

**(19) AUSTRALIAN PATENT OFFICE**

(54) Title  
Method and apparatus for discovering network devices

(51)<sup>6</sup> International Patent Classification(s)  
**G06F** 13/00 (2006.01) 12/24  
**G06F** 1/00 (2006.01) 20060101ALI2005100  
**H04L** 12/24 (2006.01) 8BMEP **H04L**  
**H04L** 12/28 (2006.01) 12/28  
**H04L** 12/56 (2006.01) 20060101ALI2006010  
**H04L** 29/08 (2006.01) 1BMKR **H04L**  
G06F 13/00 12/56  
20060101AFI2005122 20060101ALI2005122  
0BMJP **G06F** 0BMJP **H04L**  
1/00 29/08  
20060101ALI2005102 20060101ALI2007072  
7BMRU **H04L** 1BMEP

(21) Application No: 2004201420 (22) Application Date: 2004 .04 .02

(30) Priority Data

(31) Number	(32) Date	(33) Country
10/425304	2003 .04 .29	US

(43) Publication Date : 2004 .11 .18

(43) Publication Journal Date : 2004 .11 .18

(71) Applicant(s)  
Microsoft Corporation

(72) Inventor(s)  
Tabbara, Bassam

(74) Agent/Attorney  
Davies Collison Cave, 1 Nicholson Street, Melbourne, VIC, 3000

(56) Related Art  
US 2002/154606 A1

**ABSTRACT**

Multiple devices are identified in a network along with connection information associated with the multiple devices. The connection information associated with the multiple devices is converted into a hierarchical structure and  
5 communication links between the multiple devices are identified.

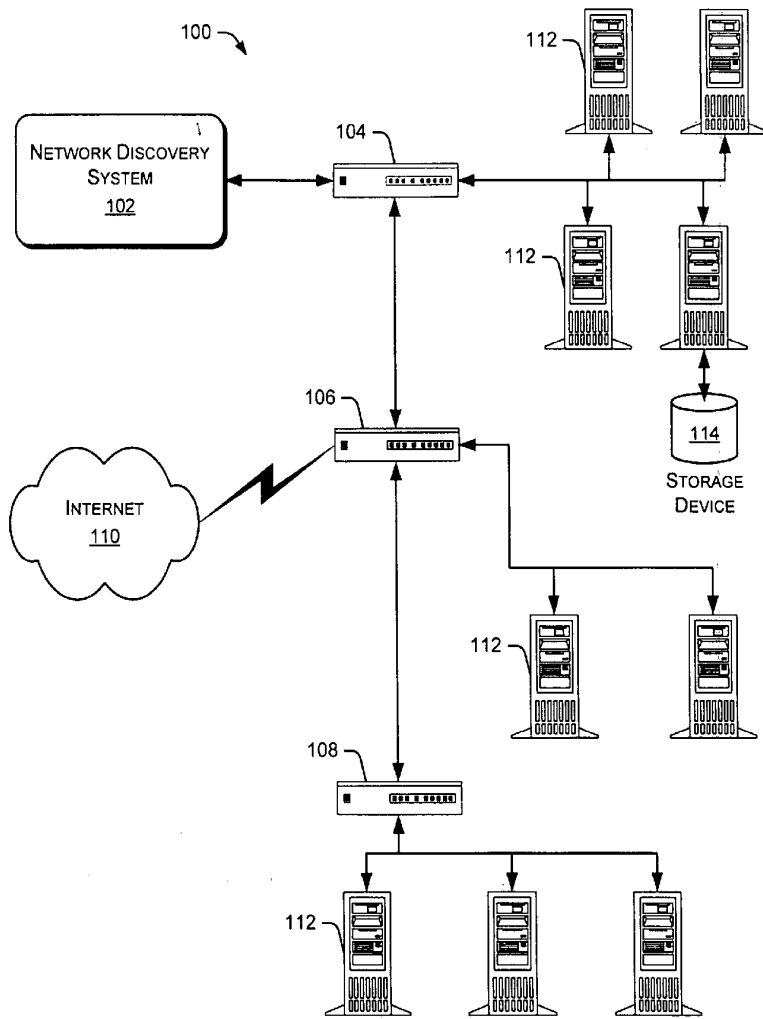


Fig. 1

**AUSTRALIA**  
**PATENTS ACT 1990**  
**COMPLETE SPECIFICATION**

NAME OF APPLICANT(S)::

**Microsoft Corporation**

ADDRESS FOR SERVICE:

**DAVIES COLLISON CAVE**  
Patent Attorneys  
1 Nicholson Street, Melbourne, 3000, Australia

INVENTION TITLE:

Method and apparatus for discovering network devices

The following statement is a full description of this invention, including the best method of performing it known to me/us:-

5102

2004201420 15 Sep 2009

- 1a -

**Technical Field**

The present invention relates to a method and apparatus for discovering devices in a network environment and determining a relationship between those devices, for  
5 example.

**Background of the Invention**

Internet usage has exploded over the past several years and continues to grow.  
10 People have become very comfortable with many services offered on the World Wide Web (or simply "Web"), such as electronic mail, online shopping, gathering news and other information, listening to music, viewing video clips, looking for jobs, and so forth. To keep pace with the growing demand for Internet-based services, there has been tremendous growth in the computer systems dedicated to  
15 hosting Web sites, providing backend services for those sites, and storing data associated with those sites.

One type of distributed computer system is an Internet data center (IDC), which is a specifically designed center that houses many computers for hosting Internet-based  
20 services. IDCs, which are also referred to as "Web farms" or "server farms", typically contain hundreds to thousands of computers in climate-controlled, physically secure buildings. These computers are interconnected to run one or more programs supporting one or more Internet services or Web sites. IDCs provide reliable Internet access, reliable power supplies, and a secure operating  
25 environment.

A particular data center may include, for example, multiple computers (such as general purpose computers), multiple storage devices and multiple data handling devices, such as routers, hubs, gateways and switches. The various devices and  
30 computers are interconnected with each other and other networks, such as the

2004201420 15 Sep 2009

- 2 -

Internet, thereby allowing the communication of data between various devices in the data center.

Existing systems utilize a significant amount of manual control of the computers and other devices in a data center. For example, when a new resource (such as a computer system) is added to a data center, certain manual steps may be performed to configure the new resource and to notify other devices in the data center of the availability of the new resource. These manual operations are time-consuming and may not be performed for some time if an administrator is busy with other tasks.

10

Accordingly, there is a need for improved techniques for discovering network topologies and managing various devices in a network.

It is generally desirable to overcome or ameliorate one or more of the above described difficulties, or to at least provide a useful alternative.

15

#### **Summary of the Invention**

In accordance with one aspect of the present invention, there is provided a method comprising:

20

- identifying a plurality of devices in a network;
- tagging each identified device in the network with a discovery identifier, wherein the discovery identifier identifies a particular discovery process,
- identifying connection information associated with the plurality of devices;
- 25 converting the connection information associated with the plurality of devices into a hierarchical structure;
- identifying communication links between the devices; and
- storing network topology data associated with the network.

