

EDGE COMPUTE ENVIRONMENT CONFIGURATION TOOL

Cross-Reference to Related Applications

[0001] This application claims the benefit of U.S. Provisional Application No. 63/132,100 filed 30 December 2020, entitled "Edge Computer Environment Configuration Tool," which is incorporated herein by reference in its entirety. This application is also related to U.S. Patent Application Serial No. 17/388,610, which is incorporated by reference herein for all purposes.

Technical Field

[0002] Embodiments of the present invention generally relate to systems and methods for implementing a telecommunications network, and more specifically for a system to automatically configure an edge compute environment at an edge site of a telecommunications network.

Background

[0003] Telecommunications networks may provide many services to customers or devices connected to the network, including transmission of communications between network devices, network services, remote computing environments, cloud services (such as storage services, networking service, compute services, etc.), and the like. Such telecommunications networks generally include interconnected devices and/or components that are configured to communicate with each other and/or customer devices to provide access to the available services from the network. Typically, configuration of devices and interconnections of the network require a network engineer to remotely or locally access network components and manually configure settings, ports, operating systems, and the like to enable the network to provide services to customers. In some instances, the configuration of network devices includes multiple network administrators and other groups to manually configure the devices. For example, configuration of some network services may require a first network administrator to log into one or more components associated with the services separately after installation of the components into the network and provide one or more inputs via a workstation or other computing device to configure the components according to a service plan as a single step in the overall service configuration process. The first network administrator may then notify another group or administrator of the completion of a step in the configuration process so that the next step in the process may be executed by the second network administrator, and so on. Such a process can be time consuming, require steps or acts from multiple network groups, and

includes multiple potential points of delay or errors that must be identified and corrected before the network service is available to the customer.

[0004] It is with these observations in mind, among other, that aspects of the present disclosure were conceived.

Summary

[0005] One aspect of the present disclosure relates to a method for managing a communications network. The method may include the operations of obtaining, from a plurality of databases, site information associated with an edge site of the communications network, wherein each of the plurality of databases is associated with a different backbone communications network and generating, by a configuration computing system and based on the site information associated with the edge site of the communications network, a configuration process for instantiating a compute service available from the edge site, the configuration process comprising a set of instructions for configuring an operational status of at least one network edge device of the edge site to provide the compute service to a requesting device in communication with the edge site. The method may also include the operation of executing the set of instructions with respect to the configuration data to cause the at least one network edge device to be configured, wherein executing the set of instructions comprises causing at least one micro-service to be executed on a micro-services system that is separate from the configuration computing system.

[0006] Yet another aspect of the present disclosure relates to non-transitory, computer-storage media storing instructions that, when executed by at least one processor, cause the at least one processor to perform a method comprising obtaining, from a plurality of databases, site information associated with an edge site of the communications network, wherein each of the plurality of databases is associated with a different backbone communications network and generating, by a configuration computing system and based on the site information associated with the edge site of the communications network, a configuration process for instantiating a compute service available from the edge site, the configuration process comprising a set of instructions for configuring an operational status of at least one network edge device of the edge site to provide the compute service to a requesting device in communication with the edge site. The method may also include the operation of executing the set of instructions with respect to the configuration data to cause the at least one network edge device to be configured, wherein

executing the set of instructions comprises causing at least one micro-service to be executed on a micro-services system that is separate from the configuration computing system.

[0007] Another aspect of the present disclosure relates to a system comprising at least one processor and memory, storing instructions that, when executed, cause the at least one processor to execute one or more operations. Such operations may include obtaining, from a plurality of databases, site information associated with an edge site of the communications network, wherein each of the plurality of databases is associated with a different backbone communications network and generating, by a configuration computing system and based on the site information associated with the edge site of the communications network, a configuration process for instantiating a compute service available from the edge site, the configuration process comprising a set of instructions for configuring an operational status of at least one network edge device of the edge site to provide the compute service to a requesting device in communication with the edge site. The operations may also include the operation of executing the set of instructions with respect to the configuration data to cause the at least one network edge device to be configured, wherein executing the set of instructions comprises causing at least one micro-service to be executed on a micro-services system that is separate from the configuration computing system.

Brief Description Of The Drawings

[0008] Figure 1A is a schematic diagram illustrating a system for configuring an edge compute environment at an edge of a network in accordance with one embodiment.

[0009] Figure 1B is a schematic diagram illustrating a system for configuring an edge compute environment at an edge of a network in accordance with one embodiment.

[0010] Figure 2 is a schematic diagram illustrating an edge compute environment of a network in accordance with one embodiment.

[0011] Figure 3A is a flowchart illustrating a method for configuring an edge compute environment at an edge of a network in accordance with one embodiment.

[0012] Figure 3B is a flowchart illustrating a method for generating the configuration process and files to configure switches of the edge compute network in a method for configuring an edge compute environment at an edge of a network in accordance with one embodiment.

[0013] Figure 4 is an example user interface for providing seeding data and information for an edge compute environment configuration in accordance with one embodiment.

[0014] Figure 5 is an example user interface for providing new site seeding data and information for an edge compute environment configuration in accordance with one embodiment.

[0015] Figure 6 is an example user interface for providing new router group seeding data and information for an edge compute environment configuration in accordance with one embodiment.

[0016] Figure 7 is an example user interface for providing new device seeding data and information for an edge compute environment configuration in accordance with one embodiment.

[0017] Figure 8 is an example user interface displaying site data and configuration progress for an edge compute environment configuration in accordance with one embodiment.

[0018] Figure 9 is an example user interface displaying edge device data and configuration progress for an edge compute environment configuration in accordance with one embodiment.

[0019] Figure 10 is an example user interface displaying data and configuration progress for multiple devices of an edge compute environment configuration in accordance with one embodiment.

[0020] Figure 11 is an example user interface displaying data and configuration progress for a service installation of an edge compute environment configuration in accordance with one embodiment.

[0021] Figure 12 is an example user interface displaying data and configuration progress for multiple interfaces associated with an edge compute environment configuration in accordance with one embodiment.

[0022] Figure 13 is an example user interface for providing data and information for updating a device configuration for an edge compute environment in accordance with one embodiment.

[0023] Figure 14 is an example user interface for providing port configuration data and information for a device of an edge compute environment configuration in accordance with one embodiment.

[0024] Figure 15 is an example user interface for importing an existing port configuration from one device of an edge compute environment to another device of the edge compute environment in accordance with one embodiment.

[0025] Figure 16 is a diagram illustrating an example of a computing system which may be used in implementing embodiments of the present disclosure.

[0026] Detailed Description

[0027] Aspects of the present disclosure involve a tool for configuring an edge compute environment of a network, such as a communications network or other computing environment. In one instance, the edge compute network configuration tool may generate a configuration process for instantiating an edge compute environment at an edge site of a network. The process may involve configuring one or more of the components of the edge compute environment. The configuration process may be based on seeding data and/or information provided to the edge compute network configuration tool via a user interface and/or from one or more databases associated with the network. The configuration process may include generating or otherwise providing automatically executed configuration instructions or steps that communicate with the devices of the edge compute environment and configure operational processes of the devices, provision communication ports and interfaces between the devices, establish one or more network addresses with the devices, and the like. In general, any operation performed at the device or in the broader network may be included in the configuration process to instantiate an edge compute service in an edge site of the network. In some instances, the edge compute network configuration tool may access and manage one or more micro-services to communicate with and control configuration of the devices of the edge compute environment. Through the configuration process, compute services may be provided to customers of the network via the network edge site. Providing compute services from devices or environments located at the edge of the network may reduce the latency/cost experienced at customer devices or networks during use of the compute services.

[0028] In one particular instance, a user interface may be associated with the edge compute network configuration tool. The user interface may include one or more portions through which seeding data and/or information for the edge compute environment may be provided. The user interface may be displayed on a computing device and may interact with a user of the computing device, the computing device itself, or a different computing device to receive the seeding information. The edge compute network configuration tool may also display information of the edge compute environment via the user interface, including information of the devices of the environment, information of communication or interfacing standards used by the environment, network infrastructure and interconnection information and the like. Alterations to such information may be accessed and provided via the user interface. In still other instances, the generated configuration process may include reporting and/or logging of a configuration status of the devices of the edge compute environment. For example, one or more of the

devices of the edge compute environment may provide a report of a configuration status to the edge compute network configuration tool indicating a successful configuration or a failed configuration. The edge compute network configuration tool may display the status information for the configuration via the user interface and one or more actions may be executed in response to the displayed status information. Through the edge compute network configuration tool, configuration of an edge compute environment of a network may be executed quickly and reliably for providing compute services to customers of a network edge site.

