

TECHNIQUES FOR NETWORK REPLICATION

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application incorporates by reference for all purposes the full disclosure of
5 U.S. Patent Application No. 12/060,074, filed on March 31, 2008, having the title
“Configuring Communications Between Computing Nodes.”

BACKGROUND

[0002] Remote computing services are often utilized in order to operate a network on
behalf of various users. An organization, for example, may remotely operate a plurality of
10 computing devices that are in communication with one another. The computing devices may
be configured in a network according to the organization’s needs, utilizing various network
devices such as routers, switches, firewalls, and the like. Further, the computing and network
devices may include virtual and/or actual computing devices. For example, an organization’s
network may include a plurality of virtual machines in communication with one another by
15 one or more virtual network devices, thereby forming a virtual network that overlays an
actual network. The computing devices may be programmed to communicate according to
the topology of the virtual network and various techniques may be used to enable the
communications to travel over the actual network which the virtual network overlays.

[0003] In some instances, it is desirable to duplicate a network, or at least a portion thereof.
20 For example, an organization may wish to test a network while still allowing the network to
function. Similarly, an organization may wish to simulate various stresses on its network
without actually subjecting the network to the stresses. The organization may wish to test
how a modification to its network may function or may wish to recreate events that caused
one or more events, such as a network failure. Generally, there are various reasons why
25 duplication of a network can be useful. Network duplication, however, can be an arduous
and time-consuming process, generally requiring substantial intervention configuring a new
network to be a duplicate of an old network.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] Figure 1 illustrates an environment in which various embodiments can be implemented.

[0005] Figure 2 illustrates an environment which may utilize various components of the environment shown in Figure 1.

[0006] Figure 3 illustrates a portion of the environment of Figure 2 in greater detail.

[0007] Figure 4 is a flowchart for an illustrative example of a process for requesting a network duplicate, in accordance with an embodiment.

[0008] Figure 5 is a flowchart for an illustrative example of a process for duplicating a network in response to receiving an instruction for duplicating a network, in accordance with an embodiment.

[0009] Figure 6 is a flowchart for an illustrative example of a process for duplicating a network, in accordance with an embodiment.

[0010] Figure 7 is a flowchart for an illustrative example of a process for creating a copy of an object for a duplicate network, in accordance with an embodiment.

[0011] Figure 8 is a diagram of a timeline that illustrates creation of a shallow copy of a network object, in accordance with an embodiment.

[0012] Figure 9 is a diagram of a timeline that illustrates creation of a deep copy of a network object, in accordance with an embodiment.

[0013] Figure 10 is a flowchart for an illustrative example of utilizing a duplicate network, in accordance with an embodiment.

[0014] Figure 11 is a flowchart for another illustrative example of utilizing a duplicate network, in accordance with an embodiment.

DETAILED DESCRIPTION

[0015] Techniques described and suggested herein include systems and methods useful for processes such as network duplication. In an embodiment, a request to duplicate at least a portion of network is received and at least that portion of the network is duplicated, responsive to the request. The request may be initiated by a user, such as a network administrator, who accesses the network remotely. Access of the network may be through various Application Programming Interfaces (APIs) of a web service layer provided for the purpose of allowing users to access their network(s). As an example, the network may be

operated on behalf of the user (and/or an organization associated with the user) by another organization that provides computing resources used for operating the network. The computing resources may be, as an example, located in one or more data centers and remotely accessed by the user using APIs of web services, although other technologies may be used. One or more of the APIs may be provided for the purpose of network duplication. The one or more data centers may additionally be used to implement networks of other users that may access their networks remotely. The computing resources may include computing devices, storage devices, switches, routers, gateways, hubs, and other devices that may be used in operating a network, including some devices described below and, generally, any device suitable for use in a network environment.

[0016] In an embodiment, the network that is duplicated is a virtual overlay network. In an embodiment, the network that is duplicated comprises a network of one or more virtual devices that overlay a network of physical devices. The physical devices may include one or more hub, switch, router, gateway, firewall, server, storage device, and/or one or more other devices that can be used in a network environment. Similarly, the virtual devices may include one or more virtual hub, virtual switch, virtual router, virtual gateway, virtual firewall, virtual server, virtual storage device, and/or one or more other devices that can be used in a network environment. As used herein, when clear from context, virtual devices may be referred to without the word “virtual.” There may not be a one-to-one correspondence between a set of devices comprising the virtual network and the set of devices comprising the physical network that is overlaid by the virtual network. As an example, one physical device may be used to implement multiple virtual devices. A computing device, for example, may be used to implement multiple virtual machines. As another example, multiple physical devices may be used to implement a single virtual device. A virtual storage device, for instance, may be implemented by multiple physical storage devices. In addition, while multiple virtual devices may be implemented by a single physical device and while multiple physical devices may be used to implement a single virtual device, some virtual devices may be implemented by a single physical device. A virtual storage volume, for instance, may be implemented by single physical storage device.

[0017] In an embodiment, a user who makes a request is provided the opportunity to specify one or more parameters regarding how a network may be duplicated. In one embodiment, the user is able to identify one or more portions of a network, such as one or more subnetworks, to copy. In addition, devices in a user’s network may be tagged with various characteristics. A tag for a device may relate to the device’s purpose in the network

(such as the type of application the device is used to support), may generally relate to any aspect of the device, or may have no apparent relation to a characteristic of the device. A user, in an embodiment, may specify one or more tags as a condition for inclusion or exclusion from the duplicate network. For example, a user may specify that only devices
5 tagged with “production” should be included in a duplicate network.

[0018] Other parameters may relate to the fidelity of objects in a duplicate network to the network to be copied. For example, if a network to be copied includes a virtual machine, the user may have various options for specifying requirements for a duplicate virtual machine. The user may, for instance, specify low fidelity for the machine whereby a duplicate network
10 may simply have a corresponding virtual machine. The user may specify higher fidelities where the capabilities of the corresponding virtual machine match the original virtual machine to various degrees. The user may specify, for instance, that the duplicate virtual machine should be implemented with minimum processing capabilities, block storage available, and/or abilities to run various software applications. Similarly, the user may
15 specify very high fidelity where, for instance, a duplicate virtual machine should be implemented on a machine using a particular processor model. Generally, any type of characteristic relating to how closely a duplicate network object should match an object being duplicated may be provided to the user for specifying copy fidelity. As another example of parameters that may be specified, a user may specify various Internet protocol addresses of
20 devices in a network that should be used for corresponding devices in a duplicate network.