25

2004201420 15 Sep 2009

- 3 -

In accordance with another aspect of the present invention, there is provided a method comprising:

- 5 creating a discovery identifier to identify a particular discovery process, tagging each identified device in the network with the discovery identifier;
- identifying a plurality of devices in a network;
- identifying connection information associated with the plurality of devices;
- identifying spanning tree information associated with the plurality of devices;
- 10 identifying communication links between the devices using the spanning tree information; and storing network topology data associated with the network.

In accordance with another aspect of the present invention, there is provided a method comprising:

- 15 creating a discovery identifier to identify a particular discovery process, tagging each identified device in the network with the discovery identifier,
- identifying a range of IP addresses to discover;
- querying each IP address in the range of IP addresses;
- if a response to the query is received:
  - 20 requesting management information base data from a responding device;
  - identifying a device type associated with the responding device; and
  - retrieving additional information from the responding device based on the device type associated with the responding device wherein the management information base data and the additional information is stored in a system
  - 25 link table.

In accordance with another aspect of the present invention, there is provided a method comprising:

- initiating a device discovery process associated with a network;

2004201420 15 Sep 2009

- 3a -

- creating a discovery identifier associated with the device discovery process;
- identifying a plurality of devices in a network;
- tagging the devices with the discovery identifier;
- maintaining a status of each of the plurality of devices;
- 5 repeatedly querying each of the plurality of devices;
- recording responses received from each of the plurality of devices;
- if no response is received from a particular device within a predetermined period of time:
  - changing the status of the particular device to inactive; and
  - 10 notifying other devices in the network of the changed status of the particular device.

In accordance with another aspect of the present invention, there is provided an apparatus comprising:

- 15 means for identifying a plurality of devices in a network;
- means for identifying connection information associated with the plurality of devices;
- means for representing the connection information associated with the plurality of devices in a hierarchical structure;
- 20 means for tagging each of the plurality of devices with a discovery identifier, wherein the discovery identifier identifies a particular discovery process; and
- means for storing network topology data associated with the network.

In accordance with another aspect of the present invention, there is provided an apparatus comprising:

- 25 a discovery engine coupled to a network, wherein the discovery engine is configured to create a discovery identifier to identify a particular discovery process,

2004201420 15 Sep 2009

- 3b -

5 identify and tag with the discovery identifier a plurality of network devices in the network and identify connection information associated with the plurality of network devices, the discovery engine further configured to identify spanning tree information associated with the plurality of network devices and identify links between the network devices using the spanning tree information; and  
a storage device coupled to the discovery engine and configured to store network topology data associated with the network.

10 In accordance with another aspect of the present invention, there is provided one or more computer-readable storage media having stored thereon a computer program that, when executed by one or more processors, causes the one or more processors to:

15 create a discovery identifier to identify a particular discovery process,  
identify a plurality of IP addresses to discover;  
query each IP address in the plurality of IP addresses;  
if a query response is received:  
tag each identified device with the discovery identifier,  
20 request management information base data from the responding device;  
identify a device type associated with the responding device; and  
retrieve additional information from the responding device based on the device type associated with the responding device wherein the management information base data and the additional information is stored  
25 in a system link table.

In accordance with another aspect of the present invention, there is provided a method comprising:

initiating a device discovery process associated with a network;

- 3c -

creating a discovery ID associated with the device discovery process;  
identifying new devices in the network by searching a range of IP addresses;  
and  
tagging the new devices with the discovery ID.

5

In accordance with another aspect of the present invention, there is provided an apparatus comprising:

means for initiating a device discovery process associated with a network;  
means for creating a discovery ID associated with the device discovery  
10 process;  
means for identifying new devices in the network by searching a range of IP addresses; and  
means for tagging the new devices with the discovery ID.

15

The systems and methods described herein preferably discover various devices in a network environment and determine relationships between those devices. The systems and methods also monitor the network devices and communication links between the devices to maintain a current database of information regarding all devices in the network.

20

In one embodiment, multiple devices in a network are identified along with connection information associated with the multiple devices. The connection information is converted into a hierarchical structure. This embodiment then identifies communication links between the devices.

25

**Brief Description of the Drawings**

Preferred embodiments of the present invention are hereafter described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

2004201420 15 Sep 2009

- 3d -

Fig. 1 shows an exemplary network environment.

Fig. 2 is a block diagram illustrating an embodiment of a network discovery system.

5

Fig. 3 is a flow diagram illustrating an embodiment of a procedure for discovering devices in a network and the topology of the network.

10 Fig. 4 is a flow diagram illustrating an embodiment of a procedure for identifying and obtaining information from network devices.

Fig. 5 illustrates a data storage structure for storing data related to various network devices and communication links between network devices.

15 Fig. 6 is a flow diagram illustrating an embodiment of a procedure for identifying link information associated with the network devices identified using the procedure of Fig. 4.

20 Fig. 7 is a flow diagram illustrating an embodiment of a procedure for discovering network devices.

Fig. 8 illustrates an example of a computing environment.

25 **Detailed Description of Preferred Embodiments of the Invention**

The systems and methods described herein relate to discovering multiple devices in a network environment and identifying a relationship between those devices. These systems and methods automate certain tasks that would otherwise be performed manually. A database maintains information about the multiple

devices in the network environment and the manner in which the devices are coupled to one another.

Various examples discussed herein refer to networks and network environments. As used herein, a “network” is any collection of two or more devices coupled to one another such that the devices can exchange data with one another. A data center, such as an Internet data center (IDC), represents one example of a network environment. The various devices in a network environment may be located in a common area or located at different geographic locations. A particular network environment may include one or more sub-networks coupled to one another. The networks discussed herein may utilize any data communication protocol and any type of data communication medium.

As used herein, a “network device” is any device or system coupled to or contained in a network. A network device may also be referred to as a network node. Example network devices include any type of computer system, storage system, bridge, router, switch, hub, network discovery system, and the like.

In particular embodiments, SNMP (Simple Network Management Protocol) is used to communicate between various network devices. SNMP is a set of protocols for managing networks. SNMP sends messages, called protocol data units (PDUs) to different parts of a network. SNMP-compliant devices (referred to as “agents”) store data about themselves in Management Information Bases (MIBs) and return this data to SNMP requesters. In alternate embodiments, other protocols may be used instead of or in addition to SNMP.

Fig. 1 shows an exemplary network environment 100. A network discovery system 102 is coupled to a bridge 104. As described herein, network discovery system 102 is capable of discovering various devices in a network environment and discovering communication links between the network devices. Additionally,

network discovery system 102 monitors the network devices for proper operation and maintains a database of information regarding the devices in the network and the manner in which those devices are coupled to one another. Bridge 104 is a device that connects two or more segments of the same network using a common  
5 protocol, such as Ethernet or Token-Ring. In one embodiment, bridge 104 contains multiple communication ports for sending and receiving data. Bridge 104 maintains a forwarding table associated with each port. Each forwarding table contains addresses (e.g., MAC addresses) of other network devices coupled to the port.

