



(19) **United States**

(12) **Patent Application Publication**
STASSINOPOULOS et al.

(10) **Pub. No.: US 2014/0036688 A1**

(43) **Pub. Date: Feb. 6, 2014**

(54) **METHOD AND SYSTEM FOR MANAGING A DISTRIBUTED NETWORK OF NETWORK MONITORING DEVICES**

Publication Classification

(71) Applicants: **Dimitris STASSINOPOULOS**, Sunnyvale, CA (US); **George Zioulas**, Sunnyvale, CA (US); **Han C. Wen**, San Jose, CA (US)

(51) **Int. Cl.**
H04L 12/26 (2006.01)
(52) **U.S. Cl.**
CPC *H04L 43/024* (2013.01); *H04L 43/0876* (2013.01); *H04L 43/062* (2013.01)
USPC **370/241**

(72) Inventors: **Dimitris STASSINOPOULOS**, Sunnyvale, CA (US); **George Zioulas**, Sunnyvale, CA (US); **Han C. Wen**, San Jose, CA (US)

(57) **ABSTRACT**

(73) Assignee: **Riverbed Technology, Inc.**, San Francisco, CA (US)


Network traffic information for nodes of a first logical hierarchy is stored at a monitoring device according to ranks of the nodes within the logical hierarchy as determined by each node's position therein and user preferences. At least some of the network traffic information stored at the network monitoring device may then be reported to another network monitoring device, where it can be aggregated with similar information from other network monitoring devices. Such reporting may occur according to rankings of inter-node communication links between nodes of different logical hierarchies of monitored nodes.

(21) Appl. No.: **14/051,813**

(22) Filed: **Oct. 11, 2013**

Related U.S. Application Data

(63) Continuation of application No. 11/092,226, filed on Mar. 28, 2005, now Pat. No. 8,589,530.








Fig. 1



METHOD AND SYSTEM FOR MANAGING A DISTRIBUTED NETWORK OF NETWORK MONITORING DEVICES

[0001] This application is a Continuation of U.S. patent application Ser. No. 11/092,226, filed 28 Mar. 2005.

FIELD OF THE INVENTION

[0002] The present invention relates to: (a) the management of data stored by individual network monitoring devices, for example where such network monitoring devices are configured to store network traffic information relating to logical groupings of network nodes, and (b) managing data stored by such network monitoring devices when arranged in a distributed network of their own (i.e., a network monitoring device network).

BACKGROUND

[0003] Today, information technology professionals often encounter a myriad of different problems and challenges during the operation of a computer network or network of networks. For example, these individuals must often cope with network device failures and/or software application errors brought about by such things as configuration errors or other causes. Tracking down the sources of such problems often involves analyzing network and device data collected by monitoring units deployed at various locations throughout the network.

[0004] Traditional network monitoring solutions group network traffic according to whether a network node is a "client" or a "server". More advanced processes, such as those described in co-pending patent application Ser. No. 10/937,986, filed Sep. 10, 2004, assigned to the assignee of the present invention and incorporated herein by reference, allow for grouping data by the role being played by a network node and/or by logical units (business units) constructed by network operators for the purpose of monitoring and diagnosing network problems. These forms of advanced monitoring techniques can yield very good results in terms of providing operators with information needed to quickly diagnose and/or solve problems.

[0005] With these advanced forms of network monitoring, however, come problems. For example, collecting and storing data for all logical groupings of nodes and inter-nodal communications paths in a network quickly becomes unmanageable as that network grows in size. Consequently, what are needed are methods and systems to facilitate centralized network monitoring for large, distributed networks.

SUMMARY OF THE INVENTION

[0006] In one embodiment of the present invention network traffic information for those of a first logical hierarchy of monitored network nodes which can be accommodated by a first network monitoring device is stored according to ranks of the monitored network nodes within the logical hierarchy as determined by a node's position therein and user preferences. At least some of the network traffic information stored at the first network monitoring device may then be reported from the first network monitoring device to a second network monitoring device of the network monitoring device network, e.g., acting as a centralized network monitoring device. For example, the second network monitoring device may receive that portion of the network traffic information stored at the

first network monitoring device according to rankings of inter-node communication links between nodes of the first logical hierarchy of monitored network nodes of the first network monitoring device and others nodes of a second logical hierarchy of monitored network nodes of a third network monitoring device of the network monitoring device network. Such rankings of inter-node communication links may be determined according to ranks of individual nodes associated with the communication links within corresponding ones of the first and second logical hierarchies of nodes, each such rank being determined according to a first distance measured from a root node of a hierarchy under consideration to a node under consideration, a second distance measured from a leaf node of the hierarchy under consideration to the node under consideration and user preferences. Also, the ranks of the monitored network nodes within the first logical hierarchy of nodes of the first network monitoring device may be determined according to a first distance measured from a root node of the hierarchy to a node under consideration and a second distance measured from a leaf node of the hierarchy to the node under consideration and user preferences.