[0029] Figure 1A is a schematic diagram illustrating a system 100 for configuring an edge compute environment at an edge or edge site of a network in accordance with one embodiment. In general, the system 100 may include an edge compute network configuration tool 106 in communication with a source of information such as one or more network databases 104a-c associated with one or more respective backbone networks 102a-c. The various backbone networks 102a-c associated with the system 100 may include both public communication networks (such as the Internet) and/or private communication networks (such as a virtual private network (VPN) or other type of private network). The edge compute network configuration tool 106 may also be in communication with a computing device 122 providing a user interface 124. The interactions and communications between the components of the edge compute configuration system 100 is described in more detail herein.

[0030] As explained above, the edge compute network configuration tool 106 may access and configure one or more network devices associated with an edge site of a network or networks to provide compute services from the edge site. By providing the compute services from the edge of the network or networks, such services may be provided with minimal or reduced latency/cost in receiving the services from the network. In some instances, the edge compute network configuration tool 106 may include an edge compute network configuration application 112 executed to perform one or more of the operations described herein. The edge compute network configuration application 112 may be stored in a computer readable media 110 (e.g., memory) and executed on a processing system 108 of the edge compute network configuration tool 106 or other type of computing system, such as that described below. For example, the edge compute network configuration application 112 may include instructions that may be executed in an operating system environment, such as a Microsoft Windows™ operating system, a Linux operating system, or a UNIX operating system environment. The computer readable medium 110 includes volatile media, nonvolatile media, removable media, non-removable media, and/or another available medium. By way of example and not limitation, non-

transitory computer readable medium 110 comprises computer storage media, such as non-transient storage memory, volatile media, nonvolatile media, removable media, and/or non-removable media implemented in a method or technology for storage of information, such as computer readable instructions, data structures, program modules, or other data.

[0031] The edge compute network configuration application 112 may also utilize a data source 120 of the computer readable media 110 for storage of data and information associated with the edge compute network configuration tool 106. For example, the edge compute network configuration application 112 may store information associated one or more networks 102a-c over which a communications or data path may be defined. In general, any data or information utilized by the edge compute network configuration application 106 may be stored and/or retrieved via the data source 120. In one particular implementation, seeding information or other data may be received via a user device 124 (described in more detail below) and/or one or more network databases 104a-c and stored in the data source 120 for use in configuring components of an edge compute environment.

[0032] The edge compute network configuration application 112 may include several components to perform one or more of the operations described herein. For example, a database translator 126 may be included in the edge compute network configuration application 112 to translate network information obtained from the databases 104 associated with one or more communication networks 102. In the example shown in Figure 1A, the edge compute network configuration tool 106 may communicate with a database 104a associated with a network A 102a, a database 104b associated with a network B 102b, and a database 104c associated with a network C 102c. The edge compute network configuration tool 106 may communicate with more or fewer such databases 104 in more or fewer networks 102. In general, each of the networks 102 comprise a collection of interconnected network components to transmit messages or packets within the networks and may include both public and/or private network portions. In some instances, however, network operators may purchase existing networks and integrate those networks into a main network to expand the geographical footprint of the main network. These purchased or obtained networks may include databases 104 with network component information, such as component locations, interconnectivity information, component types, port assignments of the components, and the like. In addition, each network 102 may maintain or store the network component information in a particular format in the associated database 104 or databases for that network. Such formats may vary from network to network such that information may be stored in a first database, such as database 104a, in a

first format while information stored in a second database, such as database 104b, may be in a second format different than the first format.

[0033] As the network information in the various network databases 104 in communication with the edge compute network configuration tool 106 may be stored in various formats, the database translator 126 of the edge compute network configuration application 112 may translate the network information from the databases 104 into a common format and store the translated information in data source 120. In some instances, the database translator 126 may translate the network information prior to storage in the data source 120. In another implementation, the network information may be stored in the data source 120 as obtained from the respective database 104 and translated when accessed by the application 112 during edge compute configuration. Regardless, the database translator 126 may be configured to identify a format of stored network information and translate the obtained stored information to the common format.

[0034] The edge compute network configuration application 112 may also include a user interface communicator 118 to communicate with a computing device 122 executing a user interface program 124. The computing device 122 may provide the user interface (e.g., a command line interface (CLI), a graphical user interface (GUI), etc.) 124 displayed on a display, such as a computer monitor, for displaying data. Through the user interface 124, a user may provide control inputs for manipulating the edge compute network configuration application 112 through one or more input devices. For example, identification of network devices, edge sites, device requests, port and bandwidth requests, and other seeding information and data may be provided via inputs to the user interface 124 and used by the edge compute network configuration application 112 to, among other things, configure components of an edge site to provide edge compute services to customers of the network, and the like. The input device for providing the inputs to the user interface 124 may include, among others and not limited to, a keyboard or a pointing device (e.g., a mouse, trackball, pen, or touch screen) to enter data into or interact with the user interface.

[0035] The edge compute network configuration tool 106 may also include a configuration generator 116 for constructing or generating a configuration plan or series of executable steps for configuring network devices or other computing devices for providing edge compute services from an edge site of a network or collection of networks. In general, the configuration generator 116 receives seed data or information for an edge compute environment via the user interface 124, the network databases 104a-c, and/or the data source 120 and generates a list of

operations or configuration steps for configuring components of the edge site to provide the compute services. One or more of the operations of the configuration generator 116 are described in more detail below with reference to the method 300 of Figure 3A. The edge compute network configuration application 112 may, in some instances, execute one or more of the configuration plan generated by the configuration generator 116 to configure the edge compute system or environment. In one example, execution of the configuration plan may include calling or executing one or more micro-services to configure components of the edge compute system. As such, the edge compute network configuration application 112 may include a micro-services manager 114 to manage the various micro-services executed during configuration of edge devices to provide the compute services. In general, the micro-services manager 114 may call, manage, respond to, and terminate micro-services executed by the edge compute network configuration application 112 during configuration of edge devices of a network or group of networks.

[0036] It should be appreciated that the components described herein are provided only as examples, and that the application 112 may have different components, additional components, or fewer components than those described herein. For example, one or more components as described in Figure 1A may be combined into a single component. As another example, certain components described herein may be encoded on, and executed on other computing systems, such as on one remotely coupled to path computation tool 106.

[0037] Figure 1B is a schematic diagram illustrating a system for configuring an edge compute environment at an edge of a network in accordance with one embodiment. In general, the system 100 may include the edge compute network configuration tool 106 in communication with a source of information such as one or more network databases 104a-c associated with one or more respective networks 102a-c as described above. The edge compute network configuration tool 106 may also be in communication with the computing device 122 providing the user interface 124. The edge compute network configuration tool may also be in communication with a micro-services system 130 and a content delivery network 132. The interactions and communications between the components of the edge compute configuration system 100 is described in more detail herein.

[0038] As explained above, the edge compute network configuration tool 106 may access and configure one or more network devices associated with an edge site of a network or networks to provide compute services from the edge site. In some instances, the edge compute network configuration tool 106 may include an edge compute network configuration application 112

configured to perform one or more of the operations described above and herein. The edge compute network configuration application 112 may include several components to perform one or more of the operations described above and herein. For example, a firmware provisioner 128 may be included in the edge compute network configuration application 112 to configure (or orchestrate configuration of) the edge compute system or environment. In one example, the firmware provisioner 128 will upload or alter firmware or operating system programs to network devices being configured in the one or more communication networks 102. In some instances, the firmware provisioner 128 may utilize the content delivery network 132 to deliver configuration data to the edge compute system or environment. For example, the edge compute network configuration application 112 may use the content delivery network 132 to upload server images, firmware, or operating system programs to network devices being configured in the one or more communication networks 102.