10 Bridge 104 is coupled to four computing devices 112 and another bridge 106. Bridge 104 is coupled to computing devices 112 via, for example, a bus or multiple individual connections between the bridge and each of the computing devices. Computing devices 112 may be any type of computer, including servers, workstations, laptop computers, portable computers, and the like. One of the  
15 computing devices 112 is coupled to a storage device 114 for storing data.

Bridge 106 is coupled to the Internet 110, two computing devices 112 and another bridge 108. Bridge 108 is coupled to three computing devices 112. The arrangement of network devices shown in Fig. 1 allows any network device to exchange data with any other network device in network environment 100.

20 Although bridges 104, 106 and 108 in Fig. 1 are shown coupled to a small number of computing devices 112, a particular bridge may be connected to any number of computing devices or other network devices. An embodiment of a bridge 104, 106 or 108 may be coupled to hundreds or thousands of different network devices. Alternate network environments may include other types of  
25 network devices, such as routers, gateways, switches, hubs, and the like.

Fig. 2 is a block diagram illustrating an embodiment of network discovery system 102. Network discovery system 102 includes a discovery engine 202, a storage device 204 to store network data, a network monitor 206 and a display device 208. Discovery engine 202 discovers network devices in a network environment, identifies communication links between the network devices and determines the topology of the network environment. The data collected and generated by discovery engine 202 is stored in storage device 204.

Network monitor 206, in combination with discovery engine 202, monitors the network environment to identify changes in the network, such as the addition of new network devices, the removal of network devices and changes to communication links between the network devices. The data stored in storage device 204 is updated by network monitor 206 to reflect any changes in network environment 100. Network monitor 206 is coupled to a display device 208, which allows a user (such as a network administrator) to view the network topology, the status of one or more network devices, or any other information retrieved or generated by network discovery system 102. A user input device 210 is also coupled to network monitor 206 and allows the user to provide input to the network monitor. Input device 210 may be, for example, a keyboard, a mouse, a track pad, a touch pad, etc. Network monitor 206 allows a user, such as a network administrator, to retrieve information about the network (e.g., network information stored in storage device 204).

A network interface 212 is coupled to discovery engine 202 and allows the discovery engine to communicate with other network devices in a network environment. Network interface 212 is coupled to one or more network devices in network environment 100.

Fig. 3 is a flow diagram illustrating an embodiment of a procedure 300 for discovering devices in a network and the topology of the network topology.

Procedure 300 can be implemented by, for example, network discovery system 102.

Initially, procedure 300 identifies devices in a network (block 302). This

5 identification can be performed, for example, by pinging various network addresses (also referred to as Internet Protocol (IP) addresses), as discussed below. The procedure then identifies connection information associated with the devices in the network (block 304). Connection information is identified by retrieving information from each of the identified network devices, as discussed below.

10 At block 306, the procedure identifies spanning information associated with the devices identified above. Spanning trees are generated by devices, such as switches, using a standard algorithm. In one embodiment, procedure 300 retrieves spanning information from one or more devices in the network. The spanning tree information is useful in computing the physical network topology. The spanning  
15 tree protocol is a link management protocol that provides path redundancy while preventing undesirable loops in a network. For example, for an Ethernet network to function properly, there should be a single active path between two network nodes. If multiple active paths exist between two network nodes, a loop may result. When a loop exists, the potential exists for a duplication of messages. The spanning tree  
20 protocol forces certain redundant data paths into a standby (i.e., blocked) state. Additionally, one or more bidirectional communication links between network nodes may be converted to unidirectional links.

After identifying spanning tree information at block 306, the procedure of Fig. 3 continues by identifying links between the network devices using information  
25 contained in the spanning tree (block 308). The procedure then computes the physical network topology (block 310). The procedure stores the network topology

data associated with the network in a database (block 312), such as storage device 204 (Fig. 2). Finally, procedure 300 generates a representation of the network topology using the stored network topology data (block 314). This representation of the network topology can be displayed on a display device, printed on a printer, 5 or otherwise rendered for a network administrator or other user. The representation of the network topology may be saved for future reference and/or communicated to another system or user.

Although particular examples discussed herein utilize information contained in one or more spanning trees, alternate embodiments may use other protocols 10 and/or processes to eliminate undesirable loops in a network.

Fig. 4 is a flow diagram illustrating an embodiment of a procedure 400 for identifying and obtaining information from network devices. Initially, procedure 400 identifies a range of IP addresses to discover (block 402). This may be a continuous range of IP addresses or any number of non-continuous IP addresses. 15 The range of IP addresses may be specified by a network administrator or the discovery engine may scan the entire IP address space. The procedure selects the first IP address in the range and pings the IP address (block 404). If a response is not received, the procedure selects the next IP address in the range and pings that IP address (block 414).

20 If a response is received after pinging an IP address, the procedure queries the network device associated with the IP address to obtain its management information base (MIB) data (block 408). In one embodiment, this query is an SNMP query that receives information from the device's system table (discussed below with respect to Fig. 5). A MIB is a database of objects that can be monitored, 25 for example, by a network management system. After receiving the MIB data, procedure 400 retrieves additional information from the network device based on its

device type (block 410). The device type information is part of the MIB data received from a network device. For example, if the device type is a computer system, the procedure retrieves information regarding the computer system (e.g., type of processor, memory capacity, hard disk storage capacity, and the like). If the device type is a bridge, the procedure retrieves additional information using a MIB that provides information specific to bridges.

After retrieving additional information from the device based on its device type, the procedure stores the MIB information and additional information in a system link table (block 412). The system link table is discussed below. The procedure then selects the next IP address in the range and pings the IP address (block 414). The procedure returns to block 406 to determine whether a response is received from the IP address.

Fig. 5 illustrates a data storage structure for storing data related to various network devices and communication links between network devices. The data storage structure of Fig. 5 is used to build a hierarchy of network devices. A "System" table 502 represents the main table for maintaining information regarding network devices. In System table 502, the "SystemID" field is an identifier associated with each system (or device). "SystemID" is unique within the data storage structure of Fig. 5. "SystemTypeID" identifies the type of device (e.g., computer, bridge, etc.). "DiscoveryID" is the identification of the discovery process during which the device was discovered. "OwnerID" identifies the owner of the device, such as a network resource manager. "ManagementIPAdr" is the IP address at which the device was discovered. "OID" refers to an object ID, which represents the manufacturer, model, etc. associated with the device. "Name" is the name of the device. "Description" represents a description of the device. "UpTime" identifies how long the device has been "up" or active. "Contact"

identifies a user or operator responsible for the device. "Location" identifies the physical location of the device (e.g., rack 4, shelf B). "Services" identify the type of services provided by the device, such as storage services, routing services, etc.