[0007] In further embodiments of the present invention, nodes of a grouping of nodes within a network are ranked, at a first network monitoring device, according to each node's position within a logical hierarchy of the nodes of the grouping and user preferences; and network traffic data associated with the nodes of the grouping of nodes is stored or not stored according to each node's rank as so determined. Thereafter, at least some of the network traffic data stored according to each node's rank may be transferred from the first network monitoring device to a second network monitoring device, for example if said rank satisfies additional ranking criteria concerning communications between nodes of different groupings.

[0008] Yet another embodiment of the present invention allows for aggregating, at a network monitoring device, network traffic information for inter-node communications between nodes of different logical groupings of nodes, said logical groupings of nodes including groupings defined in terms of other logical groupings of nodes, according to ranks of individual nodes within each of the different logical groupings associated with the inter-node communications, each such rank being determined according to a first distance measured from a root node of a logical hierarchy of a one of the logical groupings of nodes under consideration to a node thereof under consideration, a second distance measured from a leaf node of the logical hierarchy of the one of the logical groupings under consideration to the node under consideration and user preferences. Such aggregating may proceed incrementally for each branch of a logical group-to-logical group hierarchy constructed by the network monitoring device.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The present invention is illustrated by way of example, and not limitation, in the figures of the accompanying drawings in which:

[0010] FIG. 1 illustrates an example of a computer network and its associated network monitoring device;

[0011] FIG. 2 illustrates an example of a network of network monitoring devices deployed in accordance with an embodiment of the present invention; and

[0012] FIG. 3 illustrates an example of a BGO hierarchy in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0013] Described herein are methods and systems to facilitate network monitoring for computer networks. The present invention encompasses both the management of data stored by individual network monitoring devices, for example where such network monitoring devices are configured to store network traffic information relating to logical groupings of network nodes, and data stored by such network monitoring devices when arranged in a distributed network of their own (i.e., a network monitoring device network). First to be described will be the management of data stored by individual network monitoring devices. Thereafter, techniques for aggregating and managing the storage of such data among a network of monitoring devices will be presented.

[0014] For the first case, managing data stored by individual monitoring devices, consider that for large, distributed networks there may exist many (potentially hundreds or even thousands) of node-to-node communication links. Here we refer not strictly to the physical inter-node communication links (e.g., optical fibers, copper wires, and the like), but rather to the virtual or logical node-to-node connections that permit communication between two or more nodes in one or more computer networks. Because of constraints on the amount of physical memory available to a network monitoring device, it becomes at the very least impractical (and quickly impossible) to collect and store network traffic data for all of these multiple inter-node communication links for a network of any appreciable size. Indeed, the situation becomes even worse (from a need for storage viewpoint) if the nodes are grouped in some fashion, for now one must consider not only the individual node-to-node communications but also the group-to-group communications, which themselves may exist at multiple levels. Consequently, decisions about what data and/or which nodes/communication links should be monitored must be made (to meet capacity limits of the network monitoring devices); all the while remembering that network operators will still require sufficient information regarding network traffic conditions in order to make informed decisions regarding network operations and control.

[0015] In a first aspect of the present invention, these needs are addressed by a methodology for determining which nodes/links (i.e., the network traffic data associated with such monitored nodes and/or links) to track in one or a set of monitoring devices, to ensure data integrity. In this procedure, each network monitoring device collects data for designated nodes/communication links in a computer network or network of networks. Where necessary, the nodes/links are ranked and decisions are made based on such rankings if it is necessary to discard data relating to any monitored nodes/links in order not to exceed storage and/or processing capacity of a network monitoring device. As will be more fully discussed below, in one embodiment such ranking is a function of an individual node's distance from a root and/or leaf position within a logical hierarchy describing the arrangement of the nodes of the subject network as well as other factors (e.g., user preferences).

