflect the information being stored. The specific transformation of physical state can depend on various factors, in different implementations of this description. Examples of such factors can include, but are not limited to, the technology used to implement the physical storage units, whether the mass storage device 1218 is characterized as primary or secondary storage, and the like.

[0140] For example, the computer 1200 can store information to the mass storage device 1218 by issuing instructions through the storage controller 1214 to alter the magnetic characteristics of a particular location within a magnetic disk drive unit, the reflective or refractive characteristics of a particular location in an optical storage unit, or the electrical characteristics of a particular capacitor, transistor, or other discrete component in a solid-state storage unit. Other transformations of physical media are possible without departing from the scope and spirit of the present description, with the foregoing examples provided only to facilitate this description. The computer 1200 can further read information from the mass storage device 1218 by detecting the physical states or characteristics of one or more particular locations within the physical storage units.

[0141] In addition to the mass storage device 1218 described above, the computer 1200 can have access to other computer-readable storage media to store and retrieve information, such as program modules, data structures, or other data. It should be appreciated by those

skilled in the art that computer-readable storage media is any available media that provides for the non-transitory storage of data and that can be accessed by the computer 1200.

[0142] By way of example, and not limitation, computer-readable storage media can include volatile and non-volatile, removable and non-removable media implemented in any method or technology. Computer-readable storage media includes, but is not limited to, RAM, ROM, erasable programmable ROM (“EPROM”), electrically-erasable programmable ROM (“EEPROM”), flash memory or other solid-state memory technology, compact disc ROM (“CD-ROM”), digital versatile disk (“DVD”), high definition DVD (“HD-DVD”), BLU-RAY, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to store the desired information in a non-transitory fashion.

[0143] As mentioned briefly above, the mass storage device 1218 can store an operating system 1220 utilized to control the operation of the computer 1200. According to examples, the operating system comprises the LINUX operating system or one of its variants. According to another configuration, the operating system comprises the WINDOWS® SERVER operating system from MICROSOFT Corporation. According to further configurations, the operating system can comprise the UNIX operating system or one of its variants. It should be appreciated that other operating systems can also be utilized. The mass storage device 1218 can store other system or application programs and data utilized by the computer 1200.

[0144] In examples, the mass storage device 1218 or other computer-readable storage media is encoded with computer-executable instructions which, when loaded into the computer 1200, transform the computer from a general-purpose computing system into a special-purpose computer capable of implementing the configurations described herein. These computer-executable instructions transform the computer 1200 by specifying how the CPUs 1204 transition between states, as described above. According to examples, the computer 1200 has access to computer-readable storage media storing computer-executable instructions which, when executed by the computer 1200, perform the various processes described above with regard to FIGS. 1-12. The computer 1200 can also include computer-readable storage media for performing any of the other computer-implemented operations described herein.

[0145] The computer 1200 can also include one or more input/output controllers 1216 for receiving and processing input from a number of input devices, such as a keyboard, a mouse, a touchpad, a touch screen, an electronic stylus, or other type of input device.

Similarly, an input/output controller 1216 can provide output to a display, such as a computer monitor, a flat-panel display, a digital projector, a printer, or other type of output device. It will be appreciated that the computer 1200 might not include all of the components shown in FIG. 12, can include other components that are not explicitly shown in FIG. 12, or can utilize an architecture completely different than that shown in FIG. 12.

Example Clauses

[0146] A. A system comprising: one or more processors; and one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: receive, from a computing device of a user of a service provider network, a request to configure a global network that spans across a first geographic location and a second geographic location, and that includes a stand-alone network that is hosted by the service provider network and an external network that is separate from the service provider network; identify a first action to perform to configure the global network within one or more of the first geographic location, the second geographic location, the stand-alone network, and the external network, wherein the first action is selected from one or more of configuring a gateway to connect to the stand-alone network, configuring the gateway to connect to a first network in the first geographic location and a second network in the second geographic location, configuring a virtual private cloud (VPC), configuring a virtual private network (VPN), and configuring a connection from the service provider network to the external network; transmit instructions, to one or more computing devices located in one or more of the first network, the second network, the stand-alone networks, and the external network, to perform the first action; receive, from the one or more computing devices, monitoring data that indicates one or more of an occurrence of an event indicating a change to the global network and a performance of: gateways that connect different networks of the global network; VPCs; VPNs; devices within the global network; connections to the stand-alone networks; and paths between different endpoints within the global network; identify a loss of network connectivity between different network endpoints in the one or more different geographic locations based on the monitoring data; and cause one or more second actions to be performed within at least one of the different geographic locations in response to the loss of network connectivity.

[0147] B. The system of clause A, wherein the instructions further cause the one or more processors to provide for display, to the computing device of the user, a graphical user interface (GUI) that includes user interface (UI) elements relating to configuring the global

network, and querying the global network, and a graph display area that depicts a graphical representation of the global network, wherein the UI elements include a search UI element that is configured to receive a search term utilized to search the graph for one or more nodes of the graph that match the search term.

5 [0148] C. The system of clause B, wherein the instructions further cause the one or more processors to perform a static network reachability analysis to identify network connectivity between the different endpoints, wherein the static network reachability analysis utilizes formal methods.