A "SystemPort" table 504 is associated with System table 502. There is a  
5 separate SystemPort table 504 for each port on the associated system or device  
(identified by SystemID). "PortNumber" represents a particular port (e.g., network  
communication port) on the device. "Type" identifies the port type, such as 10  
Mb/sec. "PhysicalAddress" is the address of the port, such as the MAC address,  
Phy Net address, and the like. "Speed" is the communication speed of the port.  
10 "OperationStatus" is the most recent identified status of the port.

A "BridgePort" table 506 is also associated with System table 502.  
"DesignatedRoot" is the address of the designated root of the spanning tree.  
"DesignatedBridge" is the address of another bridge if this port is connected to  
another bridge. Otherwise, "DesignatedBridge" is null. "DesignatedPortNumber"  
15 is the port on the other bridge, if any. "SingletonForwardingAdr" is a single  
address for forwarding data. If the forwarding table for a specific bridge port  
contains only one address, the value of SingletonForwardingAdr is set to that  
address. Otherwise, the value of SingletonForwardingAdr is null.  
SingletonForwardingAdr is useful in identifying leaf devices coupled to a switch  
20 port.

A "Discovery" table 508 is associated with System table 502. A separate  
"DiscoveryID" is associated with each execution of the discovery process.  
"StartTime" is the time that the discovery process started and "EndTime" is the time  
that the discovery process finished. "Status" identifies whether the discovery  
25 process was successful or whether an error occurred. "StartMessage" and

“EndMessage” identify messages generated at the start of the discovery process and at the end of the discovery process.

A “SystemLink” table 510 is associated with System table 502. SystemLink table 510 contains information regarding various links coupled to the system  
5 identified by SystemID. “SourceSystemID” identifies the system that represents the source of the link and “SourceSystemPortNumber” identifies the port associated with the source of the link. “TargetSystemID” identifies the system that is the target of the link and “TargetSystemPortNumber” identifies the port associated with the target of the link.

10 A “Computer” table 512 is associated with System table 502. “SmbiosUuid” is a unique identifier associated with the computer. “Name” is the name of the computer and “Vendor” identifies the vendor of the computer as well as information about the computer (such as installed memory, processor type, and number of hard disks).

15 A “Bridge” table 514 is also associated with System table 502. “BaseBridgeAdr” is a unique identifier associated with the bridge. “NumberPorts” identifies the number of ports supported by the bridge. “BaseType” identifies the type of bridging supported, such as source level bridging. “RootBridgeAdr” is the address of the root of the spanning tree.

20 An “Owner” table 516 provides information regarding the owner of the system or device. A “SystemType” table 518 provides information specific to the type of system.

“Bridge” and “Computer” are both subtypes of “System”. Although only two subtypes (Bridge and Computer) are discussed herein, alternate embodiments  
25 include other subtypes (all of which are subtypes of the same System table), such as “Storage”, “Load Balancers”, “Hardware”, etc.

Entries in Fig. 5 shown with a "\*" are unique identifiers (such as the SystemID, the DiscoveryID, and the PortNumber). All entries in the SystemLink table 510 have an associated "\*", indicating that the entries are all part of the primary key for that table.

5           The data storage structure illustrated in Fig. 5 is used to store information discovered regarding network devices and communication links between network devices. The stored information is updated as the network topology changes (e.g., as a result of adding new devices and links, or as a result of removing existing devices and links).

10           After identifying various network devices using the procedure of Fig. 4, the network discovery system computes the various links between the identified network devices. As the link information is computed, the information is stored in the data structure of Fig. 5. In particular, the link information is stored in SystemLink table 510.

15           Fig. 6 is a flow diagram illustrating an embodiment of a procedure 600 for identifying link information associated with the network devices identified using the procedure of Fig. 4. If there are multiple spanning trees associated with the network devices, procedure 600 is performed for each spanning tree. Initially, procedure 600 identifies a root of a first spanning tree (block 602). The root of the  
20 spanning tree is marked as Level 1 (block 604). At block 606, the procedure then identifies all network devices in the next level (i.e., Level 2) that have a designated port set to the previous level (i.e., Level 1). The procedure then adds the links associated with the identified devices to the SystemLink table (block 608). The identified devices are marked as Level 2 (block 610). The procedure then  
25 determines whether there are additional levels of network devices (block 612). If there are additional levels, the current level is set to the next level (block 614), and

the procedure returns to block 606 to identify network devices having a port set to the previous level. If there are no additional levels, the procedure branches to block 616 and selects the next spanning tree, if any. After all spanning trees have been analyzed by the procedure of Fig. 6, the process of filling the SystemLink table is  
5 complete.

Fig. 7 is a flow diagram illustrating an embodiment of a procedure 700 for discovering network devices. Procedure 700 can be performed in response to a user request, after a predetermined time period, after a particular event occurs, or any other triggering criteria. In a particular embodiment, the network discovery process  
10 is performed every 10-15 minutes. The network discovery process is used to monitor the status of network devices and communication links in a network environment. By maintaining the current status of the network devices and communication links (e.g., active or inactive), the network discovery system is able to maintain an accurate understanding of the network topology.

15 The network discovery procedure is activated (or triggered) at block 702. The procedure creates a new discovery ID associated with the current network discovery process (block 704). Discovery IDs may be issued sequentially or according to another predetermined pattern. Each discovery ID is unique.

The procedure discovers new network devices and identifies new network  
20 links (block 706). These new devices may be associated with IP addresses that did not previously respond to pinging. If any new devices are discovered, they are tagged with the discovery ID (block 708). If no new devices are discovered, the procedure continues to block 710 to determine whether all devices have been discovered (i.e., the full range of IP addresses has been searched). If all currently  
25 active devices have been discovered, the procedure ends. Otherwise, the procedure continues discovering network devices (block 712) and returns to block 706 to

determine whether a new network device or a new network link has been discovered.

If a particular network device was previously discovered, but has not been discovered during recent network discovery processes, the status of the device may  
5 be changed to "inactive". In one embodiment, the status of a particular network device is changed to inactive if a response has not been received from the device or link for more than one hour.

Fig. 8 illustrates an example of a computing environment 800 within which the data carousel systems and methods, as well as the computer, network, and  
10 system architectures described herein, can be either fully or partially implemented. Exemplary computing environment 800 is only one example of a computing system and is not intended to suggest any limitation as to the scope of use or functionality of the network architectures. Neither should the computing environment 800 be  
15 interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 800.

The computer and network architectures can be implemented with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or  
20 configurations that may be suitable for use include, but are not limited to, personal computers, server computers, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, gaming consoles, distributed computing environments that include any  
25 of the above systems or devices, and the like.

The computing environment 800 includes a general-purpose computing system in the form of a computing device 802. The components of computing device 802 can include, by are not limited to, one or more processors 804 (e.g., any of microprocessors, controllers, and the like), a system memory 806, and a system bus 808 that couples various system components including the processor 804 to the system memory 806. The one or more processors 804 process various computer-executable instructions to control the operation of computing device 802 and to communicate with other electronic and computing devices.