[0019] In an embodiment, users are also able to specify at least one copy type for devices in a network to be copied. In one embodiment, users are able to specify copy types such as “shallow” or “deep” copies in a network to be copied, although other types of copies may be provided for selection. With a shallow copy of a network object, a new instance of the object
25 may be declared. The new instance may operate, retrieving information regarding the original object’s state before the copy, as needed. A shallow copy of a network object may be, for example, a copy-on-write (COW) overlay of the network object. With a deep copy of a network object, a complete copy of a store of memory comprising information about the object’s state may be made for a duplicate copy to enable the duplicate copy to operate
30 independently without the need to retrieve information regarding the original object’s state before the copy from another source.

[0020] As discussed, when the user submits a request that a network be copied, in response to the requests, the network is copied and access to the network is provided to the user, in an embodiment. Copying the network may include creating one or more virtual duplicate

devices that correspond to one or more virtual devices of the network being copied.

Virtualization or other technology used to create and/or operate the virtual devices of the network being copied may also be used to create and operate corresponding duplicate

devices. In addition, any templates used to create devices in the network being copied may

5 be used to create duplicate devices. In other words, devices and their duplicate devices may be created from common templates. Copying the network is done in a manner consistent

with any parameters specified by the user. The specified parameters may apply to the whole network being copied and/or to individual network objects being copied. In some

embodiments, users are able to specify different parameters for each network object to be

10 copied and/or subsets of the network objects to be copied.

[0021] Providing access to the duplicate network may be done in a variety of ways.

Generally, a user, such as a network administrator, is able to operate the duplicate network in various ways and for various purposes. As one example, the user may apply various stresses

on the duplicate network for testing purposes. The user may also record network events and

15 replay those events in the duplicate network for forensic purposes. As another example, the user may use the duplicate network in connection with the network being copied to provide

additional capacity or, generally, to have available the additional computing resources of the duplicate network. A user may also modify the duplicate network, such as to test how

modifying the network that was copied would affect its performance. Details regarding the

20 various embodiments, including additional examples, are described in more detail in the following paragraphs.

[0022] Figure 1 illustrates an example of an environment 100 for implementing aspects in accordance with various embodiments. As will be appreciated, although a Web-based

environment is used for purposes of explanation, different environments may be used, as

25 appropriate, to implement various embodiments. The environment includes an electronic client device 102, which can include any appropriate device operable to send and receive requests, messages, or information over an appropriate network 104 and convey information

back to a user of the device. Examples of such client devices include personal computers, cell phones, handheld messaging devices, laptop computers, set-top boxes, personal data

30 assistants, electronic book readers, and the like. The network can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network, or any other such network or combination thereof. Components used for such a system can depend

at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed

herein in detail. Communication over the network can be enabled by wired or wireless connections, and combinations thereof. In this example, the network includes the Internet, as the environment includes a Web server 106 for receiving requests and serving content in response thereto, although for other networks an alternative device serving a similar purpose could be used as would be apparent to one of ordinary skill in the art.

[0023] The illustrative environment includes at least one application server 108 and a data store 110. It should be understood that there can be several application servers, layers, or other elements, processes, or components, which may be chained or otherwise configured, which can interact to perform tasks such as obtaining data from an appropriate data store. As used herein the term "data store" refers to any device or combination of devices capable of storing, accessing, and retrieving data, which may include any combination and number of data servers, databases, data storage devices, and data storage media, in any standard, distributed, or clustered environment. The application server can include any appropriate hardware and software for integrating with the data store as needed to execute aspects of one or more applications for the client device, handling a majority of the data access and business logic for an application. The application server provides access control services in cooperation with the data store, and is able to generate content such as text, graphics, audio, and/or video to be transferred to the user, which may be served to the user by the Web server in the form of HTML, XML, or another appropriate structured language in this example. The handling of all requests and responses, as well as the delivery of content between the client device 102 and the application server 108, can be handled by the Web server. It should be understood that the Web and application servers are not required and are merely example components, as structured code discussed herein can be executed on any appropriate device or host machine as discussed elsewhere herein.

[0024] The data store 110 can include several separate data tables, databases, or other data storage mechanisms and media for storing data relating to a particular aspect. For example, the data store illustrated includes mechanisms for storing production data 112 and user information 116, which can be used to serve content for the production side. The data store also is shown to include a mechanism for storing log data 114, which can be used for reporting, analysis, or other such purposes. It should be understood that there can be many other aspects that may need to be stored in the data store, such as for page image information and to access right information, which can be stored in any of the above listed mechanisms as appropriate or in additional mechanisms in the data store 110. The data store 110 is operable, through logic associated therewith, to receive instructions from the application server 108 and

obtain, update, or otherwise process data in response thereto. In one example, a user might submit a search request for a certain type of item. In this case, the data store might access the user information to verify the identity of the user, and can access the catalog detail information to obtain information about items of that type. The information then can be returned to the user, such as in a results listing on a Web page that the user is able to view via a browser on the user device 102. Information for a particular item of interest can be viewed in a dedicated page or window of the browser.

[0025] Each server typically will include an operating system that provides executable program instructions for the general administration and operation of that server, and typically will include a computer-readable medium storing instructions that, when executed by a processor of the server, allow the server to perform its intended functions. Suitable implementations for the operating system and general functionality of the servers are known or commercially available, and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein.

[0026] The environment in one embodiment is a distributed computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components than are illustrated in Figure 1. Thus, the depiction of the system 100 in Figure 1 should be taken as being illustrative in nature, and not limiting to the scope of the disclosure.