[0016] Then, for the second case of managing the aggregation and storage of monitored data among a network of monitoring devices, we introduce the concept of "Appliances" and a "Director". As used herein, the term Appliance will be applied to those network monitoring devices assigned to collect network traffic data from designated nodes/links of one or more networks. The Director will be a central network moni-

toring device to which the Appliances send specified information concerning designated ones of the monitored nodes/links. Together, the Director and the Appliances form a network of network monitoring devices.

[0017] Just as the individual network monitoring devices (the Appliances) were limited in their ability to store network traffic data concerning all of the myriad inter-node communication links, so too is the Director limited in its ability to store network traffic data received from the Appliances. Hence, the present invention further encompasses techniques for making decisions about which data concerning the monitored nodes/links to pass from the Appliances to the Director. As was the case for the individual Appliances, such decisions involve rankings of nodes/links. In this way, network operators using the Director monitoring device may readily gain access to network diagnostic information at a single monitoring device while at the same time that monitoring device is not overwhelmed with information concerning the numerous network nodes and communication links.

[0018] As will become apparent, the ability to group various network nodes/links into logical units and to further group these logical units into higher layer units provides for many of the advantages of the present methods and will be discussed before presenting details of the various ranking algorithms used in connection with the present invention. Before doing so, however, it is important to remember that for purposes of explanation numerous specific details are set forth herein in order to provide a thorough understanding of the invention. However, it will be appreciated by one with ordinary skill in the art that these specific details need not be used to practice the present invention. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0019] The methods described herein may be used in conjunction with other techniques to allow network operators to detect problems and/or discover relevant information with respect to network application usage/performance and then isolate the problem/information to specific contributors (e.g., users, applications or network resources). More particularly, the present methods involve computations and analyses regarding many variables and are best performed or embodied as computer-implemented processes or methods (a.k.a. computer programs or routines) that may be rendered in any computer programming language including, without limitation, C#, C/C++, Fortran, COBOL, PASCAL, assembly language, markup languages (e.g., HTML, SGML, XML, VoXML), and the like, as well as object-oriented languages/environments such as the Common Object Request Broker Architecture (CORBA), Java™ and the like. In general, however, all of the aforementioned terms as used herein are meant to encompass any series of logical steps performed in a sequence to accomplish a given purpose.

[0020] In view of the above, it should be appreciated that some portions of this detailed description of the present invention are presented in terms of algorithms and symbolic representations of operations on data within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the computer science arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not

necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers or the like. It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise, it will be appreciated that throughout the description of the present invention, use of terms such as “processing”, “computing”, “calculating”, “determining”, “displaying” or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system’s registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0021] The present invention can be implemented with an apparatus to perform the operations described herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general-purpose computer, selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0022] The algorithms and processes presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method. For example, any of the methods according to the present invention can be implemented in hard-wired circuitry, by programming a general-purpose processor or by any combination of hardware and software. One of ordinary skill in the art will immediately appreciate that the invention can be practiced with computer system configurations other than those described below, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, DSP devices, network PCs, minicomputers, mainframe computers, and the like. The invention can also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. The required structure for a variety of these systems will appear from the description below.

[0023] The methods of the present invention may be implemented using computer software. If written in a programming language conforming to a recognized standard, sequences of instructions designed to implement the methods can be compiled for execution on a variety of hardware platforms and for interface to a variety of operating systems. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein. Furthermore, it is common in the art to speak of software, in one form

or another (e.g., program, procedure, application, etc.), as taking an action or causing a result. Such expressions are merely a shorthand way of saying that execution of the software by a computer causes the processor of the computer to perform an action or produce a result.

[0024] As indicated above, before describing the various ranking processes used as a basis for determining which information to store at a network monitoring device, the manner in which such information is collected will first be discussed. Turning then to FIG. 1, a computer network including multiple logical groupings (e.g., BG1, BG2) of network nodes is illustrated. Logical groupings such as BG1 and BG2 may be defined at any level. For example, they may mirror business groups, or may designate computers (or other nodes, e.g., printers, servers, image processors, scanners, or other computer equipment generally addressable within a computer network) performing similar functions, computers located within the same building, or any other aspect which a user or network operator/manager wishes to highlight. FIG. 1 shows one simple organization of a small number of computers and other network nodes, but those familiar with computer network operations/management will appreciate that the number of computers and network nodes may be significantly larger as can the number of connections (communication links) between them. Modem network configurations are mutable and complex, which is one of the reasons why the present invention is useful. Information representing the total utilization of all nodes in particular directions or activities provides much greater visibility into overall network traffic than does a large collection of individualized node information. The present invention allows for the grouping of network traffic into logical groups and groups of logical groups that a user can configure in order to allow visibility of network traffic at various hierarchical levels.