[0149] D. The system of any of clauses A–C, wherein causing the one or more second
10 actions to be performed includes executing a workflow that causes tasks to be performed within the first network and the second network to restore the network connectivity between the different network endpoints.

[0150] E. A computer-implemented method comprising: receiving, at a service provider network, data associated with a request to configure a global network that includes a stand-
15 alone network, a first network that is hosted by the service provider network within a first geographic location, and a second network that is hosted by the service provider network within a second geographic location; identifying based, at least in part, on the data, one or more actions to perform within one or more of the first network and the second network to configure the global network; and cause the one or more actions to be performed within
20 the one or more of the first network and the second network to configure the global network.

[0151] F. The computer-implemented method of clause E, wherein the one or more actions include one or more of configuring a gateway of service provider network to connect to at least one of the one or more stand-alone networks, configuring the gateway
25 to connect to the first network and the second network, configuring a virtual private cloud (VPC), configuring a virtual private network (VPN) to connect to the stand-alone network, and configuring a connection between the service provider network and the stand-alone network.

[0152] G. The computer-implemented method of clause E or F, wherein the stand-alone
30 network includes a first external network that is separate from the service provider network and that is connected to the first network, and a second external network that is separate from the service provider network and that is connected to the second network.

[0153] H. The computer-implemented method of any of clauses E–G, further comprising providing an application programming interface (API) that exposes

functionality for configuring the global network, and wherein the data is received via the API.

5 [0154] I. The computer-implemented method of any of clauses E–H, further comprising: receiving, from one or more first computing devices located in the first network, first monitoring data that indicates a first performance of first networking resources in the first network; receiving, from one or more second computing devices located in the second network, second monitoring data that indicates a second performance of second networking resources in the second network; and causing one or more second actions to be performed within one or more of the first network and the second geographical network based, at least in part, on the first monitoring data and the second monitoring data.

10 [0155] J. The computer-implemented method of clause H, wherein causing the one or more second actions to be performed includes executing a workflow within the service provider network that causes tasks to be performed within the first network and the second geographical network to increase one or more of the first performance or the second performance.

[0156] K. The computer-implemented method of any of clauses E–J, further comprising performing a static analysis of a mathematical representation of the global network to identify network connectivity between at least two network endpoints, and wherein identifying the one or more actions to perform is based at least in part on the static analysis.

20 [0157] L. The computer-implemented method of any of clauses E–K, further comprising causing a graphical user interface (GUI) to be displayed by a computing device associated with a user of the service provider network, wherein the GUI includes user interface (UI) elements relating to configuring first networking resources of the global network in the first network and second networking resources of the global network in the second network.

[0158] M. The computer-implemented method of any of clauses E–L, wherein the GUI further comprises a graph display area that displays a graphical representation of the global network.

30 [0159] N. A system comprising: one or more processors associated with a service provider network; and one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: receive, at a service provider network, data associated with configuration of a global network that includes a stand-alone network, a first network that is hosted by the service provider network within a first geographic location and a second network that is

hosted by the service provider network within a second geographic location; identify based, at least in part, on the data, one or more actions to perform within one or more of the first network, the second geographical network, and the stand-alone network, to configure the global network; and cause the one or more actions to be performed within
5 the one or more of the first network, the second geographical network, and the stand-alone network to configure the global network.

[0160] O. The system of clause N, wherein one or more actions includes one or more of configuring a gateway of service provider network to connect to at least one of the one or more stand-alone networks, configuring the gateway to connect to the first network and the
10 second network, configuring a virtual private cloud (VPC), configuring a virtual private network (VPN) to connect to the stand-alone network, and configuring a connection between the service provider network and the stand-alone network that is separate from the service provider network.

[0161] P. The system of clause N or O, wherein the stand-alone network includes a first
15 external network that is separate from the service provider network and is connected to the first network, and a second external network that is separate from the service provider network and is connected to the second network.

[0162] Q. The system of any of clauses N–P, wherein the instructions further cause the processor to provide an application programming interface (API) that exposes functionality
20 for configuring the global network, and wherein the data is received via the API.

[0163] R. The system of any of clauses N–Q, wherein the instructions further cause the processor to:

[0164] receive monitoring data that includes metrics data that indicates performance of computing resources in the global network and event data that indicates an occurrence of
25 one or more events within the global network that identify a change to a topology of the global network; and

[0165] executing one or more workflows based, at least in part, on the monitoring data, wherein the one or more workflows cause one or more actions to be performed within one or more of the first network and the second network.

30 [0166] S. The system of any of clauses N–R, wherein the instructions further cause the processor to cause a graphical user interface (GUI) to be displayed by a computing device associated with a user of the service provider network, wherein the GUI includes user interface (UI) elements relating to configuring first networking resources of the global

network in the first network and second networking resources of the global network in the second network.

[0167] T. The system of any of clauses N–S, wherein the GUI further comprises a graph display area that displays a graphical representation of the global network.