[0027] Figure 2 provides an illustrative example of an environment 200 that may make use of one or more components shown in Figure 1 and which may be used to practice various embodiments of the present disclosure. In the environment 200 pictured in Figure 2, a plurality of customer internal networks 202 are communicatively connected with a remote computing service provider network 204. While Figure 2 illustrates a plurality of customer internal networks 202, embodiments of the present disclosure may be implemented in environments in which there are one or zero customer internal networks 202. Returning to the illustrated embodiment in Figure 2, the customer internal networks 202 may be, for instance, intranets of organizations that utilize remote computing services of the remote computing service provider. As an example, a customer may utilize its own internal network for executing various applications used during the customer's operations while using computing resources provided as a service using the remote computing service provider network 204 for other applications.

[0028] In an embodiment, the remote computing service provider network 204 is implemented in a data center operated by a remote computing service provider. A data center may be a facility in which various computing resources are utilized. Computing resources include computing devices, storage devices, network devices such as routers, switches, firewalls, and the like and, generally, any device that may be used in connection with a network of computing devices. The remote computing service provider network 204 may be implemented with a plurality of data centers that communicate with one another over a communications network. The remote computing service provider may, for instance, operate a network of data centers in various geographic locations. As shown in the drawing, the remote computing service provider network 204 includes one or more customer remote networks 206 which, in this example, are networks operated on behalf of corresponding customers. Various operations of a customer may be operated using a remote customer network 206. For instance, a customer may utilize a customer internal network 202 for internal operations (such as accounting, human resources, electronic mail, and the like) while utilizing a customer remote network 206 for external operations (such as operating a network of servers and data stores to provide a website or other information resource to the public or other users external to the customer's organization). Generally, a customer may utilize a customer internal network 202 and/or customer remote network 206 for any suitable purpose. One or more of the customer remote networks 206 may be virtual overlay networks, such as those described in accordance with U.S. Patent Application No. 12/060,074, filed on March 31, 2008, having the title "Configuring Communications Between Computing Nodes," and which is incorporated herein by reference for all purposes.

[0029] In an embodiment, the remote computing service provider network 204 includes a web services layer 208 having a plurality of application programming interfaces (APIs) 210. The web services layer 208 may be implemented in a computer system that is operable to receive web service calls and respond accordingly, such as by instructing one or more other computer systems to perform an action based on the request. The web service calls may be made using various web service APIs 210 provided to users for various reasons. For instance, one API may be for the purpose of network duplication, as described in accordance with the various embodiments of the present disclosure. Other APIs may be for other purposes, such as provisioning virtual computing instances, otherwise configuring a network, and, generally, remotely operating computing resources on behalf of a user utilizing the APIs. In addition, while web service APIs are provided for the purpose of illustration, any way of

operating computing resources, either remotely or locally, may be used in accordance with various embodiments.

[0030] Figure 3 shows an illustrative example of an environment 300 which may be a sub-environment of the environment 200 illustrated above in connection with Figure 2. As illustrated in Figure 3, the environment 300 shows a portion of the environment 200 as it applies to a particular customer. In this example, a customer has its own customer internal network 302 which accesses a remote computing service provider network 304 in which a customer remote network 306 is operated by a remote computing service provider 304 on behalf of the customer. While the customer is shown to have several subnetworks 308 in Figure 3, a customer may have only one network or may have more subnetworks than shown in the drawing.

[0031] The customer remote network 306, in an embodiment, includes a plurality of customer subnetworks 308. In an embodiment, one or more of the customer subnetworks are overlay networks of a network operated by the remote computing services provider. The customer, in the illustrative embodiment shown, may access the various subnetworks 308 through a router 310 and virtual private network (VPN) gateway 312 over a secure VPN connection over the Internet 314, although other ways of accessing the customer subnetworks 308 may be used. In an embodiment, the customer, from its customer internal network 302, establishes a secure VPN connection over the Internet and interacts with the customer subnetworks 308 through web service calls, although other ways of communicating with the various computing devices operating on behalf of the customer may be used. In an embodiment, a server of the customer internal network 302 communicates encrypted information to the VPN gateway 312 which decrypts the information and passes the decrypted instructions to the router 310. The router 310 then transmits the decrypted information to an appropriate customer subnetwork 308. Information from any of the customer subnetworks 308, whether responsive to any instructions sent by the customer internal network 302 or not, may be sent to the router 310, which passes the information to the VPN gateway 312, which encrypts the information and sends the information over the secure VPN connection 314 to the customer internal network 302.

[0032] While Figure 3 shows a particular configuration for information flow for the purpose of illustration in accordance with various embodiments, other configurations may be utilized for the environment 300 or other embodiments. Generally, any suitable way of exchanging information among computing devices may be used. In addition, while several devices such as the VPN gateway 312 and router 310 are included in Figure 3 for the purpose

of illustration, other devices may be used for communication. In addition, the environment 300 shown in Figure 3 is simplified and other devices other than those illustrated may be included in the environment 300 or in other embodiments. Possible devices include, but are not limited to, switches, servers, firewalls, data stores, hubs, other gateways, and the like.

5 [0033] As noted above, various embodiments of the present disclosure provide for dynamic network replication. Accordingly, Figure 4 shows an illustrative embodiment of a process 400 which may be used by a customer to copy one or more networks. Portions of the process 400 (or any other processes described herein, or variations and/or combinations thereof) may be performed under the control of one or more computer systems configured with executable
10 instructions and may be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. The code may be stored on a computer-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors.

15 [0034] In an embodiment, a user selects 402 a network to be copied. Selection of a network to be copied may be accomplished in any suitable manner. For instance, in an embodiment that includes a remote computing services provider, a customer's network administrator or other user may utilize an interface, sometimes referred to as a management console, that allows the user to select a network that is being operated on behalf of the
20 customer by a remote computing services provider. The interface may be provided to the customer by the remote computing services provider as a web page or may be an application executing on a computing device of the customer. The network that is selected may be a subnetwork of a larger network, a collection of subnetworks, or, generally, any collection of computing resources in communication with one another.