[0025] In FIG. 1 lines between nodes and other entities are meant to indicate network communication links, which may be any mode of establishing a connection between nodes including wired and/or wireless connections. Moreover, a firewall (shown as the dashed line) surrounds a geographic collection of networked nodes and separates components of an internal network from an external network 6. A network traffic monitoring device 8 is shown at the firewall. However, the network traffic monitoring device 8 may be located within the internal network, or the external network 6 or anywhere that allows for the collection of network traffic information. Moreover, network traffic monitoring device 8 need not be “inline.” That is, traffic need not necessarily pass through network traffic monitoring device 8 in order to pass from one network node to another. The network traffic monitoring device 8 can be a passive monitoring device, e.g., spanning a switch or router, whereby all the traffic is copied to a switch span port which passes traffic to network traffic monitoring device 8.

[0026] In the example shown in FIG. 1, BG1 contains several internal network nodes N101, N102, N103, and N104 and external nodes N105, N106 and N107. Similarly, BG2 contains several internal network nodes N201, N202, N203, N204, N205, N206, and external nodes N207, N208, N209, N210 and N211. A network node may be any computer or device on the network that communicates with other computers or devices on the network. Each node may function as a client, server, or both. For example, node N103, is shown as a database which is connected to Node N104, a web application server, via a network link 10. In this configuration, it is typical

for node N104 to function as a client of node N103 by requesting database results. However N104 is also depicted as connected to the external network 6 via network link 12. In this configuration, it is typical for N104 to function as a server, which returns results in response to requests from the external network. Similarly, database node N103, which functions as a server to N104, is shown connected to node N107 via a network link 14. N107 may upload information to the database via link 14, whereby N107 is functioning as a server and N103 is functioning as a client. However, N107 is also shown connected to the external network 6 via link 16. This link could indicate that N107 is browsing the Internet and functioning as a client.

[0027] Furthermore, network nodes need not be within the internal network in order to belong to a logical group. For example, traveling employees may connect to the logical group network via a virtual private network (VPN) or via ordinary network transport protocols through an external network such as the Internet. As shown in FIG. 1, network nodes N105, N106, and N107 belong to logical group BG1, but are outside the firewall, and may be geographically distant from the other network nodes in BG1. Similarly, network nodes N207, N208, N209, N210, and N211 are members of logical group BG2, but are physically removed from the other members of group BG2. It is important to note that the firewall in this configuration is for illustrative purposes only and is not a required element in networks where the present invention may be practiced. The separation between internal and external nodes of a network may be formed by geographic distance (as described above), or by networking paths (that may be disparate or require many hops for the nodes to connect to one another regardless of their geographic proximity).

[0028] For a relatively small network such as that shown in FIG. 1, a single network monitoring device 8 may suffice to collect and store network traffic data for all nodes and communication links of interest. However, for a network of any appreciable size (or for a network of networks), this will likely not be the case. Thus decisions about what data to store for which nodes/groups of nodes and/or links/groups of links need to be made.

[0029] To further illustrate this point, consider a network monitoring device located in a data center (call it the New York (NY) datacenter) that monitors traffic between the New York (NY) office of an enterprise and its remote branch offices around the world. The NY enterprise and each of the branch offices may be organized with multiple logical groups of nodes. We will call each such logical group a “business group” or BG, however, it should be recognized that BGs could be created along any of the lines discussed above (e.g., any user-desired definition). Indeed, some of the BGs may themselves include other BGs, forming what will be termed herein business group organizations or BGOs. See, for example, FIG. 3 in which a BGO called “California” (CA) includes multiple BGs: “San Francisco” (SF), “Los Angeles” (LA), and “San Diego” (SD), each of which may themselves be made up of other BGs (e.g., the Los Angeles BG may include BGs for Santa Monica (SM), Riverside (R) and Orange County (OC)) and/or nodes. Thus, the various BGs may be grouped in various hierarchies, within which there may be many node-to-node and group-to-group (BG-to-BG and/or BGO-to-BGO) connections (representing inter-group communications).