[0039] The content delivery network 132 is a distributed platform of servers and caches that can distribute resources efficiently to devices by using servers geographically close to the devices. An example of a content delivery system is described in U.S. Patent No. 9,456,053, the contents of which are incorporated by reference herein for all purposes. The content delivery network 132 may store content, such as configuration data (e.g., server images, executable files, configuration templates, etc.) on servers and caches throughout the distributed platform. This allows the content delivery network 132 to deliver the content with relatively low latency/cost when communicating with devices geographically close to a particular server of the content delivery network 132. In some instances, the edge compute network configuration tool 106 may cause the content delivery network 132 to deliver configuration data stored within the content delivery network 132 to the edge compute environment or system. In other examples, the edge compute network configuration tool 106 may retrieve and deliver the configuration data (e.g., server images, executable files, configuration templates) needed for devices of the edge compute network to be configured. For example, the content delivery network 132 may be used because it has lower latency/cost compared to the edge compute network configuration tool 106 when communicating with the one or more network devices associated with an edge site or a network or network as will be described in more detail herein. The content delivery network 132 may also be used when the edge compute network configuration tool 106 is strained, such as when the edge compute network configuration tool 106 needs to configure edge compute environments at multiple edges of a network simultaneously. In this instance, the edge compute network configuration tool 106 may cause the content delivery network 132 to deliver at least

part of configuration data to the edge compute environments at one or more of the edges of the network needing configuration. In examples, the edge compute network configuration tool 106 may send instructions to the devices to be configured in network 102 to download the necessary configuration data from the content delivery network 132.

[0040] The edge compute network configuration tool 106 may also include a configuration generator 116 for constructing or generating a configuration plan for providing edge compute services from an edge site of a network or collection of networks as described above. The edge compute network configuration application 112 may, in some instances, execute one or more of the configuration plans generated by the configuration generator 116 to configure the edge compute system or environment. In one example, the configuration generator 116 may cause the content delivery network 132 to deliver the configuration data necessary to one or more of the configuration plans generated by the configuration generator 116 to configure the edge compute system or environment. In some instances, the content delivery network 132 will then deliver the configuration data to the edge compute system or environment, including edge sites of the one or more communication networks 102, directly.

[0041] In another example, execution of the configuration plan may include calling or executing one or more micro-services to configure components of the edge compute system. As such, the edge compute network configuration application 112 may include a micro-services manager 114 to manage the various micro-services executed during configuration of edge devices to provide the compute services. In general, the micro-services manager 114 may call, manage, respond to, and terminate micro-services executed by the edge compute network configuration application 112 during configuration of edge devices of a network or group of networks.

[0042] In some instances, the micro-services and instructions on how to communicate with the micro-services are not stored on the edge compute network configuration tool 106. In these instances, the micro-services manager 114 may call, manage, respond to, and terminate micro-services by communicating with a separate micro-services system 130. In general, the micro-services system 130 may call, manage, respond to, and terminate micro-services executed by the edge compute network configuration application 112 during configuration of edge devices of a network or group of networks when the micro-services manager 114 instructs it to do so.

[0043] In some instances, the micro-services may be stored on micro-services system 130, which is instantiated on a separate computing system from the edge computer network configuration tool 106. In these instances, the micro-services system 130 calls and executes the various micro-services executed during configuration of edge devices to provide the compute

services that the micro-services manager 114 instructs the micro-services system 130 to manage. In examples, the micro-services executed during configuration of edge devices may operate on and employ configuration data downloaded by the edge devices from the content distribution network. In examples, storing micro-services in the micro-services system 130 separately from edge compute network configuration tool 106 is useful in that the micro-services are more readily available to be updated, maintained, and used by other processes that may be unrelated to the operations of the edge compute network configuration tool 106. In examples, some micro-services may have general applicability to other processes of network 100, and the micro-services system 130 may include an API or other interface that allows the micro-services stored thereon to be called and executed by systems and processes that access the micro-services system separately. For example, micro-services system 130 may include an interface that permits micro-services manager 114 to execute micro-services for the configuration of network elements in networks 102 and also allow other processes to access such micro-services for other purposes.

[0044] It should be appreciated that the components described herein are provided only as examples, and that the application 112 may have different components, additional components, or fewer components than those described herein. For example, one or more components as described in Figure 1B may be combined into a single component. As another example, certain components described herein may be encoded on, and executed on other computing systems, such as on one remotely coupled to path computation tool 106.

[0045] Figure 2 is a schematic diagram illustrating an edge compute environment 200 of an edge site of a network in accordance with one embodiment. In general, the edge compute environment 200 of Figure 2 illustrates one example of components of an edge site of a network or collection of networks 102a-102c from which compute services may be provided to customers connected or otherwise in communication with the edge site. As mentioned above, by providing the environment 200 in an edge site of the network 102, compute services may be provided to customers with a smaller latency/cost than if the compute environment is included deeper within the network or further away from the requesting customer of the network. It should be appreciated, however, that an edge compute system may include more or fewer components than those illustrated in Figure 2 and may be connected in other configurations than shown. Rather, the system 200 of Figure 2 is but one example of an edge compute system 200 for providing compute services to devices or networks connected to or otherwise in communication with the edge compute system.

[0046] In the instance shown, the components of the system 200 may be installed or associated with a network site at the edge of one or more networks 102a-c. In general, an edge site of a network is a network site in which customer equipment may connect to the network 102 for access to services and transmission routes of the network. Further and as discussed above, the network 102 may include more than one public and/or private network interconnected to form a general network 102. Each network instance may include one or more network gateway edge devices 202 that provide gateways or ingress/egress devices for the associated network. In Figure 2, network 102a may include edge devices 202a, network 102b may include edge devices 202b, and network 102c may include edge devices 202c. Each edge device 202 of the networks 102 may connect or otherwise communicate with one or more spine switch devices 204a-b. The spine switch devices 204a-b provide an interface between the edge devices of the networks 102a-c to other components, such as host servers, that provide the compute services discussed herein. Each of the spine switch devices 204a-b may be interconnected with one or more host leaf switches 206a-b. One or more host leaf switches 206a-b may interconnect with the one or more spine switch devices 204a-b of the environment 200 to form a switch mesh for connecting to the network 102 via edge devices 202. In some instances, more or fewer spine switch devices 204 and/or host leaf switches 206 may be included in the edge compute environment 200. However, the environment 200 of Figure 2 is illustrated to provide redundancy failover services for each of the components of the environment. In particular, spine switch 2 204b may be a failover, redundant device for spine switch device 1 204a. Each host leaf switch 206a-b may similarly be a failover, redundant device for another host leaf switch.

[0047] The environment 200 may also include one or more application servers 208a-n or other types of servers connected to each host leaf switch 206a-b through a corresponding link aggregation control protocol (LACP) which combines multiple network connections in parallel to increase throughput. In one implementation, the servers 208 may host and execute applications to provide particular services to customers of the network 102. For example, the servers 208 may be configured to provide compute services (as well as other cloud computing services) to customers in communication with the servers. Further, although the environment 200 is illustrated with sixteen such servers 208, the environment 200 may include more or fewer servers for providing services to customers. The environment 200 may, in some instances, also include a host management switch 210 connected to the host leaf switches 206 for managing aspects of the switching mesh and communications to/from the servers 208. Through the

environment 200 of Figure 2, an edge compute service may be provided to customers of the network 102 requesting such services from the network 102 while reducing the latency/cost of providing the services to the customers.

[0048] Figure 3A is a flowchart illustrating a method 300 for configuring an edge compute environment at an edge of a network in accordance with one embodiment. In one implementation, one or more of the operations of the method 300 may be performed by the edge compute network configuration application 112 and, in particular, by the configuration generator 116 of the application. In other implementations, one or more of the operations may be performed by other components of the edge compute network configuration tool 106 or still other systems. The operations may be executed by hardware components of the relevant systems, software programs of the systems, or a combination of hardware and software components of the system.

[0049] Beginning in operation 302, the edge compute network configuration tool 106 may request and/or receive seeding data or information used to generate the edge compute configuration process for configuring the components of the edge compute environment 200. In general, the seeding data may include any information of the network, the edge site, and/or the installed components at the edge site for providing edge compute services to customers connected to the edge site or other location serviced by the network. In one particular instance, the seeding data or information may be provided via the user interface 124 displayed on the computing device 122. For example, Figure 4 is an example user interface 400 for providing seeding data and information for an edge compute environment 200 configuration in accordance with one embodiment. The user interface 400 may be displayed on a display device associated with computing device 122 operated by a customer, network administrator, network engineer, edge site manager or technician, and the like. Through the user interface 124, a user of the computing device 122 may provide the seeding information for generating the edge compute service infrastructure. In still other instances, another computing device may access and utilize the user interface 124 for providing the seeding information to the edge compute network configuration tool 106.