5 [0168] A1. A system comprising: one or more processors; and one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: provide for display, to a computing device of a user of a service provider network, a graphical user interface (GUI) that includes user interface (UI) elements relating to configuring a global network that includes one or
10 more stand-alone networks, the global network spanning across a first network that is hosted by the service provider network within a first geographic location and a second network that is hosted by the service provider network within a second geographic location; generate a graph that depicts the global network, wherein the graph includes first nodes that represent gateways, second nodes that represent stand-alone networks, and one or more
15 of third nodes and edges that represent connections between the first network, the second network, and the one or more stand-alone networks; provide, to the computing device of the user, the graph for display within the GUI; receive, via one or more of the UI elements of the GUI, configuration data associated with configuration of the global network; identify, based on the configuration data, an action to perform within one or more of the
20 first network, the second network, and the one or more stand-alone networks, wherein the action is selected from one or more of performing a configuring one or more of the gateways, configuring one or more of the stand-alone networks, configuring one or more of the connections, and configuring one or more of the links; transmit instructions, to one or more of the stand-alone networks, the first network, and the second network, to cause
25 the action to be performed; generate an updated graph that reflects a current configuration of the global network; and provide, to the computing device of the user, the updated graph for display within the GUI.

[0169] B1. The system of clause A1, wherein the instructions further cause the one or more processors to: receive, from the computing device of the user, second data that
30 indicates a request to generate a second graph of at least one of the stand-alone networks; generate the second graph that represents the at least one of the stand-alone networks; generate a second updated graph that includes the updated graph and the second graph; and provide, to the computing device of the user, the second updated graph for display within the GUI.

[0170] C1. The system of clause A1 or B1, wherein the instructions further cause the one or more processors to: receive, from the first network, first monitoring data that indicates first operating characteristics of first networking resources in the first network; receive, from the second network, second monitoring data that indicates second operating characteristics of second networking resources in the second network; and cause one or more second actions to be performed within one or more of the first network and the second geographical network based on one or more of the first monitoring data and the second monitoring data.

[0171] D1. The system of clauses C1, wherein the instructions further cause the one or more processors to: generate an annotated graph that reflects at least a portion of the first monitoring data and the second monitoring data; and provide for display, to the computing device of the user, the annotated graph for display within the GUI.

[0172] E1. A computer-implemented method comprising: generating a graph of a global network that includes a stand-alone network and that spans across a first network that is hosted by a service provider network within a first geographic location and a second network that is hosted by the service provider network within a second geographic location; providing for display, to a computing device of a user of the service provider network, a graphical user interface (GUI) that includes a display of at least a portion of the graph and one or more user interface (UI) elements relating to configuring the global network; receiving, from the computing device of the user, configuration data associated with configuration of the global network via one or more of an interaction with the at least the portion of the graph and the one or more UI elements; identifying, based at least in part on the configuration data, an action to perform within one or more of the first network and the second network, wherein the action is selected from one or more of configuring one or more gateways, configuring the stand-alone network, and configuring one or more connections; and causing the action to be performed within one or more of the first network and the second network.

[0173] F1. The computer-implemented method of clause E1, further comprising: generating an updated graph that reflects a current configuration of the global network; and providing, to the computing device of the user, the updated graph for display within the GUI.

[0174] G1. The computer-implemented method of clause E1 or F1, wherein generating the graph includes generating one or more first nodes that represent the one or more gateways, one or more second nodes that represent the one or more stand-alone networks,

and one or more of third nodes and edges that represent the one or more connections between the one or more stand-alone networks and the service provider network.

[0175] H1. The computer-implemented method of clause G1, further comprising: receiving, via the GUI, second data that indicates a selection of at least one of the one more
5 first nodes that represent at least one of the one or more gateways; and providing for display within the GUI, a gateway graph view that depicts connections of the at least one of the one or more gateways, wherein the connections include one or more first connections to virtual private clouds (VPCs), one or more second connections to virtual private networks (VPNs), and one or more third connections to the one or more links that connects the stand-
10 alone network.

[0176] I1. The computer-implemented method of any of clauses E1-H1, wherein causing the action to be performed includes transmitting instructions, to the stand-alone network, the first network, and the second network, to perform the action.

[0177] J1. The computer-implemented method of any of clauses E1-I1, further
15 comprising: generating a second graph that represents a second stand-alone network that is separate from the global network; generating an updated graph of the global network that includes the second graph; and providing, to the computing device of the user, the updated graph for display within the GUI.

[0178] K1. The computer-implemented method of any of clauses E1-J1, further
20 comprising: receiving, from the first network, first monitoring data that indicates first operating characteristics of first networking resources in the first network; receiving, from the second network, second monitoring data that indicates second operating characteristics of second networking resources in the second network; and causing one or more second actions to be performed based, at least in part, on one or more of the first monitoring data
25 and the second monitoring data.

[0179] L1. The computer-implemented method of any of clauses E1-K1, further comprising: generating an annotated graph that reflects at least a portion of the first monitoring data and the second monitoring data; and providing, to the computing device of the user, the annotated graph for display within the GUI.

30 [0180] M1. The computer-implemented method of any of clauses E1-L1, further comprising: receiving, via the GUI, a search term; receiving search data in response to

performing a search of the graph using the search term; and providing for display within the GUI, at least a portion of the search data.