25 [0035] In an embodiment, users are provided the ability to apply one or more filters in order to have more control over the network to be copied. In an embodiment, various devices in a customer's network have one or more assigned tags. Tags may be automatically assigned or may be user assigned. For instance, devices may be tagged with device type tags, which may identify devices as being, for instance, a server, a storage device, a switch, or any
30 other type of device. Users may also provide their own tags, in an embodiment. For example, if a customer uses a subnetwork for a different purpose than another subnetwork, the user may tag each device with a tag that corresponds to the purpose. Devices in one subnetwork may have, for instance, tags that identify the devices as part of a subnetwork used for production while other devices in another subnetwork may have tags that identify the devices

as part of a subnetwork used for another purpose, such as billing. As another example, storage devices may be tagged with tags that identify the devices as either being primary storage devices or backup storage devices. Generally, any type of tag may be used for filtering. In addition, devices may have more than one tag and devices in separate

5 subnetworks may share one or more of the same tag. Filters may also be based at least in part on the internet protocol (IP) addresses utilized by devices in the selected network.

Accordingly, in an embodiment, when the user has selected a network to be copied, the user selects 404 any filters that he or she would like to apply to the selected network.

[0036] In addition, in an embodiment, users are provided the ability to select 406 a desired resolution for one or more network objects in the selected network. A resolution selected by
10 the user may be used to determine how closely the copy of the network should match the network being copied. Characteristics of a resolution selectable by the user may include characteristics such as whether a copy is shallow or deep (as described more fully in connection with Figures 8 and 9), software installed on network objects, memory (volatile
15 and/or non-volatile) of a network object, processing capabilities, and particular hardware characteristics, such as processor models, manufacturers, and the like. As an example, a user interested in a low resolution copy may specify that he or she would like a copy of similar devices that are networked together in the same manner. A user interested in a high-resolution copy may specify that certain devices should have a certain amount of memory,
20 should include specific hardware devices, should have specific software installed on them and the like.

[0037] In an embodiment, when the user selects the copy resolution, he or she may send 408 an instruction to copy the selected network according to the resolution specified by the user and according to any filters specified by the user. In an embodiment where the network
25 to be copied is operated by a remote computing services provider, sending the instruction may include invoking a web service provided by the remote computing services provider in order to send the instruction over a communications network. Other methods of sending the instruction may be used in accordance with various embodiments.

[0038] While Figure 4 illustrates various steps performed in a particular order, variations
30 are considered as being within the scope of the present disclosure. For instance, in an embodiment, network selection, filter selection, and copy resolution selection may be performed in any order and one or more steps may be performed simultaneously. Other variations are considered as being within the scope of the present disclosure.

[0039] Figure 5 provides an illustrative example of a process 500 for copying a network. The process 500 may be performed responsive to receipt of an instruction to copy the network generated in accordance with the embodiment illustrated in connection with Figure 4. In an embodiment, a copy instruction is received 502. The instruction may identify a network to be copied, such as by an identifier of the network, a list of computing devices to be copied, or in any suitable manner. Responsive to the copy instruction, in an embodiment, a mapping of the network is accessed 504, in accordance with an embodiment. In an embodiment, a mapping of the network is one or more files or other collections of information that indicates how computing devices of the network to be copied are connected with one another.

[0040] When the mapping is accessed, in an embodiment, a new mapping is generated for a duplicate network. The new mapping may comprise a collection, such as a list, of network objects and their relationships with one another. Generating the new mapping may be accomplished by applying any user-selected filters and/or other criteria to the original mapping. Once the new mapping is generated, network objects are created according to the new mapping. In an embodiment, the new mapping may not be generated, such as when no filters or other criteria apply, and the original mapping may be used to create the duplicate network. In an embodiment, when the new mapping is generated, network objects are created 508 according to the new mapping, as described more fully below. In general, network objects that correspond to the network being copied are created according to the relationships among the network objects in the network.

[0041] Figure 6 is a flowchart for an illustrative example of a process 600 for creating and providing access to a duplicate network, in accordance with an embodiment. Portions of the process 600 may correspond to portions of the process 500 described above in connection with Figure 5. In an embodiment, a copy instruction is received 602, such as in a manner described above. A network configuration data store is examined 604 and a network object tree is fetched 606, in accordance with an embodiment. A network configuration data store may be a data store in which network objects of the network to be copied are persisted by a remote computing services provider. A network object tree, in an embodiment, is a graph comprising objects of a network and the relationships among the objects. The tree may be rooted at any suitable node.

[0042] Once the network object tree is fetched 606, in an embodiment, a determination is made 608 whether any filters apply to the network object tree. If it is determined that a filter applies to the network, the filter is applied 608 to the objects in the network object tree to

generate a network object list. For example, if a user had specified that the duplicate network should only have network objects with one or more particular tags, objects having the one or more tags are identified and added to a list (or other collection) of network objects.

Similarly, if the user had specified that only certain IP addresses are to be included in the duplicate network, only such network objects are added to a list. Generally, any criteria that the user specifies is used to determine what objects are added to the list of network objects, in accordance with an embodiment.

[0043] Once any applicable filters have been applied or if no filters apply, a directed acyclic graph (DAG) of network objects is generated 612, in accordance with an embodiment. The DAG of network objects, in an embodiment, is a graph whose nodes are network objects to be created for the duplicate network and whose edges have a direction according to dependency. The DAG may or may not be a connected graph. As an example, in an embodiment, virtual computing instances cannot be created without block storage devices being created for the virtual computing instances. Thus, in an embodiment, an edge in the DAG from a block storage device to a virtual computing instance that utilizes the block storage device for storage would have a direction from a node representing the block storage device to the virtual computing instance. While the present illustrative example uses DAGs for the purpose of illustration, any plan for duplication of a network may be used in accordance with various embodiments.

[0044] Once a DAG has been constructed, in an embodiment, a copy operation for a first network object in the DAG is invoked 614. The first network object may be any suitable network object represented by a node in the DAG, such as any node to which no edges are directed. Invoking the copy operation, in an embodiment, includes taking an action, the result of which will result in a copy of the network object being created. Invoking the copy operation, therefore, may include allocating sufficient computing resources for the network object copy and configuring the resources accordingly. For instance, if the network object is a virtual data storage volume, one or more physical storage devices may be allocated to the volume and the physical storage devices may be configured to serve the volume. If multiple physical storage devices serve the same virtual volume, a server that serves the multiple physical storage devices may likewise be allocated and configured. Generally, any computing resources may be allocated and configured in a suitable manner.