[0030] In addition to the above, for each communication link under consideration there are a host of various metrics

that might be collected by a network monitoring device. Among these are: Goodput, Payload, Throughput, Transaction Throughput, Packet Loss, Retransmission Delay, Retransmission Rate and Round Trip Time, Application Response Rate, Application Response Time, Client Reset Rate, Connection Duration, Connection Established Rate, Connection Request Rate, Connection Setup Time, Connections Failed Rate, Data Transfer Time, Server Reset Rate and Time to First Byte. These metrics can be further subdivided on the basis of the role being played by the content originator and the content requester. Thus, as this exercise should make clear, for networks (or networks of networks) of any appreciable size there are far too many data points for a single monitoring device to cope with. That is, a single device cannot reasonably store data concerning all of the various communication links within such a network (e.g., due to limits on physical storage devices, bandwidth utilization, etc.) and so decisions about what data to store and what data not to store at the monitoring device need to be made.

[0031] The solution provided by the present invention allows for such decision-making. In one embodiment of the invention, a network monitoring device consults user-supplied definitions of the BGs/BGOs for which it is responsible and “builds” a BG/BGO hierarchy. The definitions of the BGs/BGOs may be stored locally at a network monitoring device or may be stored in a central location to which a network monitoring device has access. These definitions comprise configuration files that include user-supplied instructions regarding the BGs and BGOs to be monitored and will, generally, define the types of statistics or metrics for which data is to be collected and the organization of the network nodes. The precise details of such instructions are not critical to the present invention and in some cases such instructions may be provided by manually configuring a network and its associated monitoring devices on a port-by-port level. What is important is that the network monitoring device has some means of determining which nodes it is responsible for monitoring.

[0032] With the BGO information made available, a network monitoring device constructs its relevant BGO hierarchy. In doing so, the network monitoring device considers only those links which are active; that is, those links which have active communications taking place. The BGO hierarchy may be regarded as a “tree-like” structure having a root node, branches and leaf nodes. The “root” node represents the highest hierarchical level in the BGO, while the “leaf” nodes represent the lowest such hierarchical levels (e.g., individual computer resources). “Branch” nodes may be nodes interconnecting leaf nodes to the root node and there may be any number of such branch nodes (including none) between the root node and any of the leaf nodes. A branch hierarchy is constructed by combining the network data collected for each of the leaf nodes within a branch and storing that combination with reference to the common branching node from which the leaf nodes depend. For each branch of the BGO hierarchy, and on a branch-by-branch basis (starting with the leaf nodes thereof), decisions are made about whether or not to store the monitored data for those nodes/links

[0033] It should be apparent that in the process of constructing such a hierarchy, where each higher layer includes combined statistics from lower layers, for a hierarchy of any significant depth it may not be possible to store all of the raw data and the combinations thereof for every level of the hierarchical tree. Stated differently, storage and/or processing

capabilities of the network monitoring device may demand that some of the data concerning some of the leaf nodes and/or branching nodes of the BGO hierarchy be intentionally dropped.

[0034] To accommodate this reality, the present invention provides for ranking and pruning the BGO hierarchy as it is being constructed by the network monitoring device. Importantly, this ranking and pruning process can be performed without the need for the network monitoring device to store data for each node of the entire BGO structure. Thus, the BGO hierarchy can be constructed “on-the-fly”, with each branch being pruned as needed during that process so as to accommodate the network monitoring device’s storage and/or processing limitations.

[0035] The ranking algorithm used by the network monitoring device as it constructs each branch of the BGO hierarchy may be any such process as permits the above-described operations. In one embodiment, the algorithm used is:

$$R_{composite} = F_{devices}(a) + F_{rank}(r) + F_{depth}(d) \quad (1)$$

where $R_{composite}$ is a composite ranking of the node/link under consideration.

[0036] In equation (1), $F_{devices}(a)$ is a constant that is proportional to a , the number of monitoring devices designated to have their data for this associated BG or BGO (henceforth referred to as a “node”) aggregated onto a central monitoring device (the Director) as discussed in further detail below. This $F_{devices}(a)$ constant is intended to give prioritization of the highest ranking to those nodes that need to have their data aggregated to the central monitoring device. Put differently, the $F_{devices}(a)$ factor ensures that nodes for which there is at least one monitoring device contributing to the group (i.e., $a > 0$) receive the highest ranking. In one embodiment of the present invention, $F_{devices}(a)$ is a monotonically increasing function of a , to ensure that preference is given to nodes with higher a values.