[0181] N1. A system comprising: one or more processors associated with a service provider network; and one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: generate a graph of a global network that includes a stand-alone network and that spans across a first network that is hosted by a service provider network within a first geographic location and a second network that is hosted by the service provider network within a second geographic location; provide for display, to a computing device of a user of the service provider network, a graphical user interface (GUI) that includes a display of at least a portion of the graph and one or more user interface (UI) elements relating to configuring the global network; receive, from the computing device of the user, configuration data associated with configuration of the global network via one or more of an interaction with the at least the portion of the graph and the one or more UI elements; identify, based at least in part on the configuration data, an action to perform within one or more of the first network, the second network, and the stand-alone network, wherein the action is selected from one or more of configuring one or more gateways, configuring the stand-alone network, and configuring one or more connections; and cause the action to be performed within one or more of the first network and the second network.

[0182] O1. The system of clause N1, wherein the instructions further cause the one or more processors to: generate an updated graph that reflects a current configuration of the global network; and provide, to the computing device of the user, the updated graph for display within the GUI.

[0183] P1. The system of clause N1 or O1, wherein generating the graph includes generating one or more first nodes that represent the one or more gateways, one or more second nodes that represent the stand-alone network, and one or more of third nodes and edges that represent the one or more connections between one or more of the stand-alone network and the service provider network and the first network and the second network.

[0184] Q1. The system of any of clauses N1-P1, wherein the instructions further cause the one or more processors to: receive, via the GUI, second data that indicates a selection of at least one of the one more first nodes that represent at least one of the one or more gateways; and provide for display within the GUI, a gateway graph view that shows connections of the at least one of the one or more gateways, wherein the connections include one or more first connections to virtual private clouds (VPCs), one or more second

connections to virtual private networks (VPNs), and one or more third connections to one or more links.

[0185] R1. The system of any of clauses N1-Q1, wherein the instructions further cause the one or more processors to: generate a second graph that represents a second stand-alone
5 network that is separate from the global network; generate an updated graph of the global network that includes the second graph; and provide, to the computing device of the user, the updated graph for display within the GUI.

[0186] S1. The system of clause N1, wherein the instructions further cause the one or more processors to: receive, from the first network, first monitoring data that indicates first
10 operating characteristics of first networking resources in the first network; receive, from the second network, second monitoring data that indicates second operating characteristics of second networking resources in the second network; and cause one or more second actions to be performed based, at least in part, on one or more of the first monitoring data and the second monitoring data.

15 [0187] T1. The system of any of clauses N1-S1, wherein the instructions further cause the one or more processors to: generate an annotated graph that reflects at least a portion of the first monitoring data and the second monitoring data; and provide, to the computing device of the user, the annotated graph for display within the GUI.

[0188] A2. A system comprising: one or more processors; and one or more computer-
20 readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: identify network endpoints within a global network to monitor for connectivity, wherein each of the network endpoints are at different locations within the global network, and wherein the global network includes stand-alone networks and spans across a first network that is hosted by a service provider
25 network within a first geographic location and a second network that is hosted by the service provider network within a second geographic location; transmit instructions, to the stand-alone networks, the first network, and the second network, to cause monitoring of metrics to be performed for the network endpoints, wherein the metrics includes a flow metric that indicates a flow of data between a gateway that connects the first network and
30 the second network; receive monitoring data that includes the metrics, from the stand-alone networks, the first network and the second network; identify, based on the monitoring data, a loss of network connectivity in a path between at least a first network endpoint and a second network endpoint of the network endpoints; and cause an action to be performed based on the loss of network connectivity, wherein the action is one or more of causing a

workflow to be performed and providing network connectivity data to a computing device associated with a user of the service provider network.

- 5 [0189] B2. The system of clause A2, wherein the first network endpoint is located in the first network and the second network endpoint is located in the second network or one of the stand-alone networks and wherein the instructions further cause the one or more processors to perform a static network reachability analysis to identify at least a portion of the network connectivity data indicating network connectivity between the network endpoints, wherein the static network reachability analysis utilizes formal methods that includes a mathematical representation of the global network.
- 10 [0190] C2. The system of clause A2 or B2, wherein the instructions further cause the one or more processors to identify a different path between the first network endpoint and the second network endpoint, wherein the workflow, when performed by the service provider network, selects the different path to utilize between the first network endpoint and the second network endpoint.
- 15 [0191] D2. The system of any of clauses A2–C2, wherein the instructions further cause the one or more processors to receive, from a computing device of the user, network reachability data that identifies at least a portion of the network endpoints to monitor for connectivity.
- 20 [0192] E2. A computer-implemented method comprising: identifying network endpoints within a global network to monitor for connectivity, wherein the global network includes a stand-alone network and spans across a first network that is hosted by the service provider network within a first geographic location and a second network that is hosted by the service provider network within a second geographic location; identifying, based at least in part on monitoring data including metrics associated with the network endpoints, a loss of network connectivity for a path between at least a first network endpoint and a second network endpoint of the network endpoints; causing an action to be performed based, at least in part, on loss of the network connectivity, wherein the action is one or more of causing a workflow to be performed and providing network connectivity data to a computing device associated with a user of the service provider network.
- 25 [0193] F2. The computer-implemented method of clause E2, further comprising: transmitting instructions, to the stand-alone network, the first network, and the second network, to cause monitoring to be performed to generate the monitoring data; receiving a first portion of the monitoring data from the stand-alone network; receiving a second
- 30

portion of the monitoring data from the first network; and receiving a third portion of the monitoring data from the second network.