[0050] The user interface 400 of Figure 4 may be a landing page for the interface. To provide the seeding data and information, a selection or other indication of a “Basics” tab 403 in a row of control tabs 402 may be received. Upon selection of the Basics tab 403, one or more active buttons 404-408 for providing seeding information may be displayed. In one particular implementation, the buttons may include an “Add Site” button 404 for adding a new edge site

identifier or other information associated with a new edge site of the network, an “Add PE’s” button 406 for adding a new provider edge device or devices, and an “Add Device” button for adding other new devices at the edge site to the network. More or fewer buttons may be included in the user interface 400 as needed to request or receive seeding information for generating an edge compute system configuration. In one example and upon selection of the Add Site button 404, the user interface 122 may display screenshot 500 as illustrated in Figure 5 and which includes one or more information boxes through which site-specific information may be provided to add to the data source 120 of the edge compute network configuration tool 106 and associated with an edge site of a network.

[0051] The new site screen 500 may, in one instance, include one or more boxes for a user or computing device to enter site identification information. For example, the interface 500 may include a box 502 for entry of a Common Language Location Identifier (CLLI) code associated with a location of the new site. In general, CLLI codes may be used by networks to identify a geographic location associated with the network, such as a geographic location of a network device and/or site. A box 504 for entry of a site name and a box 506 for an identification of a location (such as city and state identifiers, a physical address, geographic coordinate values, etc.) of the site may also be included. Additional site information may also be provided via the user interface 500 used to manage at least portions of the site components. For example, a management subnet box 508 may be included for providing a range of Internet Protocol (IP) addresses used to access and manage components of the edge site. The range of IP addresses may therefore be associated with devices of the identified new site and through which such devices may be accesses (such as by a network administrator) and configured. In some instances, a secondary management subnet box 510 may also be included for identifying a second IP address range used to access devices of the edge site.

[0052] Still other IP address ranges may be specified or provided via the user interface 500 as associated with a new edge site of the network being input through the user interface, such as an IP address range included in a box 512 for identifying a range of IP addresses for a GMSN subnet of the edge site and/or a box 514 for identifying an IP address range for a secondary GMSN subnet. The user interface 500 may also include one or more boxes 516 for providing an indication of the installation of mobility as a service (MAAS) services for the edge site. In some instances, one or more of the boxes of the user interface 500 may be pre-populated with information or data obtained from network databases 104 associated with networks 102 available through the edge site. For example, upon providing the CLLI code for the edge site,

the edge compute configuration tool 106 may access one or more network databases 104 to obtain associated management subnet IP addresses and display the retrieved information in the user interface 500. Such information may or may not be editable by a user of the interface 500 once populated. In another example, the edge compute network configuration tool 106 may obtain site specific information or data from the data source 120 of the tool. In general, any information displayed via the user interface 500 may be pre-populated by the edge compute network configuration tool 106 based on information obtained from one or more databases. Information provided in the boxes 502-516 of the user interface 500 may be received and processed or stored by the edge compute network configuration tool 106 through selection of a "Create site" button 518.

[0053] In another example and upon selection of the Add PE's button 406 illustrated in the user interface 400 of Figure 4, the user interface 124 may display screenshot 600 as illustrated in Figure 6 and which includes one or more information boxes through which information for adding a provider edge or other edge device to the indicated edge site may be input. As mentioned above with reference to Figure 2, an edge site may include one or more edge devices 202a-c that provide an interface for components of the edge site and the broader network 102. Further, various networks 102 accessible through the edge site may include different edge devices 202. New PE screen 600, also referred to as "New Router Group" screen, may therefore include one or more boxes for a user or a computing device to enter identification information for edge devices 202 of the edge site. In one particular instance as shown in Figure 6, the user interface 600 may include a drop-down menu 608 for selecting a previously entered site identifier, such as through new site screen 500 discussed above. Upon selection of a site, one or more boxes for entering edge device information may be presented in the user interface 600. In one instance, a first portion 602 of the user interface 600 may be displayed for entering information for edge devices of a first network 102a, a second portion 604 may be displayed for entering information for edge devices of a second network 102b, and a third portion 606 may be displayed for entering information for edge devices of a third network 102c. More or fewer such portions 602-606 may be included in the user interface 600 corresponding to the networks available at the selected edge site.

[0054] For each available network 602-606, information may be entered associated with edge devices 202 of a corresponding network 102 indicated through the selected site 608. In some instances, each available network 102 may include two edge devices at the edge site for redundancy and load balancing to the associated network. Thus, the first portion 602 may

include entry boxes 610 for a first edge device and entry boxes 612 for a second edge device associated with network A 102a, second portion 604 may include entry boxes 614 for a first edge device and entry boxes 616 for a second edge device associated with network B 102b, and third portion 606 may include entry boxes 618 for a first edge device and entry boxes 620 for a second edge device associated with network C 102c. For each edge device, the user interface 600 may include boxes or drop-down menus for receiving or displaying information associated with the corresponding edge device. Further, the boxes illustrated may vary depending upon a network database from which device information is obtained. For example, section 602 may be associated with network A 102a, which has a first type of edge device (designated as “EAR” devices herein) at the network edge. The edge compute configuration tool 106 may obtain the type of edge device used by network A 102a and, based on identifying the edge device type, populate section 602 with input boxes that correspond to the identified edge device type for that network. In a similar manner, the edge compute configuration tool 106 may identify a different edge device for network 102b (designated as “VAR” devices herein) and, based on obtaining the identification of the edge device type from network database 104b, may populate section 604 with input boxes that may differ than the input boxes of section 602. Thus, the labels and input boxes presented in screenshot 600 may be determined and generated by the edge compute configuration tool 106 based on edge device information obtained from the corresponding network database 104.

[0055] In one example of section 602 of the screenshot 600, a hostname box 622 is provided for entering or displaying a hostname for the corresponding edge device, a drop-down menu 624 for selecting a connection type, and one or more interface boxes 626 may be provided for entering or displaying interface information for the corresponding edge device. Upon entry of information or data for one or more of the edge devices of the edge compute system 200 (including edge devices from the various networks 102 connected to the edge site), a user or computing device may submit the information to the edge compute network configuration tool 106 through selection of a submit button 628. The edge compute network configuration tool 106 may then store the received information in the data source 120 for use in generating an edge compute configuration process as explained in more detail below.

[0056] In addition to providing site information and edge device information, the user interface may be utilized to provide device information to the edge compute configuration tool 106 for inclusion in an edge compute environment 200. In particular, Figure 7 is an example user interface 700 for providing new device seeding data and information for an edge compute

environment configuration in accordance with one embodiment. Similar to the information provided above, the new device information may be received by the edge compute network configuration tool 106 via the user interface 124 and stored in the data source 120 for use in generating an edge compute configuration process. In the example user interface 700 shown, a user or computing device may provide or select a site identifier 702 through a drop down menu or entry box. Other device specific information may also be provided, such as a function 704 of the device to be added (in some instances selectable from a drop-down listing of functions, a serial number 706 associated with the device to be added, a rack identifier 708 identifying a rack at the edge site in which the device may be located, a rack unit identifier 710 at which the device is located in the edge site, a terminal server hostname 712 for the device, and/or a management switch hostname 714 associated with the device. Similar to above, one or more of the interfaces for providing information may be pre-populated by the edge compute network configuration tool 106 based on other provided information, such as a site identifier 702 and based on information obtained from one or more network databases 104 associated with one or more networks 102. Further, upon selection of a submit button 716, the entered information of the new device may be provided to the edge compute network configuration tool 106 for storage in the data source 120.

[0057] Returning now to Figure 3A, the edge compute network configuration tool 106 may, in operation 304, determine one or more configuration templates to use in configuring an edge compute environment based on the information provided above via the user interface 124. In one example, one or more configuration templates corresponding to devices of an edge compute environment 200 may be stored in the data source 120 or database 104 associated with a network 102. The configuration templates provide a generalized process for configuring a network device to perform a defined function. The configuration templates may receive as input information or data associated with a device to be configured, such as that information provided above via the user interface 124. For example, a switch may be configured with an IP address based on switch identification information provided via the user interface 124 or from information stored in a network database 104. The edge compute network configuration tool 106 may identify one or more such configuration templates from a group of stored configuration templates for use in building a configuration process for the edge compute environment 200. For example, the edge compute network configuration tool 106 may use the seed data to identify a manufacturer of a device or a device type and select a configuration template associated with that device type for use in building the configuration process.