[0037] The $F_{rank}(r)$ term is a function whose value monotonically decreases with increasing values of r , the distance (measured in the number of nodes) between the associated root or “top” node of the hierarchy and the node associated with $R_{composite}$ (e.g., the number of hops within the BGO hierarchical tree). The $F_{rank}(r)$ term is intended to give preference to nodes that are higher in the BGO tree hierarchy.

[0038] The $F_{depth}(d)$ term is a function whose value monotonically increases with increasing values of d , the distance (measured in the number of nodes) between the leaf or “bottom” node of the hierarchy and the node associated with $R_{composite}$. The $F_{depth}(d)$ term is intended for scenarios where there are nodes with the same values for $F_{devices}(a)$ and $F_{rank}(r)$, to give preference to those nodes that have “deeper” tree hierarchies, as reflected by the value of d . In one embodiment of the present invention the relative magnitudes of the three terms $F_{devices}(a)$, $F_{rank}(r)$ and $F_{depth}(d)$ may be expressed as $F_{devices}(a) \gg F_{rank}(r) \gg F_{depth}(d)$, for expected values of $a = (0 \text{ to } \sim 100)$, “ r ” = (1 to ~ 20) and “ d ” = (1 to ~ 20).

[0039] The rank ($R_{composite}$) of each node/link is recorded in a database maintained by the network monitoring device. Thereafter, as each branch of the hierarchy is constructed, nodes/links may be pruned (i.e., decisions may be made to drop data collected for such nodes/links) according to such ranks and the storage/processing capacity of the network monitoring device. Alternatively, or in addition, decisions about pruning may be based on thresholds for the number of nodes for which to store data as configured by a user.

[0040] The foregoing has thus addressed the need to determine which nodes/links (i.e., the network traffic data associated with such monitored nodes and/or links) to track in an individual network monitoring device. The discussion now turns to the second aspect of the present invention: the case of managing the aggregation and storage of such monitored data among a network of monitoring devices. In this discussion we refer to different types of network monitoring devices, namely Appliances and a central Director.

[0041] Earlier it was noted that the network monitoring device **8** illustrated in FIG. **1** may be capable of storing all relevant network traffic information for a relatively small network. However, when the network became large, this was no longer true and so decisions had to be made about what data to store and what data not to store. Consider now the case where not only is the network (or network of networks) under consideration large, but also where more than a single network monitoring device is used.

[0042] Returning to the earlier example, such a situation may arise where, for example, in addition to the NY datacenter, an additional datacenter is located in California (CA). Just like NY, the datacenter in CA sends/receives traffic to/from the same set of remote branch offices distributed throughout the world. However, the monitoring device (call it Appliance 1) in NY does not “see” any of this data being transferred through the CA datacenter. That is, Appliance 1 does not capture the traffic to/from the CA datacenter. Therefore, a separate monitoring device (Appliance 2) is deployed in the CA datacenter to monitor traffic to/from that datacenter.

[0043] But now if a network operator wants to assess the total traffic between the London branch office and each of the NY and CA datacenters, then somehow the information collected by each of the Appliances must be aggregated. In accordance with the present invention, this aggregation is performed at a Director—a central network monitoring device. Collectively, the Director and the various Appliances make up a network of network monitoring devices and FIG. **2** illustrates an example thereof.

[0044] Within network **20**, central network monitoring device **22** receives and aggregates network traffic information from two individual network monitoring devices **24_a** and **24_b**. Monitoring device **24_a** is responsible for collecting network traffic data associated with a first network **26_a**. Monitoring device **24_b** is responsible for collecting network traffic data associated with a second network **26_b**. Networks **26_a** and **26_b** may each include multiple nodes, interconnected with one another and/or with nodes in the other respective network by a myriad of communication links, which may include direct communication links or indirect communication links (e.g., which traverse other networks not shown in this illustration). Thus, each of the network monitoring devices **24_a** and **24_b** may be responsible for collecting data concerning multiple groupings (logical and/or physical) of nodes in their associated networks **26_a** and **26_b**. That is, the network operator may, for convenience, define multiple logical and/or physical groupings of nodes in each of the networks **26_a** and **26_b** and configure the respective network monitoring devices **24_a** and **24_b** to store and track network traffic information accordingly. The total number of monitored nodes/links may be quite large.