[0194] G2. The computer-implemented method of clause E2 or F2, further comprising performing static network reachability analysis to identify at least a portion of the network connectivity data indicating network connectivity between at least a portion of the network endpoints, wherein the static network reachability analysis utilizes a mathematical representation of the global network.

[0195] H2. The computer-implemented method of any of clauses E2–G2, wherein the workflow, when performed, changes the path to a second path to utilize between the first network endpoint and the second network endpoint.

[0196] I2. The computer-implemented method of any of clauses E2–H2, further comprising providing an input mechanism to specify at least a portion of the network endpoints, and metrics to monitor within the stand-alone network, the first network, and the second network, wherein the input mechanism is one or more of a graphical user interface (GUI), a command line interface (CLI), and an application programming interface (API) that expose functionality.

[0197] J2. The computer-implemented method of any of clauses E2–I2, wherein the monitoring data includes metrics data associated with a flow of data between a gateway that connects the first network and the second network.

[0198] K2. The computer-implemented method of any of clauses E2–J2, wherein identifying network endpoints within the global network to monitor for connectivity comprises receiving, from the computing device of the user, network reachability data that identifies at least a portion of the network endpoints to monitor for connectivity.

[0199] L2. The computer-implemented method of any of clauses E2–K2, further comprising: generating a graph that includes network connectivity data indicating the network connectivity data for at least a portion of the global network; and providing, to the computing device of the user, the graph for display within a graphical user interface (GUI).

[0200] M2. A system comprising: one or more processors associated with a service provider network; and one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: identifying network endpoints within a global network to monitor for connectivity, wherein the global network includes a stand-alone network and spans across a first network that is hosted by the service provider network within a first geographic location and a second network that is hosted by the service provider network within a

second geographic location; identifying, based at least in part on monitoring data associated with the network endpoints, a loss of network connectivity for a path between at least a first network endpoint and a second network endpoint of the network endpoints, wherein the monitoring data includes one or more of metrics data and events data; and causing an action to be performed based, at least in part, on the loss of network connectivity, wherein the action is one or more of causing a workflow to be performed and providing network connectivity data to a computing device associated with a user of the service provider network.

5
10 [0201] N2. The system of clause M2, wherein the instructions further cause the one or more processors to: transmit instructions, to the stand-alone network, the first network, and the second network, to cause monitoring of metrics to be performed to generate the monitoring data; receive a first portion of the monitoring data from the stand-alone network; receive a second portion of the monitoring data from the first network; and receive a third portion of the monitoring data from the second network.

15 [0202] O2. The system of clause N2, wherein the first network endpoint is located in the first network and the second network endpoint is located in the second network or the stand-alone network.

[0203] P2. The system of any of clauses M2–O2, wherein the monitoring data includes metrics data associated with a flow of data between a gateway that connects the first network and the second network.

20 [0204] Q2. The system of any of clauses M2–P2, wherein the workflow, when performed, changes a path to utilize between the first network endpoint and the second network endpoint.

[0205] R2. The system of any of clauses M2–Q2, wherein the instructions further cause the one or more processors to provide one or more of a graphical user interface (GUI), a command line interface (CLI), and an application programming interface (API) that expose functionality for specifying the network endpoints, and metrics to monitor within the stand-alone network, the first network, and the second network.

25
30 [0206] S2. The system of any of clauses M2–R2, wherein identifying the loss of network connectivity includes performing a static network reachability analysis that utilizes a mathematical representation of the global network and a dynamic network analysis that utilizes metrics data.

[0207] T2. The system of any of clauses M2–S2, wherein the instructions further cause the one or more processors to: generate a graph that includes network connectivity data

indicating the network connectivity data for at least a portion of the global network; and provide, to the computing device of the user, the graph for display within a graphical user interface (GUI).

[0208] A3. A system comprising: a monitoring service configured to monitor a global
5 network that spans across a first geographic location and a second geographic location and that includes first stand-alone networks hosted by a service provider network and second stand-alone networks that are external from the service provider network; a workflow service configured to perform one or more workflows; an event service configured to generate events in response to a change to the global network; a metrics service configured
10 to generate metrics that indicate a performance of gateways that connect different networks of the global network, virtual private clouds (VPCs), virtual private networks (VPNs), devices within the global network, connections to the first stand-alone networks, connections to the second stand-alone networks, and paths between different endpoints within the global network; one or more processors; and one or more computer-readable
15 media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: receive, from a computing device of a user of the service provider network, a request to monitor the global network for metrics associated with; transmit instructions, via the monitoring service to the first and second stand-alone networks, a first network within the first geographic location, and a second
20 network within the second geographic location, to cause monitoring of the metrics to be performed; receive monitoring data for the metrics and the events, from the first and second stand-alone networks, the first network, and the second network; identify, based on the monitoring data, one or more of an occurrence of an event that indicates a network topology change to the global network and a trigger of an alarm that indicates a value of one or more
25 of the metrics has changed and a decrease of performance of the global network has occurred; and cause an action to be performed within the global network, wherein the action is one or more of causing a workflow to be performed by the event service to adjust a configuration of the global network and providing data, indicating the one or more of the occurrence of the event and the trigger of the alarm, to the computing device associated
30 with the user.