[0045] Once the copy operation has been invoked for the first network object in the DAG, a determination is made 616 whether other objects remain to be copied. If objects remain to be copied, the DAG is traversed to another node and a copy operation is invoked 618 for the

next network object in the DAG. This process of traversing the DAG and determining whether other objects remain to be copied may continue until the DAG has been completely traversed. When the DAG is completely traversed, in an embodiment, access is provided 620 to a duplicate network consisting of the copied objects, in accordance with an embodiment.

- 5 Providing access to the duplicate network may include providing one or more users of the customer the ability to direct operation of the network, configure the network, or otherwise access the network using a web service or otherwise.

[0046] Figure 7 shows an illustrative example of a process 700 for copying a network object, in accordance with an embodiment. The process 700 may be performed, for example, upon invocation of a copy operation as described in connection with Figure 6. In an 10 embodiment, when an object is to be copied, an object type is identified 702, in accordance with an embodiment. An object type may be one of various object types that can be used in a network environment. Examples include machines, storage devices, routers, switches, firewalls, and the like. In an embodiment, once the object type is identified, a determination 15 is made 704 whether the object is stateless. An object may be stateless if the object whose configuration does not change as the object operates in a network. A server that treats each request it receives as independent and unrelated to any previous request may be a stateless object. Similarly, a switch whose configuration does not change may also be stateless. In an embodiment, objects that are not stateless are stateful, that is their configurations change 20 during operation. A server whose operations depend on previous actions taken in the network, for example, is not stateless. Similarly, a storage volume whose configuration changes as data is written and/or deleted is not stateless.

[0047] If the object is stateless, then a duplicate object may be created 706 immediately. For example, resources may be provisioned and configured for the object so as to emulate the 25 object. If, however, the object is not stateless, then more complex processes for copying the object may be necessary. For example, a determination may be made 708 regarding a copy type. In an embodiment, a copy may be one of two types, shallow or deep, although other types of copies may be made. A shallow copy of an object, in an embodiment, is an object that is at least an initially incomplete copy of an original object. A lazy copy is a type of 30 shallow copy. With a shallow copy, information regarding the object copy's state may be available from the original object until that information is to change due to operation of the original object, at which point the information may be copied to the object copy before it is changed. An example is a copy-on-write (COW) overlay of a storage volume. When the COW overlay is initiated, a block of data for the COW overlay may be stored in the original

storage volume until a write operation is requested for the block of data. When the write operation is requested for the block of data, the block is copied to the COW overlay and then performed for the original storage volume. In this manner, a copy may be made that is quickly and initially available for use. Similar types of copies may be made for any stateful object, such as virtual machines and, generally, any object whose state changes during operation.

[0048] A deep copy, on the other hand, may be a complete copy of an object such that, once the copy is made, it is not necessary to refer to the original object during the copy's operation. As an example, a deep copy of a storage volume may be made by copying exactly all of the data in the storage volume into a new volume such that any data that was stored in the original volume before the copy was made is accessible via the copy. Deep copies of virtual machines and other stateful objects may also be made. In this manner, the copy may become completely independent from the original object.

[0049] If, in an embodiment, a determination is made that the copy type is to be a shallow copy, the object is paused 710. Once the object is paused, an object overlay is created as a duplicate object 712 and the original object's operation is resumed 714. If, however, a determination is made that the copy is to be a deep copy, a deep copy of the object is made accordingly. For example, the object may be paused 716 and the memory state of the object may be copied 718. The copied memory state, in an embodiment, is used to create 720 a duplicate object and the operation of the original object is resumed 722.

[0050] Figures 8 and 9 provide respective illustrative examples of how shallow and deep copies of network objects may be made in accordance with the various embodiments. Beginning with Figure 8, the figure illustrates a timeline and an object's state with respect to the timeline. A network object 802 is operated during the normal course of operations and a request to duplicate a network of which the object is part is received. In an embodiment, network duplication proceeds, such as in a manner described above. At a point in time during network duplication, the network object 802 is paused. An overlay of the paused network object 804 is declared and the paused network object 804 is resumed. As illustrated in the drawing, when the network object is resumed, both the network object and the network object overlay proceed from a common state, the state of the network object at pause. When operations of the network object overlay require information regarding the state of the network object prior to network duplication, such information may be obtained from storage of the state of the network object at pause. Similarly, data may be written to the network object overlay as necessary, such as in the case of the network object overlay being a COW

overlay, or at another convenient time, such as during times when the network traffic is relatively light.

[0051] Proceeding to Figure 9, the figure similarly illustrates a timeline and an object's state with respect to the timeline. A network object 902 is operated during the normal course of operations and a request to duplicate a network of which the object is part is received. In an embodiment, network duplication proceeds, such as in a manner described above. At a point in time during network duplication, the network object 902 is paused. Information about the network object at pause 904 is copied to produce a copy of the network object at pause 906. For instance, all information regarding the state of the network object at pause 904 may be duplicated and stored in a separate location. In the example of the network object 902 being a virtual storage volume, the information stored in the virtual storage volume may be copied to one or more physical storage devices different from one or more physical storage devices used to store the virtual storage volume, or at least different locations on one or more of the same storage devices. In an embodiment, once the copy of the network object at pause is made, the network object at pause 904 is resumed. The network object and copy of the network object then proceed separately and independently.

[0052] While Figures 8 and 9 show illustrative examples of manners in which shallow and deep copies may be made, respectively, other methods may be used. As an example, the state of an object may be used as a basis for both the original object and the duplicate object. In other words, an object that has been copied may, after the copy, operate in a manner similar to the way the duplicate object operates, as described, referring to stored information about the object's state at the time of the copy, as necessary. As another example, when making a copy, techniques may be used that do not require pausing network objects if enough care is taken to ensure data integrity. For example, if care is made to record and account for changes to a network object during the duplication process, pausing may not be necessary. Generally, any technique for copying network objects, whether shallow or deep, may be used in accordance with the various embodiments of the present disclosure.