The system bus 808 represents any number of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, such architectures can include an Industry Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

Computing environment 800 typically includes a variety of computer-readable media. Such media can be any available media that is accessible by computing device 802 and includes both volatile and non-volatile media, removable and non-removable media. The system memory 806 includes computer-readable media in the form of volatile memory, such as random access memory (RAM) 810, and/or non-volatile memory, such as read only memory (ROM) 812. A basic input/output system (BIOS) 814, containing the basic routines that help to transfer information between elements within computing device 802, such as during start-up, is stored in ROM 812. RAM 810 typically contains data and/or program

modules that are immediately accessible to and/or presently operated on by the processing unit 804.

Computing device 802 can also include other removable/non-removable, volatile/non-volatile computer storage media. By way of example, a hard disk drive 5 816 is included for reading from and writing to a non-removable, non-volatile magnetic media (not shown), a magnetic disk drive 818 for reading from and writing to a removable, non-volatile magnetic disk 820 (e.g., a "floppy disk"), and an optical disk drive 822 for reading from and/or writing to a removable, non-10 volatile optical disk 824 such as a CD-ROM, DVD, or any other type of optical media. The hard disk drive 816, magnetic disk drive 818, and optical disk drive 822 are each connected to the system bus 808 by one or more data media interfaces 826. Alternatively, the hard disk drive 816, magnetic disk drive 818, and optical disk drive 822 can be connected to the system bus 808 by a SCSI interface (not shown).

15 The disk drives and their associated computer-readable media provide non-volatile storage of computer-readable instructions, data structures, program modules, and other data for computing device 802. Although the example illustrates a hard disk 816, a removable magnetic disk 820, and a removable optical disk 824, it is to be appreciated that other types of computer-readable media which 20 can store data that is accessible by a computer, such as magnetic cassettes or other magnetic storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical storage, random access memories (RAM), read only memories (ROM), electrically erasable programmable read-only memory (EEPROM), and the like, can also be utilized to implement the exemplary 25 computing system and environment.

Any number of program modules can be stored on the hard disk 816, magnetic disk 820, optical disk 824, ROM 812, and/or RAM 810, including by way of example, an operating system 826, one or more application programs 828, other program modules 830, and program data 832. Each of such operating system 826, 5 one or more application programs 828, other program modules 830, and program data 832 (or some combination thereof) may include an embodiment of the systems and methods for a test instantiation system.

Computing device 802 can include a variety of computer-readable media identified as communication media. Communication media typically embodies 10 computer-readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" refers to a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not 15 limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared, and other wireless media. Combinations of any of the above are also included within the scope of computer-readable media.

A user can enter commands and information into computing device 802 via 20 input devices such as a keyboard 834 and a pointing device 836 (e.g., a "mouse"). Other input devices 838 (not shown specifically) may include a microphone, joystick, game pad, controller, satellite dish, serial port, scanner, and/or the like. These and other input devices are connected to the processing unit 804 via input/output interfaces 840 that are coupled to the system bus 808, but may be 25 connected by other interface and bus structures, such as a parallel port, game port, and/or a universal serial bus (USB).

A monitor 842 or other type of display device can also be connected to the system bus 808 via an interface, such as a video adapter 844. In addition to the monitor 842, other output peripheral devices can include components such as speakers (not shown) and a printer 846 which can be connected to computing  
5 device 802 via the input/output interfaces 840.

Computing device 802 can operate in a networked environment using logical connections to one or more remote computers, such as a remote computing device 848. By way of example, the remote computing device 848 can be a personal computer, portable computer, a server, a router, a network computer, a peer device  
10 or other common network node, and the like. The remote computing device 848 is illustrated as a portable computer that can include many or all of the elements and features described herein relative to computing device 802.

Logical connections between computing device 802 and the remote computer 848 are depicted as a local area network (LAN) 850 and a general wide  
15 area network (WAN) 852. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet. When implemented in a LAN networking environment, the computing device 802 is connected to a local network 850 via a network interface or adapter 854. When implemented in a WAN networking environment, the computing device 802  
20 typically includes a modem 856 or other means for establishing communications over the wide network 852. The modem 856, which can be internal or external to computing device 802, can be connected to the system bus 808 via the input/output interfaces 840 or other appropriate mechanisms. It is to be appreciated that the illustrated network connections are exemplary and that other means of establishing  
25 communication link(s) between the computing devices 802 and 848 can be employed.

In a networked environment, such as that illustrated with computing environment 800, program modules depicted relative to the computing device 802, or portions thereof, may be stored in a remote memory storage device. By way of example, remote application programs 858 reside on a memory device of remote  
5 computing device 848. For purposes of illustration, application programs and other executable program components, such as the operating system, are illustrated herein as discrete blocks, although it is recognized that such programs and components reside at various times in different storage components of the computer system 802, and are executed by the data processor(s) of the computer.

10 Although the description above uses language that is specific to structural features and/or methodological acts, it is to be understood that the invention defined in the appended claims is not limited to the specific features or acts described. Rather, the specific features and acts are disclosed as exemplary forms of  
implementing the invention.

15 Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any form of suggestion that that prior art forms part of the common general knowledge in Australia.

2004201420 15 Sep 2009

- 20 -

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method comprising:  
identifying a plurality of devices in a network;  
5 tagging each identified device in the network with a discovery identifier,  
wherein the discovery identifier identifies a particular discovery process,  
identifying connection information associated with the plurality of devices;  
converting the connection information associated with the plurality of  
10 devices into a hierarchical structure;  
identifying communication links between the devices; and  
storing network topology data associated with the network.
  
2. A method comprising:  
creating a discovery identifier to identify a particular discovery process,  
15 tagging each identified device in the network with the discovery identifier,  
identifying a plurality of devices in a network;  
identifying connection information associated with the plurality of devices;  
identifying spanning tree information associated with the plurality of  
20 devices;  
identifying communication links between the devices using the spanning tree  
information; and storing network topology data associated with the network.
  
3. A method as recited in claim 1 or claim 2, wherein storing network topology  
25 data includes storing network topology data in at least one system link table.
  
4. A method as recited in claim 1 or claim 2, wherein storing network topology  
data includes storing device identities, connection identities and links  
between devices.

2004201420 15 Sep 2009

- 21 -

5. A method as recited in claim 1 or claim 2, further comprising generating a representation of the network topology.
6. A method as recited in claim 1 or claim 2, wherein the devices include network routing devices.
7. A method as recited in claim 1 or claim 2, wherein the devices include computing systems.
8. A method as recited in claim 1 or claim 2, wherein identifying links between the hardware devices includes:  
identifying a range of IP addresses to analyze;  
querying each IP address in the range of IP addresses; and  
if a response to the query is received, requesting management information base data from the responding device.
9. A method as recited in claim 8 further comprising retrieving additional information about the responding device if a response to the query is received.
10. A method as recited in claim 1 wherein converting the connection information associated with the plurality of devices into a hierarchical structure includes generating a spanning tree associated with the plurality of devices.
11. One or more computer-readable storage media containing a computer program that is executable by a processor to perform the method recited in claim 1 or claim 2.