[0209] B3. The system of clause A3, wherein receiving the request to monitor the global network for metrics includes receiving the request via an application programming interface (API), and wherein causing the action to be performed comprises transmitting, via the API, an instruction to the event service to perform the workflow to adjust the

configuration of the global network, wherein adjusting the configuration of the global network includes to change one or more paths to use between at least two network endpoints within the global network.

5 [0210] C3. The system of clause A3 or B3, wherein the instructions further cause the one or more processors to receive, from the computing device of the user, alarm data to configure the alarm to trigger based on a value of one or more of the metrics, wherein the value of the one or more metrics is based on a flow of packets through one or more of the gateways.

10 [0211] D3. The system of any of clauses A3–C3, wherein receiving the monitoring data includes receiving a first portion of the monitoring data from a first manager in the first network, receiving a second portion of the monitoring data from a second manager in the second network, and a third portion of the monitoring data from one or more computing device in the first and second stand-alone networks, and wherein providing the data, indicating the one or more of the occurrence of the event and the trigger of the alarm, to
15 the computing device associated with the user includes utilizing an application programming interface (API) to provide the data.

[0212] E3. A computer-implemented method comprising: monitoring a global network to generate monitoring data that includes one or more of metrics associated with the performance of the global network and events that indicate a change to the global network,
20 wherein the global network spans across a first geographic location and a second geographic location and includes a first stand-alone network hosted by a service provider network and a second stand-alone network that is external from the service provider network; receiving the monitoring data from the stand-alone networks, a first network within the first geographic location, and a second network within the second geographic
25 location; identifying, based at least in part on the monitoring data, one or more of an occurrence of an event and a trigger of an alarm; and cause an action to be performed within the global network, wherein the action is one or more of causing a workflow to be performed within the service provider network and providing data, indicating one or more that the alarm is triggered and the occurrence of the event, to a computing device associated
30 with a user of the service provider network.

[0213] F3. The computer-implemented method of clause E3, further comprising receiving, from the computing device of the user, a request to monitor the global network to identify one or more of the one or more metrics and the events.

[0214] G3. The computer-implemented method of clause F3, wherein receiving the request to monitor the global network includes receiving the request via an application programming interface (API), and wherein causing the action to be performed comprises transmitting, via the API, an instruction to an event service hosted by the service provider network to perform the workflow.

[0215] H3. The computer-implemented method of clause F3 or G3, wherein the one or more metrics indicate a performance of one or more of gateways that connect different networks of the global network, virtual private clouds (VPCs) within the global network, virtual private networks (VPNs) within the global network, devices within the global network, connections to one or more of the first and second stand-alone networks, and paths between different endpoints within the global network.

[0216] I3. The computer-implemented method of any of clauses E3–H3, further comprising transmitting instructions to the first and second stand-alone networks, the first network, and the second network to cause monitoring of the one or more metrics to be performed.

[0217] J3. The computer-implemented method of any of clauses E3–I3, further comprising receiving, from the computing device of the user, alarm data to configure the alarm to trigger based on a value of the one or more metrics.

[0218] K3. The computer-implemented method of clause J3, wherein the alarm data specifies the value of the one or more metrics associated with a gateway that connects the first network and the second network.

[0219] L3. The computer-implemented method of any of clauses E3–K3, wherein providing the data, indicating that the alarm is triggered, to the computing device associated with the user includes utilizing an application programming interface (API) to provide the data.

[0220] M3. The computer-implemented method of any of clauses E3–L3, further comprising providing an input mechanism to specify the metrics and the events to monitor within the global network, wherein the input mechanism is one or more of a graphical user interface (GUI), a command line interface (CLI), and an application programming interface (API).

[0221] N3. A system comprising: one or more processors associated with a service provider network; and one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to: monitor a global network to generate monitoring data indicating one or more

of a change to the global network and a performance of the global network, wherein the global network spans across a first geographic location and a second geographic location and includes a first stand-alone network hosted by a service provider network and a second stand-alone network that is external from the service provider network; receive the
5 monitoring data from the first and second stand-alone networks, a first network within the first geographic location, and a second network within the second geographic location; identify, based on the monitoring data, that one or more of an alarm is triggered, or an occurrence of an event; and cause an action to be performed within the global network, wherein the action includes causing a workflow to be performed within the global network.

10 **[0222]** O3. The system of clause N3, wherein the instructions further cause the one or more processors to receive, from the computing device of the user, a request to monitor the global network.

[0223] P3. The system of clause O3, wherein receiving the request to monitor the global network to identify the one more metrics includes receiving the request via an application
15 programming interface (API), and wherein causing the action to be performed comprises transmitting, via the API, an instruction to an event service within the service provider network to perform the workflow.

[0224] Q3. The system of clause O3 or P3, wherein the one or more metrics indicate a performance of one or more of gateways that connect different networks of the global
20 network, virtual private clouds (VPCs) within the global network, virtual private networks (VPNs) within the global network, devices within the global network, connections to the first and second stand-alone networks, and paths between different endpoints within the global network.