[0058] In operation 306, the edge compute network configuration tool 106 may generate the configuration process, which may include one or more executable files, to configure the devices of the edge compute environment 200 based on the obtained templates, seed information, and/or system variables. The configuration process may configure or otherwise “build” one or more components or systems of the environment 200, such as switches 208, spine switches 204, host leafs 206, and the like to process data packets according to the received information and provide compute services to customers connected to the edge devices. For example, the configuration process may upload or alter firmware or an operating system program of a device of the edge compute environment 200. The firmware/operating system program may configure the device according to the template, such as by making communication ports open and available, establishing a network address (such as an IP address) for the device, populating routing tables and/or communication paths through and between devices, and the like. To execute the configuration process, the edge compute configuration tool 106 may perform one or more steps to request or provision port assignments for components of the environment 200 for communications between the devices. In still another example, the configuration process may call a Dynamic Host Configuration Protocol (DHCP) routine that assigns one or more IP addresses or other network configuration parameters to the components of the edge compute environment 200. Also, as explained in more detail below, the configuration process may include instructions to the devices of the environment 200 to report or provide status updates, notices of changes to configurations, check-ins, and other reporting process, including but not limited to information obtained through domain name server records, simple network management protocol, network time protocol, syslog, and/or terminal access controller access control system. In some instances, some steps of the configuration process may wait or be paused until the edge compute network configuration tool 106 receives notification of completion of a prior configuration step from one or more of the edge compute devices. In general, any configuration process or operation used to prepare a device of the edge compute environment 200 may be automated by the configuration process generated by the edge compute network configuration tool 106. An example of element 306 in Fig. 3A is further depicted in Fig. 3B, discussed below.

[0059] In one particular implementation, the configuration process may include calls to one or more micro-services maintained and/or managed by micro-services manager 114 of the edge compute network configuration tool 106. In other examples, the micro-services may be maintained in a separate micro-services system, such as micro-services system 130 in Fig. 1B,

and invoked or otherwise managed by micro-services manager 114 of the edge computer network configuration tool 106. In general, micro-services may include small, containerized, REST-based API services used to interface with various external systems, such as the devices of the edge environment 200 being configured. Thus, in operation 308, the edge compute network configuration tool 106 may call one or more micro-services to execute at least a portion of the generated configuration process. Such micro-services may be utilized to provide zero touch provisioning of the components of the edge compute environment 200, thereby removing the need for a network administrator to accomplish each step of the configuration process. The micro-services may communicate with one or more of the devices of the edge compute environment 200 to instruct or control the devices to perform an operation of the configuration process, including accessing the corresponding device, translating a step of the configuration process into a language used for communicating with the corresponding device, and transmitting the translated step to the device. Through the micro-services, one or more of the steps of the configuration plan may be executed on the devices of the edge compute environment 200.

[0060] As mentioned above, one aspect of the configuration process generated by the edge compute network configuration tool 106 includes instructing a device of the environment 200 to report a status of the configuration step at the device, including a successful configuration or a failed configuration. As such, in operation 310, the edge compute network configuration tool 106 may receive a report from a configured device and, based on the report, determine completion or failure of a configuration step of the configuration process. In some instances, the configuration step may be controlled or executed utilizing a micro-service call. In addition, the edge compute network configuration tool 106 may report the determined success or failure of the configuration step to the user interface 124 for display on the computing device 122. Display of the status of the configuration of devices of the edge compute environment 200 is discussed in more detail below with respect to Figures 8-12.

[0061] In operation 312, the edge compute network configuration tool 106 may determine if additional devices of the edge compute environment 200 is to be configured. For example, the above operations may be executed to configure a first switch of the edge compute environment 200 and receive a response from the first switch that the configuration is successful. The edge compute network configuration tool 106 may then determine, based on the configuration process, that a second switch device is also to be configured. If the edge compute network configuration tool 106 determined more devices of the environment 200 is to be configured,

another micro-service may be called or executed in operation 308 to perform the additional environment 200 configuration. If no additional devices of the environment 200 are to be configured, the edge compute network configuration tool 106 may report the success or failure of the configuration of the edge compute environment 200 to the user interface for display on a computing device 122.

[0062] Through the method 300 of Figure 3A, the edge compute network configuration tool 106 may automatically configure the devices or components of the edge compute environment 200 for use by customers to the edge site. In particular, upon entry of the seed information identifying the devices for configuration, the edge compute network configuration tool 106 may generate and execute, via micro-services in some instances, a configuration process that prepares communications between the devices and allows for the environment 200 to provide compute services to customers connected or otherwise in communication with the edge compute environment. The edge compute network configuration tool 106 may therefore eliminate or reduce instances of mistakes in configuring the devices and improve the time consumed to make edge compute services available from the edge site. For example, the edge compute configuration tool 106 may remove the need for multiple network administrators or technicians to configure the various components separately. Rather, compute services may be provided to customers of the edge site through the configuration process described herein in a fast and efficient manner.

[0063] Figure 3B is a flowchart illustrating, in one example, the method for generating and delivering the configuration data needed to configure devices of the edge compute network. The method of Figure 3B may, in some examples, comprise at least part of element 306 in Fig. 3A. Beginning in operation 322, the edge compute network configuration tool 106 may determine its own latency/cost in communicating with and configuring one or more devices in the edge compute environment 200, such as the particular device to be configured. In examples, latency/cost may be expressed in terms of time, financial cost, bandwidth usage, processing overhead, or otherwise. In operation 324, the edge compute network configuration tool 106 may determine the latency/cost between a content delivery network (e.g., content delivery network 132) and the devices in the edge compute environment 200, such as the particular device to be configured. The determined latencies/costs in communicating with the devices in the edge compute environment 200 may result in either the edge compute network configuration tool 106 or the content delivery network 132 being able to deliver configuration data to the edge compute environment 200 faster or more efficiently than the other.

[0064] In operation 326, it is determined whether to use the content delivery network to deliver configuration data to the devices to be configured. In one example, if the content delivery network has a higher latency/cost when communicating with the edge compute environment 200 than the edge compute network configuration tool 106, flow branches “No,” and the edge compute network configuration tool 106 is used to directly obtain and deliver the configuration data. The method then continues to operation 328, where the edge compute network configuration tool 106 obtains and delivers the configuration data.

[0065] In another example, if the content delivery network has a lower latency/cost when communicating with the edge compute environment 200 than the edge compute network configuration tool 106, flow branches “Yes,” and the method continues to operation 330. In operation 330, the content delivery network is used to deliver the configuration data to the edge compute environment 200. In examples, if a content delivery network 132 is to be used to deliver configuration data to the edge compute environment 200, the edge compute environment 200 (or a particular network element that is being configured within the edge compute environment 200) may be instructed to request the configuration data using a uniform resource locator serviced by the content delivery network 132. Content delivery network 132 will then deliver the configuration data using the delivery techniques described, e.g., at U.S. Patent No. 9,456,053 (or otherwise). In other examples, the edge compute network configuration tool 106 may instruct the content delivery network 132 directly to deliver the configuration data to the edge compute network environment 200.

[0066] In some examples, operation 324 may be optional in that the edge compute network configuration tool 106 may not have access to, or be able to determine, the latency/cost associated with using the content delivery network 132. In those instances, the decision 326 of whether to use the content delivery network 132 may be based on either an assumed latency/cost of using the content delivery network 132, or a threshold latency/cost of using the edge compute network configuration tool 106. For example, if the latency/cost of using the edge compute network configuration tool 106 is above a preset threshold, then the content delivery network 132 may be used regardless of any determination of an actual or estimated latency/cost for the content delivery network 132.

[0067] In examples, the decision 326 whether to use a content delivery network 132 may be based on other factors beyond latency/cost. For example, the edge compute network configuration tool may decide to employ the content delivery network to deliver configuration files to the edge compute environment when the computing system upon which the edge

compute network configuration tool is operating is experiencing a high volume of requests. Further, in examples, the latency/cost may be measured or estimated using any of a variety of methods. For example, the latency/cost may be based on geographic distance between network elements, logical distance between network elements, average latency/cost of recent previous requests/responses, etc.

[0068] The edge compute network configuration tool 106 and the content delivery network 132 may cooperate to configure the edge compute network 200 in the ways described by the examples above. In examples, edge compute network configuration tool 106 and the content delivery network 132 may both have the firmware and/or an operating system program that need to be uploaded or altered. If the content delivery network has a lower latency/cost than the edge compute network configuration tool 106, the content delivery network is instructed to upload the updated firmware or operating system program to the edge compute environment.