2004201420 15 Sep 2009

- 22 -

12. A method comprising:  
creating a discovery identifier to identify a particular discovery process,  
tagging each identified device in the network with the discovery identifier,  
identifying a range of IP addresses to discover;  
5 querying each IP address in the range of IP addresses;  
if a response to the query is received:  
    requesting management information base data from a responding  
device;  
    identifying a device type associated with the responding device; and  
10 retrieving additional information from the responding device based on  
the device type associated with the responding device wherein the  
management information base data and the additional information is stored  
in a system link table.
- 15 13. A method as recited in claim 12 wherein the identified device types include  
computing devices and bridges.
14. A method comprising:  
initiating a device discovery process associated with a network;  
20 creating a discovery identifier associated with the device discovery process;  
identifying a plurality of devices in a network;  
tagging the devices with the discovery identifier;  
maintaining a status of each of the plurality of devices;  
repeatedly querying each of the plurality of devices;  
25 recording responses received from each of the plurality of devices;  
if no response is received from a particular device within a predetermined  
period of time:  
    changing the status of the particular device to inactive; and

2004201420 15 Sep 2009

- 23 -

notifying other devices in the network of the changed status of the particular device.

- 5
15. A method as recited in claim 14 wherein repeatedly querying each of the plurality of devices in the network includes querying each of the plurality of devices in the network at periodic intervals.
- 10
16. A method as recited in claim 14 wherein the predetermined period of time is on the order of one hour.
- 15
17. A method as recited in claim 14 further comprising notifying other application programs of the changed status of the particular device.
18. A method as recited in claim 14 further comprising:  
15 identifying new communication links in the network; and  
tagging the new communication links with the discovery identifier.
- 20
19. A method as recited in claim 14 wherein tagging the new devices with the discovery identifier includes storing the discovery identifier with other network topology data in a storage device.
20. A method as recited in claim 14 wherein each discovery identifier is unique with respect to other discovery identifiers.
- 25
21. An apparatus comprising:  
means for identifying a plurality of devices in a network;

2004201420 15 Sep 2009

- 24 -

means for identifying connection information associated with the plurality of devices;

means for representing the connection information associated with the plurality of devices in a hierarchical structure;

5 means for tagging each of the plurality of devices with a discovery identifier, wherein the discovery identifier identifies a particular discovery process; and means for storing network topology data associated with the network.

10 22. An apparatus as recited in claim 21 wherein the plurality of devices in the network are identified by pinging a plurality of IP addresses.

15 23. An apparatus comprising:  
a discovery engine coupled to a network, wherein the discovery engine is configured to create a discovery identifier to identify a particular discovery process,  
identify and tag with the discovery identifier a plurality of network devices in the network and identify connection information associated with the plurality of network devices, the discovery engine further configured to  
20 identify spanning tree information associated with the plurality of network devices and identify links between the network devices using the spanning tree information; and  
a storage device coupled to the discovery engine and configured to store network topology data associated with the network.

25 24. An apparatus as recited in claim 23 wherein the discovery engine is further configured to periodically update the status of network devices in the network.

2004201420 15 Sep 2009

- 25 -

25. An apparatus as recited in claim 23 further comprising a network monitor coupled to the discovery engine and configured to generate a display representing the network topology.
- 5 26. One or more computer-readable storage media having stored thereon a computer program that, when executed by one or more processors, causes the one or more processors to:
- 10 create a discovery identifier to identify a particular discovery process, identify a plurality of IP addresses to discover; query each IP address in the plurality of IP addresses; if a query response is received:
- 15 tag each identified device with the discovery identifier, request management information base data from the responding device; identify a device type associated with the responding device; and retrieve additional information from the responding device based on the device type associated with the responding device wherein the management information base data and the additional information is stored in a system link table.
- 20 27. One or more computer-readable storage media as recited in claim 26 wherein if no query response is received, the next IP address is queried.
28. One or more computer-readable storage media as recited in claim 26 wherein the device type is a bridge.
- 25 29. One or more computer-readable storage media as recited in claim 26 wherein the device type is a computing device.

2004201420 15 Sep 2009

- 26 -

30. One or more computer-readable storage media as recited in claim 26 wherein the plurality of IP addresses correspond to the IP addresses associated with a network environment.
- 5 31. A method comprising:  
initiating a device discovery process associated with a network;  
creating a discovery ID associated with the device discovery process;  
identifying new devices in the network by searching a range of IP addresses;  
and  
10 tagging the new devices with the discovery ID.
32. A method as recited in claim 31 further comprising:  
identifying new communication links in the network by requesting  
management information base data from the identified devices; and  
15 tagging the new communication links with the discovery ID.
33. A method as recited in claim 32, further comprising deriving a device type  
for each of the identified devices from their requested management base  
data, and retrieving additional information from the identified devices based  
20 on their device type.
34. A method as recited in claim 33, wherein the management information base  
data and the additional information is stored in a system link table.
- 25 35. A method as recited in claim 33, wherein the identified device types include  
computing devices and bridges.
36. A method as recited in claim 31, wherein tagging the new devices with the  
discovery ID includes storing the discovery ID with other network topology

2004201420 15 Sep 2009

- 27 -

data in a storage device.

37. A method as recited in claim 31, wherein each discovery ID is unique with respect to other discovery IDs.
- 5
38. A method as recited in claim 31, further comprising repeating the method, and if a particular device previously identified as a new device in the network cannot be identified when repeating the method, then changing a status of the particular device to inactive.
- 10
39. A method as recited in claim 38, further comprising notifying other devices in the network of the changed status of the particular device.
40. A method as recited in claim 39, further comprising generating a spanning tree associated with the identified devices.
- 15
41. A method as recited in claim 31, wherein searching a range of IP addresses comprises pinging each IP address included in the range of IP addresses.
- 20
42. An apparatus comprising:  
means for initiating a device discovery process associated with a network;  
means for creating a discovery ID associated with the device discovery process;  
means for identifying new devices in the network by searching a range of IP  
addresses; and  
25 means for tagging the new devices with the discovery ID.
43. An apparatus as recited in claim 42, further adapted to identify new communication links in the network by requesting management information

2004201420 15 Sep 2009

- 28 -

base data from the identified devices, and to tag the new communication links with the discovery ID.

- 5
44. An apparatus as recited in claim 42, further comprising means for generating a spanning tree associated with the identified devices.
45. An apparatus as recited in claim 42, further comprising means for storing the discovery ID of each tagged device with other network topology data.
- 10 46. One or more computer-readable media containing a computer program that is executable by a processor to perform the method recited in claim 31.
47. A method substantially as hereinbefore described, with reference to the accompanying drawings.
- 15 48. An apparatus substantially as hereinbefore described, with reference to the accompanying drawings.