[0053] Figure 10 shows an illustrative example of a process 1000 that may be used to take advantage of various embodiments of the present disclosure. The process 1000, for instance, may be used forensically to investigate various events that occur on a network. In an embodiment, network events are recorded 1002. Network events may be any events that occur in connection with the network including, but not limited to, transfer of information from one device to another, administrative changes to the network including administrative changes made using APIs, network auto-scaling events, operations on data sets such as read,

write, or delete operations, and the like. Recording the network events may include storing information corresponding to the events in a manner such that the events may be at least partially reconstructed. Events may be recorded in a variety of ways. For example, in an embodiment, customers of a remote computing service provider may make API calls that affect the customers' networks in various ways, as determined by the API calls. These customer-initiated API calls (or, generally, any API calls that affect a network being copied) may be recorded. As another example, embodiments of the present disclosure used in connection with embodiments described in U.S. Patent Application No. 12/060,074, noted above, may involve routing network communications through one or more systems that are able to record packets traveling through the network, as well as routing information for the packets. While Figure 10 shows recording network events as a discrete step in the process 1000, recording (and generally other actions of the various disclosed processes, and variations thereof) may take place continuously during performance of other actions, and/or in an order different from that which is illustrated.

[0054] In an embodiment, a particular network event is detected. A network administrator, for example, may notice that the network behaves apparently abnormally at a regular time, such as the same time every day. One or more applications may detect one or more anomalies in the network performance, such as an unusually high latency, traffic volume, or other characteristics regarding the performance of one or more devices in the network. Other network events that may be detected include auto-scaling events, device malfunction, and the like. Generally, any way of detecting any type of particular network event, or combination of events, may be used. When a particular network event is detected, a duplicate network may be created 1006, in accordance with an embodiment. For example, a network administrator may utilize the interface of a network management console to submit a request that the network be duplicated. A duplicate network may be created 1006 responsive to the request. Duplication of the network may be performed in a manner described above and may take into account various copy-type, resolution, and other criteria specified by the request. The duplicate network may include all of the objects in the original network or may have a different amount.

[0055] In an embodiment, when the duplicate network is created, recorded network events may be replayed 1008 on the duplicate network one or more times in order to allow analysis of the events. For example, events may be replayed and recorded on the duplicate network in order to identify the cause of the detected network events or, simply, to gain a further understanding of the network capabilities. In this manner, analysis with regard to the

network operation may be performed on a duplicate network so as to allow the original network to continue normal operation. In other words, the duplicate network may be used for analysis while allowing the original network to operate normally, thereby allowing analysis of the network with minimal compromise to the network's performance.

5 [0056] Figure 11 illustrates another process 1100 that may be used in order to take advantage of the various embodiments of the present disclosure. The process 1100 may be used, for example, to allow a network administrator to study how a network will operate if certain changes are made before actually making the changes. In an embodiment, an instruction to duplicate a network is received 1102. The instruction to duplicate the network
10 may be received in accordance with the above description. When the instruction to duplicate the network is received, a duplicate network is created 1104, such as in a manner described above.

[0057] With a duplicate network, a network administrator or other user may modify 1106 the duplicate network for various purposes. Modification of the duplication network, for
15 example, may include reconfiguring devices in the network, adding and/or removing devices to/from the network, increasing bandwidth among various devices of the network, and the like. In an embodiment, upon modification of the duplicate network, the modified duplicate network may be used. For instance, the modified duplicate network may be used for simulation 1108 of how the network may respond under various stresses. A user may, for
20 example, may send requests to various devices in the network in order to stress the network and measure how the modified duplicate network is able to handle the stress. For example, a user may stress a network to determine whether the modified duplicate network auto-scales one or more portions of the modified duplicate network in order to accommodate the additional stress. Generally, the user may cause various loads to be generated for the
25 modified duplicate network. Loads may include requests to access information, to store information, to modify information stored in one or more data stores, to perform computations to be performed, and the like, multiple simultaneous connections to servers, and, in general, any action that can be performed by a network.

[0058] Further, portions of the processes described in connection with Figures 10 and 11,
30 and other processes described herein, may be combined. For instance, network events may be recorded and then used for simulations on a modified duplicate network to enable determining whether any modifications made improved network performance. Other variations are also contemplated as being within the spirit of the present disclosure.

[0059] Clauses:

Clause 1. A computer-implemented method for network duplication, comprising:
under the control of one or more computer systems configured with executable instructions,

5 operating a first network based at least in part on instructions provided remotely by a customer, the first network including a plurality of virtual network devices in a configuration that overlays one or more physical networks;

receiving, from the customer, a web service request to duplicate the first network;
responsive to the web service request:

10 retrieving configuration information about the first network from a data store;
generating a network duplication plan from the retrieved configuration information;
creating, according to the network duplication plan, a plurality of virtual duplicate devices, each of the virtual duplicate devices corresponding to a device of the first network;

configuring, according to the network duplication plan, the created virtual duplicate
15 devices in a duplicate network that overlays the one or more physical networks in a manner similar to the manner in which the first network overlays the one or more physical networks;
and

operating the duplicate network based at least in part on instructions provided
remotely by the customer.

20 Clause 2. The computer-implemented method of clause 1, wherein the first network includes one or more devices selected from the group consisting of servers, storage volumes, switches, hubs, firewalls, bridges, and gateways.

Clause 3. The computer-implemented method of clause 1, wherein one or more of the
virtual duplicate devices operate dependently from one or more stored state representations
25 used by one or more corresponding devices of the first network.

Clause 4. The computer-implemented method of clause 1, wherein one or more of the
virtual duplicate devices are complete copies of corresponding devices of the first network.

Clause 5. The computer-implemented method of clause 1, wherein the first network is a
sub-network of a larger network, wherein receiving the request to duplicate the first network
30 includes receiving one or more filtering criteria, and wherein the method further includes
applying the one or more filtering criteria to the larger network to determine the first network.

Clause 6. The computer-implemented method of clause 1, wherein at least one of the
virtual network devices of the first network was created using a template and wherein

creating the plurality of virtual duplicate devices includes using the template to create at least one virtual duplicate device corresponding to said at least one of the virtual network devices.