[0225] R3. The system of any of clauses N3–Q3, wherein the instructions further cause
25 the one or more processors to transmit instructions to one or more computing devices associated with the stand-alone networks, the first network, and the second network to cause monitoring to identify the one or more metrics to be performed.

[0226] Based on the foregoing, it should be appreciated that technologies for configuration and management of global private networks have been disclosed herein.
30 Moreover, although the subject matter presented herein has been described in language specific to computer structural features, methodological acts, and computer readable media, it is to be understood that the invention defined in the appended claims is not necessarily limited to the specific features, acts, or media described herein. Rather, the

specific features, acts, and media are disclosed as example forms of implementing the claims.

[0227] The subject matter described above is provided by way of illustration only and should not be construed as limiting. Furthermore, the claimed subject matter is not limited
5 to implementations that solve any or all disadvantages noted in any part of this disclosure. Various modifications and changes can be made to the subject matter described herein without following the example configurations and applications illustrated and described, and without departing from the true spirit and scope of the present invention, which is set forth in the following claims.

CLAIMS

WHAT IS CLAIMED IS:

1. A computer-implemented method comprising:
generating a graph of a global network that includes a stand-alone network and that
5 spans across a first network that is hosted by a service provider network within a first
geographic location and a second network that is hosted by the service provider network
within a second geographic location;
providing for display, to a computing device of a user of the service provider
network, a graphical user interface (GUI) that includes a display of at least a portion of the
10 graph and one or more user interface (UI) elements relating to configuring the global
network;
receiving, from the computing device of the user, configuration data associated
with configuration of the global network via one or more of an interaction with the at least
the portion of the graph and the one or more UI elements;
15 identifying, based at least in part on the configuration data, an action to perform
within one or more of the first network and the second network, wherein the action is
selected from one or more of configuring one or more gateways, configuring the stand-
alone network, and configuring one or more connections; and
causing the action to be performed within one or more of the first network and the
20 second network.
2. The computer-implemented method of claim 1, further comprising:
generating an updated graph that reflects a current configuration of the global
network; and
providing, to the computing device of the user, the updated graph for display within
25 the GUI.
3. The computer-implemented method of claim 1, wherein generating the
graph includes generating one or more first nodes that represent the one or more gateways,
one or more second nodes that represent the one or more stand-alone networks, and one or
more of third nodes and edges that represent the one or more connections between the one
30 or more stand-alone networks and the service provider network.
4. The computer-implemented method of claim 3, further comprising:
receiving, via the GUI, second data that indicates a selection of at least one of the
one more first nodes that represent at least one of the one or more gateways; and

providing for display within the GUI, a gateway graph view that depicts connections of the at least one of the one or more gateways, wherein the connections include one or more first connections to virtual private clouds (VPCs), one or more second connections to virtual private networks (VPNs), and one or more third connections to the one or more links that connects the stand-alone network.

5 5. The computer-implemented method of claim 1, wherein causing the action to be performed includes transmitting instructions, to the stand-alone network, the first network, and the second network, to perform the action.

6. The computer-implemented method of claim 1, further comprising:
10 generating a second graph that represents a second stand-alone network that is separate from the global network;
 generating an updated graph of the global network that includes the second graph;
 and
 providing, to the computing device of the user, the updated graph for display within
15 the GUI.

7. The computer-implemented method of claim 1, further comprising:
 receiving, from the first network, first monitoring data that indicates first operating characteristics of first networking resources in the first network;
 receiving, from the second network, second monitoring data that indicates second
20 operating characteristics of second networking resources in the second network; and
 causing one or more second actions to be performed based, at least in part, on one or more of the first monitoring data and the second monitoring data.

8. The computer-implemented method of claim 7, further comprising:
 generating an annotated graph that reflects at least a portion of the first monitoring
25 data and the second monitoring data; and
 providing, to the computing device of the user, the annotated graph for display within the GUI.

9. The computer-implemented method of claim 1, further comprising:
 receiving, via the GUI, a search term;
30 receiving search data in response to performing a search of the graph using the search term; and
 providing for display within the GUI, at least a portion of the search data.

10. A system comprising:
 one or more processors associated with a service provider network; and

one or more computer-readable media storing computer-executable instructions that, when executed by the one or more processors, cause the one or more processors to:

5 generate a graph of a global network that includes a stand-alone network and that spans across a first network that is hosted by a service provider network within a first geographic location and a second network that is hosted by the service provider network within a second geographic location;

10 provide for display, to a computing device of a user of the service provider network, a graphical user interface (GUI) that includes a display of at least a portion of the graph and one or more user interface (UI) elements relating to configuring the global network;

receive, from the computing device of the user, configuration data associated with configuration of the global network via one or more of an interaction with the at least the portion of the graph and the one or more UI elements;

15 identify, based at least in part on the configuration data, an action to perform within one or more of the first network, the second network, and the stand-alone network, wherein the action is selected from one or more of configuring one or more gateways, configuring the stand-alone network, and configuring one or more connections; and

20 cause the action to be performed within one or more of the first network and the second network.