Fig. 1

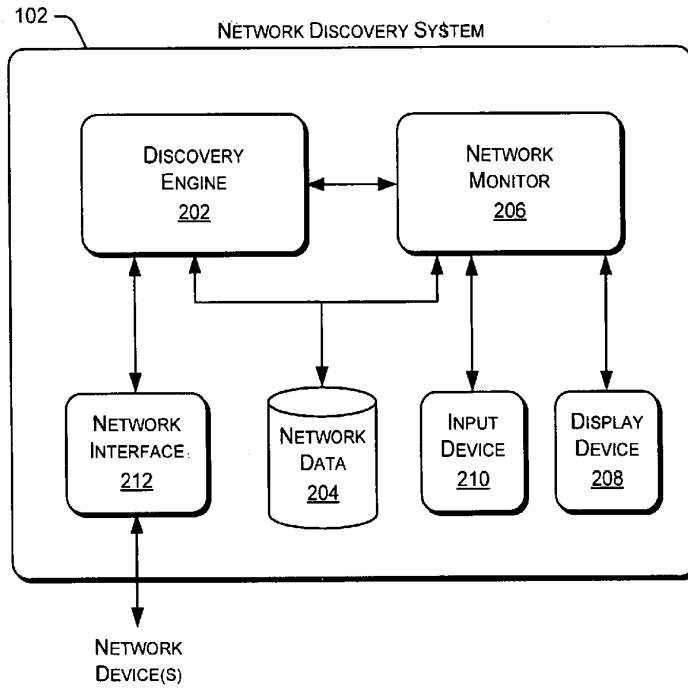
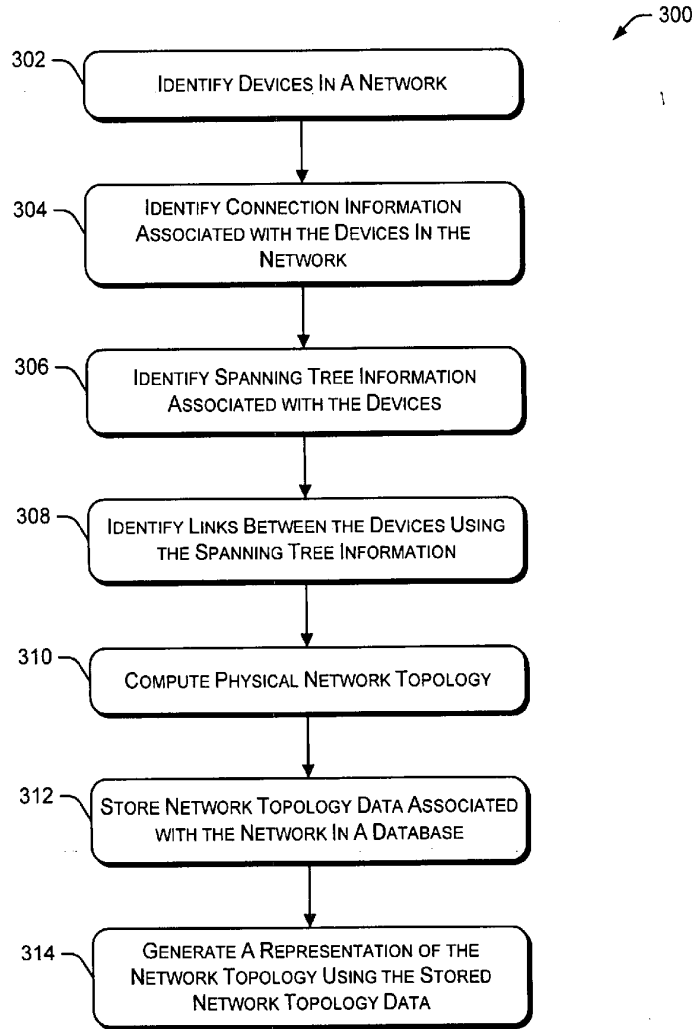