Clause 7. A computer-implemented method for network duplication, comprising:

under the control of one or more computer systems configured with executable
5 instructions,

receiving a request to duplicate a first network that includes one or more virtual
devices and that overlays one or more physical networks;

responsive to receipt of the request, creating a plurality of virtual duplicate devices
each corresponding to a device in the first network and arranging the created virtual duplicate
10 devices in a duplicate network to overlay the one or more physical networks in a manner
similar to the manner in which the first network overlays the one or more physical networks;
and

providing access to the duplicate network.

Clause 8. The computer-implemented method of clause 7, wherein providing access to
15 the duplicate network includes:

receiving instructions for performing a test of the duplicate network;
performing the test of the duplicate network; and
providing results of the test.

Clause 9. The computer-implemented method of clause 7, further comprising:

20 recording one or more network events for the first network; and
wherein providing access to the duplicate network includes causing the recorded
one or more network events to occur on the duplicate network.

Clause 10. The computer-implemented method of clause 9, wherein recording the one or
more network events includes recording one or more API calls affecting the first network and
25 wherein causing the recorded one or more network events to occur on the duplicate network
includes causing the one or more API calls to be made with respect to the duplicate network.

Clause 11. The computer-implemented method of clause 7, wherein providing access to
the duplicate network includes:

modifying the duplicate network; and
30 performing a test on the modified duplicate network.

Clause 12. The computer-implemented method of clause 11, wherein modifying the
duplicate network includes at least one of: changing capacity of at least one device of the
duplicate network; causing a change in packet loss to the duplicate network; causing a change

in failure rate in the duplicate network; and replacing one or more virtual duplicate devices with another virtual duplicate device that performs differently than the one or more virtual duplicate devices being replaced.

Clause 13. The computer-implemented method of clause 7, wherein at least one of the virtual network devices of the first network was created using a template and wherein creating the plurality of virtual duplicate devices includes using the template to create at least one virtual duplicate device corresponding to said at least one of the virtual network devices.

Clause 14. The computer-implemented method of clause 7, wherein creating the plurality of devices includes:

for a particular device of the first network, identifying whether the particular device is stateful device; and

when determined that the particular device is a stateful device, creating an copy of the particular device that operates operate dependently from one or more stored state representations used by the particular device.

Clause 15. One or more computer-readable storage media having collectively stored thereon instructions that, when executed by one or more computer systems, cause the one or more computer systems to collectively at least:

receive a request to duplicate a first network that includes one or more virtual devices and that overlays one or more physical networks;

responsive to receipt of the request, create a duplicate network that overlays the one or more physical networks with a topology consistent with the first network and that includes one or more virtual duplicate devices that each correspond to one of the one or more virtual devices of the first network; and

provide access to the duplicate network.

Clause 16. The one or more computer-readable storage media of clause 15, wherein providing access to the duplicate network includes:

receiving instructions for performing a test of the duplicate network;

performing the test of the duplicate network; and

providing results of the test.

Clause 17. The one or more computer-readable storage media of clause 15, wherein the instructions further cause the one or more computer systems to collectively at least:

record one or more network events for the first network; and

wherein providing access to the duplicate network includes causing the recorded one or more network events to occur on the duplicate network.

Clause 18. The one or more computer-readable storage media of clause 15, wherein creating the duplicate network includes configuring at least a subset of the duplicate virtual devices to have identifiers identical to identifiers of a set of the one or more virtual devices of the first network.

Clause 19. The one or more computer-readable storage media of clause 15, wherein the identifiers are internet protocol addresses.

Clause 20. A computer system operable to duplicate networks, comprising:

one or more processors; and
memory, including executable instructions that, when executed by the one or more processors, cause the computer system to at least:

receive a request to duplicate a first network that includes one or more virtual devices and that overlays one or more physical networks;

responsive to receipt of the request, create a plurality of virtual duplicate devices each corresponding to a device in the first network and arranging the created devices in a virtual duplicate network to overlay the one or more physical networks in a manner similar to the manner in which the first network overlays the one or more physical networks; and
provide access to the duplicate network.

Clause 21. The computer system of clause 20, wherein at least a subset of the created plurality of virtual duplicate devices are copy-on-write overlays of devices in the first network.

Clause 22. The computer system of clause 20, wherein providing access to the duplicate network includes:

receiving instructions for performing a test of the duplicate network;
performing the test of the duplicate network; and
providing results of the test.

Clause 23. The computer system of clause 22, further comprising modifying the duplicate network prior to performing the test.

Clause 24. The computer system of clause 20, wherein the first network includes one or more devices selected from the group consisting of servers, storage volumes, switches, hubs, firewalls, bridges, and gateways.

Clause 25. A computer-implemented method of network duplication, comprising:

under the control of one or more computer systems configured with executable instructions,

receiving a request to duplicate a first network that overlays one or more physical
5 networks; and

responsive to receipt of the request:

creating a plurality of virtual devices and configuring the plurality of devices in a
virtual network such that the plurality of virtual devices are communicatively connected with
one another in a manner so as to overlay the one or more physical networks in a manner
10 similar to the manner in which the first network overlays the one or more physical networks,
the plurality of devices including one or more virtual devices that correspond to one or more
devices of the first network.

Clause 26. The computer-implemented method of clause 25, wherein receiving the
request includes receiving information specifying a fidelity for one or more devices of the
15 first network; and

wherein creating the new virtual network includes creating one or more duplicate
virtual devices according to the specified fidelity.

Clause 27. The computer-implemented method of clause 25, wherein the request is
received remotely from a user; and

20 wherein creating the plurality of virtual devices is performed in a facility in which one or
more other networks for one or more other users are operated.

Clause 28. The computer-implemented method of clause 25, wherein receiving the
request to duplicate the first network includes receiving one or more filtering criteria, and
wherein the method further includes applying the one or more filtering criteria to the larger
25 network to determine said at least a subset of the first network.

Clause 29. The computer-implemented method of clause 25, wherein the first network is
a virtual overlay network.

[0060] The various embodiments further can be implemented in a wide variety of operating
environments, which in some cases can include one or more user computers, computing
30 devices, or processing devices which can be used to operate any of a number of applications.
User or client devices can include any of a number of general purpose personal computers,
such as desktop or laptop computers running a standard operating system, as well as cellular,
wireless, and handheld devices running mobile software and capable of supporting a number

of networking and messaging protocols. Such a system also can include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These devices also can include other electronic devices, such as dummy terminals, thin-clients, gaming systems, and other devices capable of communicating via a network.