11. The system of claim 10, wherein the instructions further cause the one or more processors to:

25 generate an updated graph that reflects a current configuration of the global network; and

provide, to the computing device of the user, the updated graph for display within the GUI.

30 12. The system of claim 11, wherein generating the graph includes generating one or more first nodes that represent the one or more gateways, one or more second nodes that represent the stand-alone network, and one or more of third nodes and edges that represent the one or more connections between one or more of the stand-alone network and the service provider network and the first network and the second network.

13. The system of claim 11, wherein the instructions further cause the one or more processors to:

receive, via the GUI, second data that indicates a selection of at least one of the one or more first nodes that represent at least one of the one or more gateways; and

provide for display within the GUI, a gateway graph view that shows connections of the at least one of the one or more gateways, wherein the connections include one or more first connections to virtual private clouds (VPCs), one or more second connections to virtual private networks (VPNs), and one or more third connections to one or more links.

14. The system of claim 11, wherein the instructions further cause the one or more processors to:

generate a second graph that represents a second stand-alone network that is separate from the global network;

generate an updated graph of the global network that includes the second graph; and

provide, to the computing device of the user, the updated graph for display within the GUI.

15. The system of claim 11, wherein the instructions further cause the one or more processors to:

receive, from the first network, first monitoring data that indicates first operating characteristics of first networking resources in the first network;

receive, from the second network, second monitoring data that indicates second operating characteristics of second networking resources in the second network; and

cause one or more second actions to be performed based, at least in part, on one or more of the first monitoring data and the second monitoring data.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

500

Global Private Networks 502

Select Global Network 504

SEARCH GLOBAL NETWORKS

514

☒ IDENTIFIER

GLOBAL NETWORK 1

US, EUROPE

REGIONS

STATUS

☒ AVAILABLE

EDIT GLOBAL NETWORK 506

VIEW GLOBAL NETWORK 508

DELETE GLOBAL NETWORK 510

CREATE GLOBAL NETWORK 512

FIG. 5A

520

Global Private Networks

GATEWAYS 522

DEVICES 524

REGIONS 526

CONNECTIONS 528

CONFIGURATION 530

Select Gateway to Register 532

SEARCH BOX FOR GATEWAY

534

	IDENTIFIER	REGION STATE	
☒	GATEWAY 2 – US-EAST	US-EAST	☒ AVAILABLE
☒	GATEWAY 3 – US-WEST	US-WEST	☒ AVAILABLE
☒	GATEWAY 12 – EUROPE	EUROPE	☒ AVAILABLE

CANCEL 536

REGISTER GATEWAY 538

FIG. 5B




FIG. 5C




FIG. 5D

592




FIG. 5E




FIG. 6

11/16




FIG. 7




FIG. 8

13/16




FIG. 9





FIG. 10






FIG. 12

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2020/062347

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L12/24

ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2018/018082 A1 (SARBIN BRIAN [US] ET AL) 18 January 2018 (2018-01-18) paragraph [0003] - paragraph [0091]; figures 1-5 -----	1,3-5,7, 10,12, 13,15
X	US 10 063 427 B1 (BROWN ALLEN HERBERT [US]) 28 August 2018 (2018-08-28) column 2, line 11 - column 21, line 57; figures 1-8 -----	1-3,5,7, 10-12,15
X	US 10 469 304 B1 (KEMPE GREGORY JONATHAN [ZA] ET AL) 5 November 2019 (2019-11-05) column 3, line 1 - column 29, line 23; figures 1-21 ----- -/--	1-3,5,7, 10-12,15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

1 March 2021

Date of mailing of the international search report

10/03/2021

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2

NL - 2280 HV Rijswijk

Tel. (+31-70) 340-2040,

Fax: (+31-70) 340-3016

Authorized officer

Lupia, Sergio

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2020/062347

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2019/268218 A1 (THOMASSON BENJAMIN [US] ET AL) 29 August 2019 (2019-08-29) paragraph [0017] - paragraph [0107]; figures 1-12 -----	1-3,5,7, 10-12,15
X	WO 2016/209317 A1 (DELL PRODUCTS LP [US]) 29 December 2016 (2016-12-29) paragraph [0002] - paragraph [0049]; figures 1-13B -----	1,3,5-7, 10,12, 14,15
X	US 2017/085446 A1 (ZHONG QIANJIE [CN] ET AL) 23 March 2017 (2017-03-23) paragraph [0002] - paragraph [0414]; figures 1-22 -----	1,3,5, 7-10,12, 15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2020/062347

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2018018082 A1	18-01-2018	NONE	
US 10063427 B1	28-08-2018	NONE	
US 10469304 B1	05-11-2019	US 10469304 B1 US 2020127881 A1	05-11-2019 23-04-2020
US 2019268218 A1	29-08-2019	US 2019268218 A1 US 2020295987 A1	29-08-2019 17-09-2020
WO 2016209317 A1	29-12-2016	CN 107710241 A EP 3314550 A1 US 2016380850 A1 WO 2016209317 A1	16-02-2018 02-05-2018 29-12-2016 29-12-2016
US 2017085446 A1	23-03-2017	US 2017085446 A1 US 2020287807 A1	23-03-2017 10-09-2020