[0045] Such a network of network monitoring devices poses several challenges. For example, if the network traffic information associated with the various BG/BGO-to-BG/BGO communications sought by the Director exceeds the

storage capacity of the Director, what information for which group-to-group communications should be kept? Also, in order not to overwhelm the available bandwidth within the network of network monitoring devices, how can the volume of information being sent between the Appliances and the Director be kept manageable? Finally, how can one ensure completeness (i.e., integrity) of the information for the various aggregations being performed? For example, if an operator wants all the traffic between London and the two datacenters aggregated, how can the operator be certain that each Appliance has stored traffic between its datacenter and London if each of the Appliances is pruning the number of nodes/links for which it stores traffic in accordance with the above-described processes? The present invention addresses these issues by employing a global ranking process somewhat similar to that discussed above with reference to a single network monitoring device.

[0046] Once the individual BGO hierarchies have been constructed by the Appliances, decisions about which data to transfer to the Director can be made. Because the Director will also have limits on the amount of data which it can store/process, a ranking algorithm, which may include a bias for ensuring that any nodes/links which the network operator has indicated should be tracked at this level are always included, for determining what data to store and what data not to store is used. One example of such a ranking algorithm used to select the links to be transferred to the Director is:

$$R'_{composite}(r_1, r_2) = F_{devices}(\max(a_1, a_2)) + F_{rank}(r_1, r_2) + F_{depth} \quad (2)$$

where “r”, “d” and “a” denote the same metrics as above and the subscripts 1 and 2 indicate the values associated with the different BGs/BGOs which the link under consideration interconnects. For example, in a CA-to-NY BGO-to-BGO example, subscript 1 might designate a node within the CA Appliance hierarchy and subscript 2 might designate a node within the NY Appliance hierarchy.

[0047] Thus, based on the rankings of the BGO-to-BGO hierarchical trees on the distributed edge monitoring devices (i.e., the Appliances), a central network monitoring device (the Director) can construct a composite BGO-to-BGO hierarchy encompassing the traffic seen by all the distributed edge monitoring devices. Indeed, this process can be repeated for multiple level network monitoring device hierarchies, which each monitoring device at successively higher layers of the hierarchy receiving data from lower layer devices and pruning BGO-to-BGO hierarchical trees accordingly.

[0048] Importantly, the ranking and pruning processes described herein may be implemented at network monitoring devices at any level within a network of network monitoring devices (e.g., one in which first layer Appliances report up to second layer Appliances, which in turn report up to higher layer Appliances, until finally reports are made to a Director). That is, network monitoring devices at any point within a network of such devices may employ such methods to keep network traffic information related to group-to-group communications bounded.

[0049] Thus, methods and systems to facilitate centralized network monitoring for distributed networks have been described. Although these methods and systems were discussed with reference to particular embodiments of the present invention, such embodiments should not be read as unnecessarily limiting the scope of the invention. Instead, the invention should only be measured in terms of the claims, which follow.

1. A network monitoring system comprising:
 - a plurality of network monitoring devices that monitor network traffic data from a plurality of nodes of a network, each network monitoring device being configured to collect network traffic data from an assigned subset of the nodes in the network, and
 - a central network monitoring device that is configured to receive at least a portion of the network traffic data collected by the network monitoring devices;
 - wherein at least one of the network monitoring devices is configured to select fewer nodes than its assigned subset of nodes for collecting network traffic data, based on a capacity of the network monitoring device and a priority associated with each node of its assigned subset of nodes.
2. The network monitoring system of claim 1, wherein the priority associated with at least one of the nodes is based on a number of network monitoring devices that provide network traffic data associated with this node to the central network monitoring device.
3. The network monitoring system of claim 1, wherein each subset of assigned nodes includes a root node, each of the nodes of the subset being hierarchically related to the root node, and the priority associated with each node is based on a hierarchical distance of the node from the root node.
4. The network monitoring system of claim 1, wherein each subset of assigned nodes includes leaf nodes and branch nodes arranged in a hierarchy, and the priority associated with each branch node is based on a hierarchical distance of the branch node from a hierarchically-closest leaf node.
5. The network monitoring system of claim 1, wherein the priority associated with each link between nodes is dependent upon a number of network monitoring devices that provide network traffic data associated with each node of the link to the central network monitoring device.
6. The network monitoring system of claim 1, wherein the central network monitoring device controls the portion of the network traffic data received from the network monitoring devices based on a capacity of the central network monitoring device and a priority associated with links between the nodes in the network.
7. The network monitoring system of claim 1, wherein each subset of assigned nodes includes a root node, each of the nodes of the subset being hierarchically related to the root node, and the priority associated with each link is based on a hierarchical distance of each node of the link from its root node.
8. The network monitoring system of claim 1, wherein each subset of assigned nodes includes leaf nodes and branch nodes arranged in a hierarchy, and the priority associated with each link is based on a hierarchical distance of each node of the link from a hierarchically-closest leaf node, the hierarchical distance of a leaf node from a hierarchically-closest leaf node being zero.
9. A method comprising:
 - assigning, via a central monitoring device, a subset of nodes in a network to each of a plurality of network monitoring devices that are configured to monitor network traffic data of the assigned subset of nodes,
 - selecting, at at least one network monitoring device, fewer nodes to monitor than its assigned subset of nodes, based on a capacity of the at least one network monitoring device and a priority associated with each node of its