[0069] In one implementation, the configuration generation process may require the edge compute network configuration tool 106 and the content delivery network 132 to configure the edge compute environment. For example, the edge compute network configuration tool 106 may not have access to firmware or an operating system program that needs to be uploaded or altered in the edge compute environment 200. In this example, the content delivery network 132 is instructed to upload the firmware or operating system program even though it has a higher latency/cost when communicating with the edge compute environment 200 than the edge compute network configuration tool 106. Similarly, if the content delivery network 132 does not have the firmware or operating system program needed, the edge compute network configuration tool 106 may be used to deliver the configuration data to the network device at the edge compute environment 200 even though it has a higher latency/cost. The edge compute network configuration tool 106 may be used to do any of the other configuration steps required, including orchestrating configuration of the device using configuration data provided by the content delivery network.

[0070] Finally, the method moves to operation 332, where the edge compute network configuration tool 106 instructs the edge compute environment 200 to provide status updates, including whether particular configuration data has been successfully downloaded to the device(s) of the edge compute environment 200. As discussed, the method of Fig. 3B is optional and, in some examples, the content delivery network 132 may always be used to deliver configuration data, and in other examples, the content delivery network 132 may be omitted or never used to deliver configuration data.

[0071] As discussed above, the user interface 124 may display information on the status and/or progress results of aspects of the configuration process for an edge compute environment 200. For example, Figure 8 is an example user interface 800 displaying site data and configuration progress for an edge compute environment configuration in accordance with one embodiment. In one instance, the user interface 800 may be accessed through selection of a “Sites” tab from the control tabs 402 of the interface. Upon selection, the user interface 800 may display site information for sites entered via the New Site interface 500 discussed above. In the example shown, site information for a first site 802 and a second site 804 is illustrated. The site information may include, but is not limited to, a site CLLI code, a site name, a geographical location identifier, management subnet IP addresses for accessing or otherwise communicating with devices of the site, secondary management subnet IP addresses, GMSN subnet IP addresses, and secondary GMSN subnet IP addresses, among others. Such information may be provided to the user interface 124 via the New Site interface 500 discussed above. In one instance, the site information 802-804 may include a status 808 of edge devices located at the edge site. For example, the user interface 800 may include an “unconfigured” status for devices, such as a provider edge device, of the site. Further, the user interface 800 may include one or more action buttons 806 that allow actions to be performed on the site information. For example, the action buttons 806 may allow for editing or deletion of the site information. In another example, selection of a “scan site” action button 806 may cause generation of one or more instructions to scan or communicate with devices at the edge site to ping the devices, request operational information from the devices, and update the site information. Such instructions may cause the known devices of the edge site to generate an operational status report and transmit the report to the edge compute network configuration tool 106, which may process the status reports and update the site information displayed in the user interface 800 accordingly.

[0072] In some instances, the site information 802-804 may include the latency/cost between the edge compute network configuration tool 106 and the edge devices located at the edge sites. The displayed latency/cost may allow a user viewing the user interface 800 to make decisions on how a site is configured or reconfigured. Additionally, it may indicate to the user that another instance of the edge compute network tool 106 should be installed geographically closer to sites having latency/cost higher than the user desires. In some instances, the latency/cost between the sites and the management (MGMT) subnet and the GMSN subnet shown in site information 802-804 is displayed.

[0073] In other instances, the latency/cost value may not be displayed in the user interface 800, but the latency/cost may be automatically evaluated by the edge compute network configuration tool 106 to determine how a site will be configured and whether a new instance of the edge compute network tool 106 should be installed geographically closer to sites having latency/cost higher than desired. Examples of how a site is configured may include, but is not limited to, determining connections between edge devices 202a-c and spine switch devices 204a-b and connections between spine switch devices 204a-b and host leaf switches 206a-b as shown in Figure 2. The latency/cost between edge devices 202, spine switch devices 204, and host leaf switches 206 may allow the edge compute network configuration tool 106 to determine that more or fewer spine switch devices 204 and host leaf switches 206 should be included in the edge compute environment 200. Further examples of how a site is configured may include using content delivery networks instead of the edge compute network configuration tool 106 to configure the edge compute environment as described above.

[0074] In some instances, information corresponding to feedback from the devices during a configuration of the edge compute environment 200 may be included in the user interface 800. For example, aspects of the site information 802-804 may be color coded to indicate verified information (perhaps illustrated in a green color) and unverified information (perhaps illustrated in a red color). As shown in Figure 8, IP addresses of subnets of the site may be verified and displayed in green to illustrate the verified status of the IP addresses, while unverified IP addresses may be displayed in red. Verification of the site information 802-804 may be based on one or more reports of configuration status and/or operational status of devices of the edge site, including devices of the edge compute environment 200. Actions may be taken on the components based on the information in the report to verify the site information 802-804, in some instances.

[0075] A similar interface 900 may be displayed by the user interface for edge devices of the edge site. In particular, Figure 9 is an example user interface 900 displaying edge device information accessed through selection of a "Site PE's" tab from the control tabs 402 of the interface. Upon selection, the user interface 900 may display site information for edge devices at selected sites entered via the New Router Group interface 600 discussed above. In the example shown, edge device information 902 for a selected site is illustrated. The edge device information 902 may include, but is not limited to, an edge device identifier and a connection type for each edge device associated with the selected site. Such information may be provided via the New Router Group interface 600 discussed above. In one instance, the user interface

900 may also include one or more action buttons 904 that allow actions to be performed on the edge device information, such as editing or deletion of the edge device information.

[0076] Selection of the “Devices” tab from the control tabs 402 may cause display of the user interface 1000 of Figure 10, providing information associated with devices of the edge compute environment 200 of one or more edge sites. Similar to above, the user interface 1000 may display device information for devices of the environment 200 entered via the New Device interface 700 discussed above. In the example shown, device information 1002-1004 for multiple devices of an edge site is illustrated. The device information 1002-1004 may include, but is not limited to, an edge site identifier, a hostname of the device, an interface standard for communicating with the device, a device model identifier, a serial number or identifier of the device, and/or a management IP address for accessing the device. In one instance, the user interface 1000 may also include one or more action buttons 1006 that allow actions to be performed on the device information, such as editing or deletion of the device information.

[0077] In some instances, information corresponding to feedback from the devices during a configuration of the edge compute environment 200 may be included in the user interface 1000. For example, aspects of the device information 1002-1004, such as a management IP address associated with the device, may be color coded to indicate verified information (perhaps illustrated in a green color) and unverified information (perhaps illustrated in a red color). Other verification of the device information may be based on one or more reports of configuration status and/or operational status of devices of the edge site, including devices of the edge compute environment 200 and may be displayed as verified or failed 1008.

[0078] A similar interface 1100 may be displayed by the user interface for MAAS installations of the edge site. In particular, Figure 11 is an example user interface 1100 displaying MAAS installations information accessed through selection of a “MAAS Installations” tab from the control tabs 402 of the interface. Upon selection, the user interface 1100 may display information for MAAS installations at an edge site. In the example shown, MAAS installations information 1102 for a selected site is illustrated. The MAAS installations information may include, but is not limited to, an MAAS installation identifier, a DHCP forward IP address, and an endpoint IP address. In one instance, the user interface 1100 may also include one or more action buttons 1104 that allow actions to be performed on the MAAS installations information, such as editing or deletion of the MAAS installations information.

[0079] A user interface 1200 illustrating interface standards utilized by an edge compute environment 200 may also be displayed. For example, selection of the “Interface Standards”

tab from the control tabs 402 may cause display of the user interface 1200 of Figure 12, providing information associated with interface standard of the devices configured in an edge compute environment 200 of one or more edge sites. In the example shown in Figure 12, interface standard information 1202-1204 may be illustrated and associated with devices of the edge compute environment 200 which may include, but is not limited to, a function of a device using the interface standard, a device interface identifier, a remote interface identifier, a remote function for the standard, and the like. In one instance, the user interface 1200 may also include one or more action buttons 1206 that allow actions to be performed on the interface information, such as editing or deletion of the interface information.

[0080] In addition to adding and configuring new devices in an edge compute environment of the network, the edge compute configuration tool may also provide and/or display a user interface for providing data and information for updating a device configuration for an edge compute environment. An example of such a user interface is illustrated in Figure 13. Similar to the information provided above, the update device information may be received by the edge compute configuration tool 106 via the user interface 124 and stored in the data source 120 for use in generating an edge compute configuration process. In the example user interface 1300 shown, a user or computing device may provide or select a site identifier 1302 through a drop down menu or entry box. Other device specific information may also be provided to aid in identifying the device to be updated, such as a device function 1304 (in some instances selectable from a drop-down listing of functions), and a network device standard 1306 corresponding to the device to be updated. A rack location 1310 and/or rack unit identifier 1312 may also be provided. Upon selection of a submit button 716, the entered information of the new device may be provided to the edge compute configuration tool 106 for storage in the data source 120 for updating the information associated with a network device.