Fig. 2



*Fig. 3*

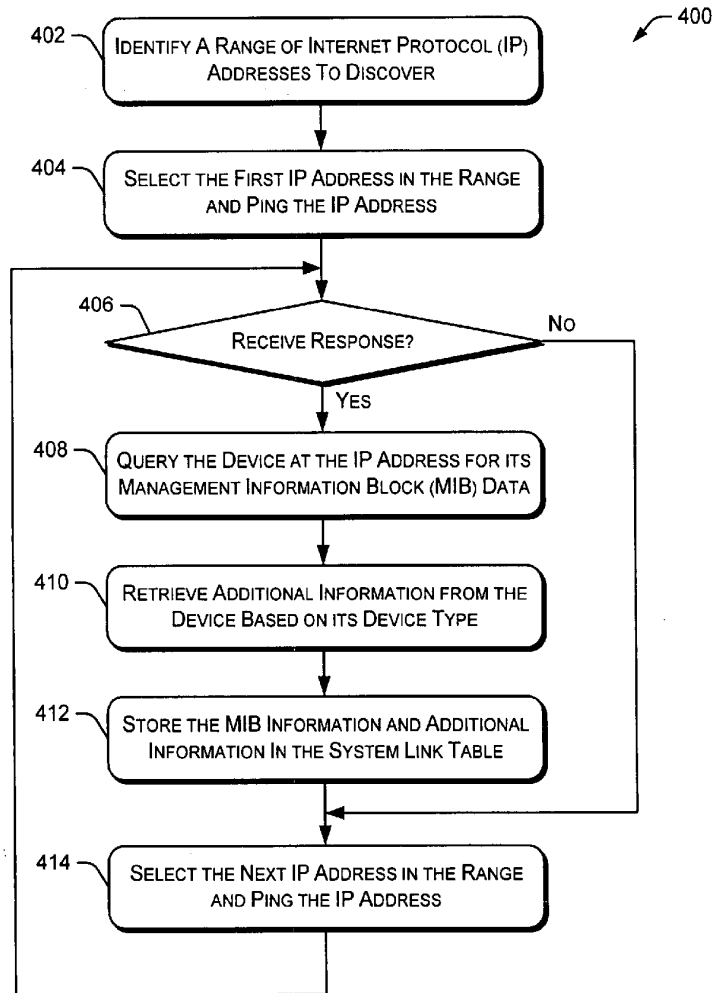


Fig. 4

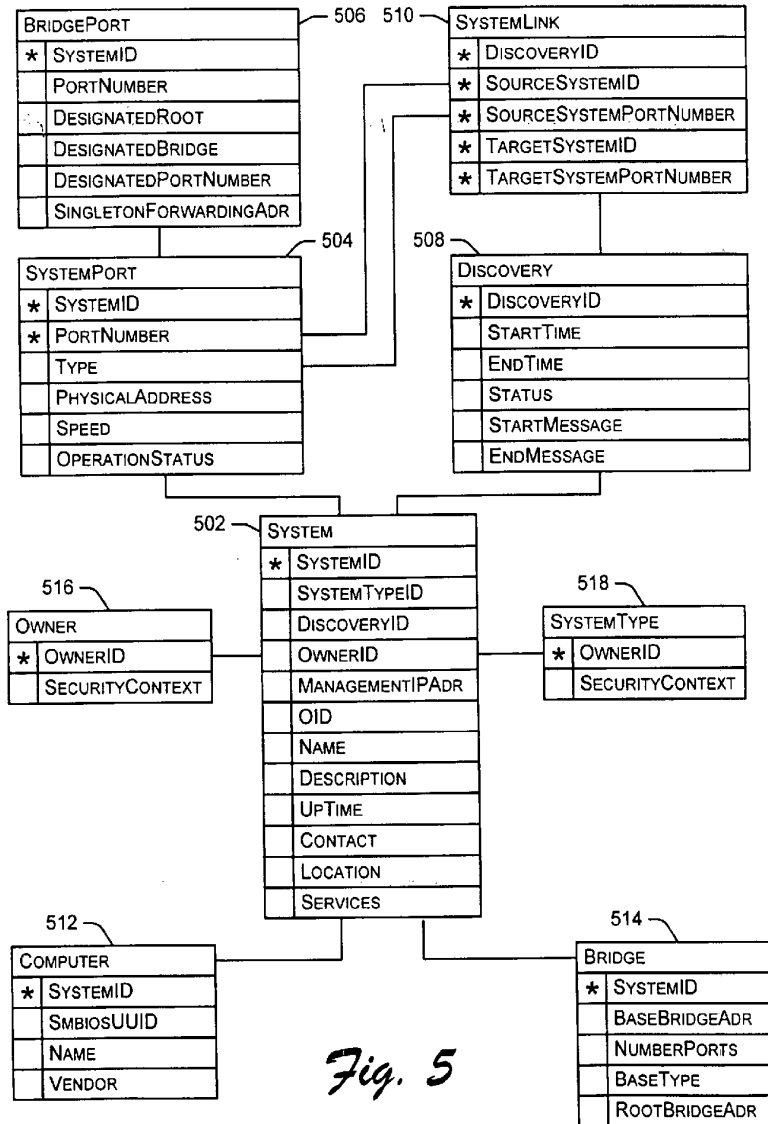


Fig. 5

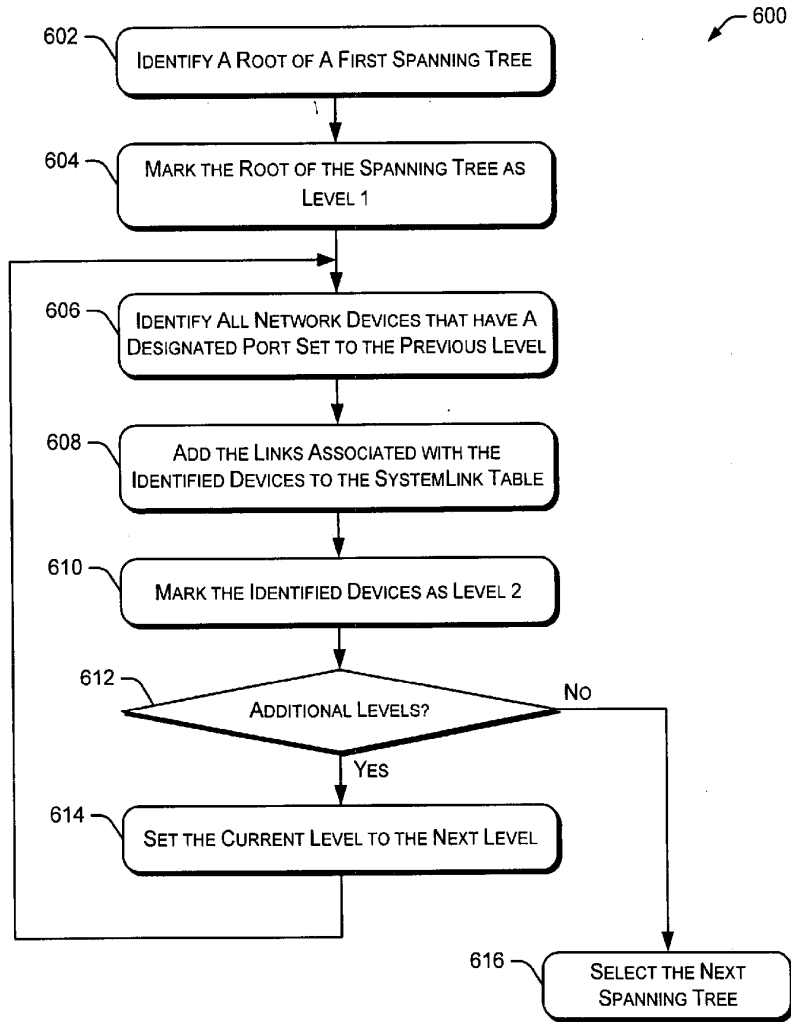


Fig. 6

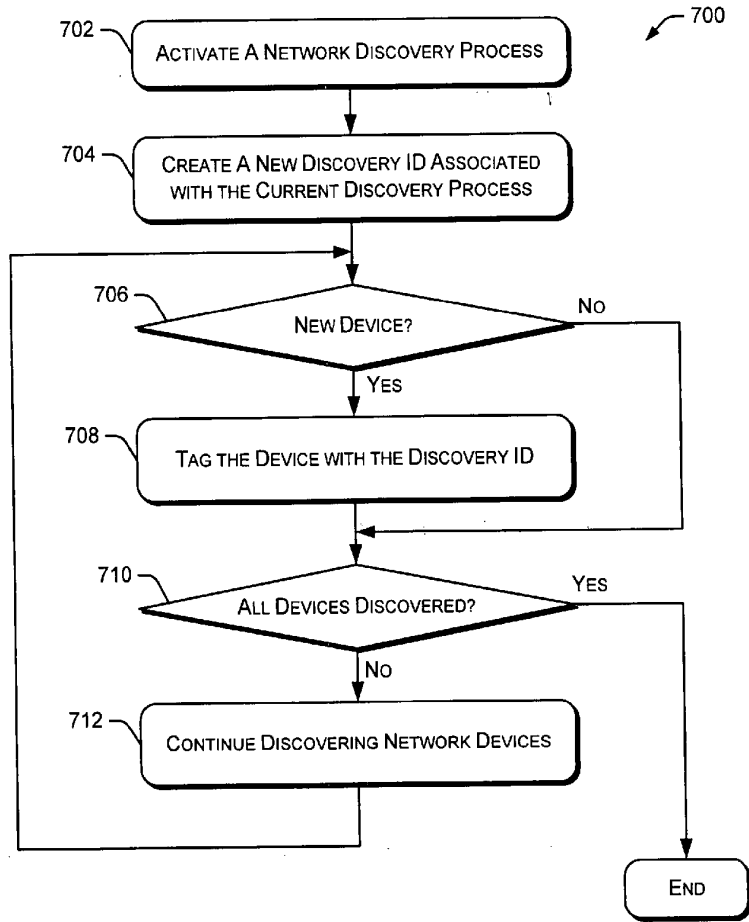


Fig. 7




Fig. 8