[0061] Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as TCP/IP, OSI, FTP, UPnP, NFS, CIFS, and AppleTalk. The network can be, for example, a local area network, a wide-area network, a virtual private network, the Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network, and any combination thereof.

[0062] In embodiments utilizing a Web server, the Web server can run any of a variety of server or mid-tier applications, including HTTP servers, FTP servers, CGI servers, data servers, Java servers, and business application servers. The server(s) also may be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java[®], C, C# or C++, or any scripting language, such as Perl, Python, or TCL, as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle[®], Microsoft[®], Sybase[®], and IBM[®].

[0063] The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (“SAN”) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers, or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (*e.g.*, a mouse, keyboard, controller, touch screen, or keypad), and at least one output device (*e.g.*, a display device, printer, or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices, and solid-state storage devices such as random access memory (“RAM”) or

read-only memory ("ROM"), as well as removable media devices, memory cards, flash cards, etc.

[0064] Such devices also can include a computer-readable storage media reader, a communications device (*e.g.*, a modem, a network card (wireless or wired), an infrared communication device, etc.), and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed, and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services, or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or Web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including portable software, such as applets), or both. Further, connection to other computing devices such as network input/output devices may be employed.

[0065] Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules, or other data, including RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0066] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

[0067] Other variations are within the spirit of the present disclosure. Thus, while the disclosed techniques are susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the

5 invention to the specific form or forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions, and equivalents falling within the spirit and scope of the invention, as defined in the appended claims.

[0068] The use of the terms “a” and “an” and “the” and similar referents in the context of describing the disclosed embodiments (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. The term “connected” is to be construed as partly or wholly contained within, attached to, or joined together, even if there is something

15 intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context.

20 The use of any and all examples, or exemplary language (*e.g.*, “such as”) provided herein, is intended merely to better illuminate embodiments of the invention and does not pose a limitation on the scope of the invention unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the invention.

25 [0069] Preferred embodiments of this disclosure are described herein, including the best mode known to the inventors for carrying out the invention. Variations of those preferred embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. The inventors expect skilled artisans to employ such variations as appropriate, and the inventors intend for the invention to be practiced otherwise than as specifically described herein. Accordingly, this invention includes all modifications and

30 equivalents of the subject matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the invention unless otherwise indicated herein or otherwise clearly contradicted by context.

[0070] All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

WHAT IS CLAIMED IS:

1 1. A computer-implemented method for network duplication, comprising:
2 under the control of one or more computer systems configured with executable
3 instructions,
4 receiving a request to duplicate a first network that includes one or more
5 virtual devices and that overlays one or more physical networks;
6 responsive to receipt of the request, creating a plurality of virtual duplicate
7 devices each corresponding to a device in the first network and arranging the created virtual
8 duplicate devices in a duplicate network to overlay the one or more physical networks in a
9 manner similar to the manner in which the first network overlays the one or more physical
10 networks; and
11 providing access to the duplicate network.

1 2. The computer-implemented method of claim 1, wherein providing
2 access to the duplicate network includes:
3 receiving instructions for performing a test of the duplicate network;
4 performing the test of the duplicate network; and
5 providing results of the test.

1 3. The computer-implemented method of claim 1, further comprising:
2 recording one or more network events for the first network; and
3 wherein providing access to the duplicate network includes causing the
4 recorded one or more network events to occur on the duplicate network.

1 4. The computer-implemented method of claim 3, wherein recording the
2 one or more network events includes recording one or more API calls affecting the first
3 network and wherein causing the recorded one or more network events to occur on the
4 duplicate network includes causing the one or more API calls to be made with respect to the
5 duplicate network.

1 5. The computer-implemented method of claim 1, wherein providing
2 access to the duplicate network includes:
3 modifying the duplicate network; and
4 performing a test on the modified duplicate network.

1 6. The computer-implemented method of claim 5, wherein modifying the
2 duplicate network includes at least one of: changing capacity of at least one device of the
3 duplicate network; causing a change in packet loss to the duplicate network; causing a change
4 in failure rate in the duplicate network; and replacing one or more virtual duplicate devices
5 with another virtual duplicate device that performs differently than the one or more virtual
6 duplicate devices being replaced.

1 7. The computer-implemented method of claim 1, wherein at least one of
2 the virtual network devices of the first network was created using a template and wherein
3 creating the plurality of virtual duplicate devices includes using the template to create at least
4 one virtual duplicate device corresponding to said at least one of the virtual network devices.

1 8. The computer-implemented method of claim 1, wherein creating the
2 plurality of devices includes:

3 for a particular device of the first network, identifying whether the particular
4 device is stateful device; and

5 when determined that the particular device is a stateful device, creating an
6 copy of the particular device that operates operate dependently from one or more stored state
7 representations used by the particular device.

1 9. A computer system operable to duplicate networks, comprising:

2 one or more processors; and

3 memory, including executable instructions that, when executed by the one or
4 more processors, cause the computer system to at least:

5 receive a request to duplicate a first network that includes one or more virtual
6 devices and that overlays one or more physical networks;

7 responsive to receipt of the request, create a plurality of virtual duplicate
8 devices each corresponding to a device in the first network and arranging the created devices
9 in a virtual duplicate network to overlay the one or more physical networks in a manner

10 similar to the manner in which the first network overlays the one or more physical networks;

11 and

12 provide access to the duplicate network.

1 10. The computer system of claim 9, wherein at least a subset of the
2 created plurality of virtual duplicate devices are copy-on-write overlays of devices in the first
3 network.

1 11. The computer system of claim 9, wherein providing access to the
2 duplicate network includes:

3 receiving instructions for performing a test of the duplicate network;
4 performing the test of the duplicate network; and
5 providing results of the test.

1 12. The computer system of claim 11, further comprising modifying the
2 duplicate network prior to performing the test.

1 13. The computer system of claim 9, wherein the first network includes
2 one or more devices selected from the group consisting of servers, storage volumes, switches,
3 hubs, firewalls, bridges, and gateways.