[0081] As the user interface 1300 of Figure 13 may be used for an existing device already configured into the edge compute network, one or more of the ports associated with the device may be configured when the device is added to the edge compute environment. However, the user interface 1300 may provide for reconfiguring the ports of the identified device through the reconfigure revenue ports selection button 1308 of the interface. Selection of the reconfigure revenue ports button 1308 may cause the computing device 122 to display user interface 1400 of Figure 14 to receive port configuration information of a deployed edge compute device. In particular, Figure 14 is an example user interface 1400 for providing port configuration data and information for a device of an edge compute environment configuration in accordance with one

embodiment. Similar to the user interfaces above, the edge compute configuration tool 106 may provide and/or display the user interface 1400 for displaying and/or updating a port configuration for a device of an edge compute environment. As such, the user interface 1400 may include a first section 1402 providing information on a current port configuration for an existing edge compute environment. The existing port configuration section 1402 may include an identification 1404 of the edge compute environment and a summary of the port configuration for the site 1406, including a port configuration type, a starting port value, and a number of ports included in the port configuration. In the example illustrated in Figure 14, the edge compute environment includes a port configuration type A that starts at port value 0 and includes 8 ports. The port configuration for the edge compute environment may be edited through the user interface 1400 through an “Edit” button or may be deleted through a “Delete” button. An interface map portion 1410 of the user interface 1400 may also be displayed that provides particular information for the individual ports configured for the edge compute environment. Such information may include an interface identifier, the configuration type, and a location identifier of the port, such as a chassis number, rack identifier, device identifier, and the like. More or less port information may be displayed in the user interface 1400 for use by a user of the interface.

[0082] In addition to the port summary information, additional or supplemental port configurations may be selected through a port selection portion 1408 of the user interface 1400. For example, a user may delete configuration type A from the existing port configuration illustrated in portion 1402 and select to add another port configuration type from port selection portion 1408. In another example, a user may simply select to add an additional port configuration type from the port selection portion 1408 to add to the edge compute environment identified in the user interface 1400. In this manner and through the user interface 1400, the port configuration for the edge compute environment may be configured or edited.

[0083] In some instances, the configuration types listed in the user interface 1400 may be based on the types of devices added to the edge compute environment. For example, the configuration types available for a compute device selected for inclusion in the edge compute environment may be different than configuration types available for a storage device as those devices are configured to perform different operations for the edge compute environment. The selection of the available configuration types made available for inclusion in the edge compute environment may be executed or performed by the edge compute configuration tool 106. For example, the edge compute configuration tool 106 may store a plurality of port configuration

types cross-referenced to edge compute devices, such as compute-based devices, storage-based devices, network-based devices, and the like. Upon addition of a particular type of device for inclusion in the edge compute environment, the edge compute configuration tool 106 may determine which port configuration types correlate to the selected devices for population of the configuration selection portion 1408 of the user interface 1400.

[0084] In still further instances, the edge compute configuration tool 106 may also automatically configure one or more ports of the selected devices. For example, the edge compute configuration tool 106 may determine that a particular configuration type correlates to a selected device for the edge compute environment and apply the determined configuration type to configure the ports of the identified device. Such application of the configuration may also include configuring one or more ports of other devices of the network, such as upstream devices of a network connected to the edge compute environment. The configuration of the ports may also include requesting one or more interface identifiers for the ports and maintaining or storing the port configuration information, such as in the data source 120. In this manner, the edge compute configuration tool 106 may automatically configure one or more ports of the devices of the edge compute environment based on the type of devices selected for inclusion in the edge compute environment. The port configuration flexibility provides for an edge compute environment in which all types of devices, such as compute devices and storage devices, may be located within the same rack of the edge compute environment and which are automatically configured with one or more port configurations.

[0085] The user interface also provides a mechanism through which a port configuration may be copied from one device of the edge compute environment to another device. In particular, Figure 15 is an example user interface 1500 for importing an existing port configuration from one device of an edge compute environment to another device of the edge compute environment in accordance with one embodiment. The user interface 1500 includes a drop-down menu 1502 through which a particular device of an edge compute environment may be selected for copying. Upon selection, the port configuration for the selected device may be displayed 1504 for verification by a user of the user interface. A duplicate settings button 1506 may also be included that, upon activation, causes the port configuration of the selected device 1502 to be imported to a second device, such as device B. The user interface 1500 of Figure 15 may thus be launched in response to a selection of button 1308 of user interface 1300 such that the device to which the port configuration is to be duplicated is known by the edge compute configuration tool 106. Through the user interface 1500, the port configuration of another edge

compute environment device may be imported or duplicated to another edge compute environment device to reduce the port configuration operations performed by a user of the edge compute configuration tool 106.

[0086] Figure 16 is a block diagram illustrating an example of a computing device or computer system 1600 which may be used in implementing the embodiments of the components of the network disclosed above. For example, the computing system 1600 of Figure 16 may be the computing device 122 discussed above. The computer system (system) includes one or more processors 1602-1606. Processors 1602-1606 may include one or more internal levels of cache (not shown) and a bus controller or bus interface unit to direct interaction with the processor bus 1612. Processor bus 1612, also known as the host bus or the front side bus, may be used to couple the processors 1602-1606 with the system interface 1614. System interface 1614 may be connected to the processor bus 1612 to interface other components of the system 1600 with the processor bus 1612. For example, system interface 1614 may include a memory controller 1618 for interfacing a main memory 1616 with the processor bus 1612. The main memory 1616 typically includes one or more memory cards and a control circuit (not shown). System interface 1614 may also include an input/output (I/O) interface 1620 to interface one or more I/O bridges or I/O devices with the processor bus 1612. One or more I/O controllers and/or I/O devices may be connected with the I/O bus 1626, such as I/O controller 1628 and I/O device 1630, as illustrated.

[0087] I/O device 1630 may also include an input device (not shown), such as an alphanumeric input device, including alphanumeric and other keys for communicating information and/or command selections to the processors 1602-1606. Another type of user input device includes cursor control, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to the processors 1602-1606 and for controlling cursor movement on the display device.

[0088] System 1600 may include a dynamic storage device, referred to as main memory 1616, or a random access memory (RAM) or other computer-readable devices coupled to the processor bus 1612 for storing information and instructions to be executed by the processors 1602-1606. Main memory 1616 also may be used for storing temporary variables or other intermediate information during execution of instructions by the processors 1602-1606. System 1600 may include a read only memory (ROM) and/or other static storage device coupled to the processor bus 1612 for storing static information and instructions for the processors 1602-1606.

The system set forth in Figure 16 is but one possible example of a computer system that may employ or be configured in accordance with aspects of the present disclosure.

[0089] According to one embodiment, the above techniques may be performed by computer system 1600 in response to processor 1604 executing one or more sequences of one or more instructions contained in main memory 1616. These instructions may be read into main memory 1616 from another machine-readable medium, such as a storage device. Execution of the sequences of instructions contained in main memory 1616 may cause processors 1602-1606 to perform the process steps described herein. In alternative embodiments, circuitry may be used in place of or in combination with the software instructions. Thus, embodiments of the present disclosure may include both hardware and software components.

[0090] A machine readable medium includes any mechanism for storing or transmitting information in a form (e.g., software, processing application) readable by a machine (e.g., a computer). Such media may take the form of, but is not limited to, non-volatile media and volatile media and may include removable data storage media, non-removable data storage media, and/or external storage devices made available through a wired or wireless network architecture with such computer program products, including one or more database management products, web server products, application server products, and/or other additional software components. Examples of removable data storage media include Compact Disc Read-Only Memory (CD-ROM), Digital Versatile Disc Read-Only Memory (DVD-ROM), magneto-optical disks, flash drives, and the like. Examples of non-removable data storage media include internal magnetic hard disks, SSDs, and the like. The one or more memory devices 1606 may include volatile memory (e.g., dynamic random access memory (DRAM), static random access memory (SRAM), etc.) and/or non-volatile memory (e.g., read-only memory (ROM), flash memory, etc.).

[0091] Computer program products containing mechanisms to effectuate the systems and methods in accordance with the presently described technology may reside in main memory 1616, which may be referred to as machine-readable media. It will be appreciated that machine-readable media may include any tangible non-transitory medium that is capable of storing or encoding instructions to perform any one or more of the operations of the present disclosure for execution by a machine or that is capable of storing or encoding data structures and/or modules utilized by or associated with such instructions. Machine-readable media may include a single medium or multiple media (e.g., a centralized or distributed database, and/or

associated caches and servers) that store the one or more executable instructions or data structures.

[0092] Embodiments of the present disclosure include various steps, which are described in this specification. The steps may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware, software and/or firmware.

[0093] Various modifications and additions can be made to the exemplary embodiments discussed without departing from the scope of the present invention. For example, while the embodiments described above refer to particular features, the scope of this invention also includes embodiments having different combinations of features and embodiments that do not include all of the described features. Accordingly, the scope of the present invention is intended to embrace all such alternatives, modifications, and variations together with all equivalents thereof.

[0094] While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without parting from the spirit and scope of the disclosure. Thus, the following description and drawings are illustrative and are not to be construed as limiting. Numerous specific details are described to provide a thorough understanding of the disclosure. However, in certain instances, well-known or conventional details are not described in order to avoid obscuring the description. References to one or an embodiment in the present disclosure can be references to the same embodiment or any embodiment; and, such references mean at least one of the embodiments.

[0095] Reference to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the disclosure. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, various features are described which may be exhibited by some embodiments and not by others.

[0096] The terms used in this specification generally have their ordinary meanings in the art, within the context of the disclosure, and in the specific context where each term is used. Alternative language and synonyms may be used for any one or more of the terms discussed

herein, and no special significance should be placed upon whether or not a term is elaborated or discussed herein. In some cases, synonyms for certain terms are provided. A recital of one or more synonyms does not exclude the use of other synonyms. The use of examples anywhere in this specification including examples of any terms discussed herein is illustrative only, and is not intended to further limit the scope and meaning of the disclosure or of any example term.

Likewise, the disclosure is not limited to various embodiments given in this specification.

[0097] Without intent to limit the scope of the disclosure, examples of instruments, apparatus, methods and their related results according to the embodiments of the present disclosure are given below. Note that titles or subtitles may be used in the examples for convenience of a reader, which in no way should limit the scope of the disclosure. Unless otherwise defined, technical and scientific terms used herein have the meaning as commonly understood by one of ordinary skill in the art to which this disclosure pertains. In the case of conflict, the present document, including definitions will control.

[0098] Additional features and advantages of the disclosure will be set forth in the description which follows, and in part will be obvious from the description, or can be learned by practice of the herein disclosed principles. The features and advantages of the disclosure can be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the disclosure will become more fully apparent from the following description and appended claims, or can be learned by the practice of the principles set forth herein.

CLAIMS

We claim:

1. A method for managing a communications network, the method comprising:
obtaining, from a plurality of databases, site information associated with an edge site of the communications network, wherein each of the plurality of databases is associated with a different backbone communications network;
generating, by a configuration computing system and based on the site information associated with the edge site of the communications network, a configuration process for instantiating a compute service available from the edge site, the configuration process comprising a set of instructions for configuring an operational status of at least one network edge device of the edge site to provide the compute service to a requesting device in communication with the edge site; and
executing the set of instructions with respect to the configuration data to cause the at least one network edge device to be configured, wherein executing the set of instructions comprises:
causing at least one micro-service to be executed on a micro-services system that is separate from the configuration computing system.
2. The method of claim 1 further comprising:
receiving the site information via a user interface displayed on a computing device, the site information comprising an edge site identifier, an edge device identifier, and an identifier of the network device.
3. The method of claim 1, further comprising:
determining, by the configuration computing system, whether to use a content delivery network to deliver the configuration data to the at least one network edge device;
when it is determined not to use the content delivery network to deliver the configuration data, obtaining, by the configuration computing system, the configuration data and delivering the configuration data to the at least one network edge device; and

when it is determined to use the content delivery network to deliver the configuration data, causing the configuration data to be delivered to the at least one network edge device by the content delivery network.

4. The method of claim 1, wherein determining whether to use the content delivery network to deliver the configuration data comprises determining a first latency/cost between the content delivery network and the at least one network edge device and a second latency/cost between the configuration computing system and the at least one network edge device.

5. The method of claim 1, wherein the micro-service is operable to translate the set of instructions and communicate the translated set of instructions to the at least one network edge device.

6. The method of claim 1 wherein the at least one network edge device is a bare metal server storing an application configured to provide the portion of the compute service to the requesting device.

7. The method of claim 1 wherein generating the configuration process comprises:
obtaining a configuration template comprising a portion of the set of instructions; and
altering the configuration template with the site information.

8. A non-transitory, computer-storage medium storing instructions that, when executed by at least one processor, cause the at least one processor to perform a method for managing a communications network, the method comprising:

obtaining, from a plurality of databases, site information associated with an edge site of the communications network, wherein each of the plurality of databases is associated with a different backbone communications network;

generating, by a configuration computing system and based on site information associated with an edge site of a network, a configuration process for instantiating a compute service available from the edge site, the configuration process comprising a set of instructions for configuring, using configuration data, an operational status of a network device to provide a portion of the compute service to a requesting device; and

executing the set of instructions with respect to the configuration data to cause the network device to be configured, wherein executing the set of instructions comprises:
causing at least one micro-service to be executed on a micro-services system that is separate from the configuration computing system.

9. The non-transitory, computer-storage medium of claim 8 further comprising:
receiving the site information via a user interface displayed on a computing device, the site information comprising an edge site identifier, an edge device identifier, and an identifier of the network device.
10. The non-transitory, computer-storage medium of claim 8, further comprising:
determining, by the configuration computing system, whether to use a content delivery network to deliver the configuration data to the at least one network edge device;
when it is determined not to use the content delivery network to deliver the configuration data, obtaining, by the configuration computing system, the configuration data and delivering the configuration data to the at least one network edge device; and
when it is determined to use the content delivery network to deliver the configuration data, causing the configuration data to be delivered to the at least one network edge device by the content delivery network.
11. The non-transitory, computer-storage medium of claim 8, wherein determining whether to use the content delivery network to deliver the configuration data comprises determining a first latency/cost between the content delivery network and the at least one network edge device and a second latency/cost between the configuration computing system and the at least one network edge device.
12. The non-transitory, computer-storage medium of claim 8, wherein the micro-service is operable to translate the set of instructions and communicate the translated set of instructions to the at least one network edge device.
13. The non-transitory, computer-storage medium of claim 8, wherein the at least one network edge device is a bare metal server storing an application configured to provide the portion of the compute service to the requesting device.

14. The non-transitory, computer-storage medium of claim 8 wherein generating the configuration process comprises:
- obtaining a configuration template comprising a portion of the set of instructions; and
 - altering the configuration template with the site information.
15. A system for managing a communications network, the system comprising:
- at least one processor;
 - memory, operatively connected to the at least one processor and storing instructions that, when executed by the at least one processor, cause the system to perform a method, the method comprising:
 - obtaining, from a plurality of databases, site information associated with an edge site of the communications network, wherein each of the plurality of databases is associated with a different backbone communications network;
 - generating, by a configuration computing system and based on site information associated with an edge site of a network, a configuration process for instantiating a compute service available from the edge site, the configuration process comprising a set of instructions for configuring, using configuration data, an operational status of a network device to provide a portion of the compute service to a requesting device; and
 - executing the set of instructions with respect to the configuration data to cause the network device to be configured, wherein executing the set of instructions comprises:
 - causing at least one micro-service to be executed on a micro-services system that is separate from the configuration computing system.
16. The system of claim 15, wherein the method further comprises:
- receiving the site information via a user interface displayed on a computing device, the site information comprising an edge site identifier, an edge device identifier, and an identifier of the network device.
17. The system of claim 15, wherein the method further comprises:
- determining, by the configuration computing system, whether to use a content delivery network to deliver the configuration data to the at least one network edge device;

when it is determined not to use the content delivery network to deliver the configuration data, obtaining, by the configuration computing system, the configuration data and delivering the configuration data to the at least one network edge device; and

when it is determined to use the content delivery network to deliver the configuration data, causing the configuration data to be delivered to the at least one network edge device by the content delivery network.

18. The system of claim 15, wherein determining whether to use the content delivery network to deliver the configuration data comprises determining a first latency/cost between the content delivery network and the at least one network edge device and a second latency/cost between the configuration computing system and the at least one network edge device.

19. The system of claim 15, wherein the micro-service is operable to translate the set of instructions and communicate the translated set of instructions to the at least one network edge device.

20. The system of claim 15, wherein the at least one network edge device is a bare metal server storing an application configured to provide the portion of the compute service to the requesting device.