assigned subset of nodes, and collecting network traffic data from the selected fewer nodes,
 receiving, at the central monitoring device, at least a portion of the network traffic data collected by the plurality of network monitoring devices, and
 reporting, by the central monitoring device, one or more statistics based on the received network traffic data.

10. The method of claim **9**, wherein the priority associated with at least one of the nodes is based on a number of network monitoring devices that provide network traffic data associated with this node to the central network monitoring device.

11. The method of claim **9**, wherein each subset of assigned nodes includes a root node, each of the nodes of the subset being hierarchically related to the root node, and the priority associated with each node is based on a hierarchical distance of the node from the root node.

12. The method of claim **9**, wherein each subset of assigned nodes includes leaf nodes and branch nodes arranged in a hierarchy, and the priority associated with each branch node is based on a hierarchical distance of the branch node from a hierarchically-closest leaf node.

13. The method of claim **9**, including selecting, by the central network monitoring device, the portion of the network traffic data to be received from the network monitoring devices based on a capacity of the central network monitoring device and a priority associated with links between the subsets of nodes in the network.

14. The method of claim **13**, wherein the priority associated with each link between subsets is dependent upon a number of network monitoring devices that provide network traffic data associated with each node of the link to the central network monitoring device.

15. The method of claim **13**, wherein each subset of assigned nodes includes a root node, each of the nodes of the subset being hierarchically related to the root node, and the priority associated with each link is based on a hierarchical distance of each node of the link from its root node.

16. The method of claim **13**, wherein each subset of assigned nodes includes leaf nodes and branch nodes arranged in a hierarchy, and the priority associated with each link is based on a hierarchical distance of each node of the link from a hierarchically-closest leaf node, the hierarchical distance of a leaf node from a hierarchically-closest leaf node being zero.

17. A non-transitory computer readable medium that includes a computer program that, when executed at a network monitoring device, causes the device to:

receive an assignment of a subset of nodes in a network to monitor for network traffic data,

select fewer nodes to monitor than the assigned subset of nodes, based on a capacity of the network monitoring device and a priority associated with each node of its assigned subset of nodes,

collect network traffic data from the selected fewer nodes, and

communicate at least a portion of the collected network traffic data to a central monitoring device.

18. The medium of claim **17**, wherein the priority associated with at least one of the nodes is based on a number of other network monitoring devices that provide network traffic data associated with this node to the central network monitoring device.

19. The medium of claim **17**, wherein the subset of assigned nodes includes a root node, each of the nodes of the subset being hierarchically related to the root node, and the priority associated with each node is based on a hierarchical distance of the node from the root node.

20. The medium of claim **17**, wherein the subset of assigned nodes includes leaf nodes and branch nodes arranged in a hierarchy, and the priority associated with each branch node is based on a hierarchical distance of the branch node from a hierarchically-closest leaf node.

21. A non-transitory computer readable medium that includes a computer program that, when executed at a central monitoring device, causes the device to:

assign a subset of nodes of a network to each of a plurality of network monitoring devices, each network monitoring device being configured to collect network traffic data from the subset of nodes, and

receive a portion of the network traffic data from the network monitoring devices based on a capacity of the central network monitoring device and a priority associated with links between the subsets of nodes.

22. The medium of claim **21**, wherein the priority associated with each link between subsets is dependent upon a number of network monitoring devices that provide network traffic data associated with each node of the subset to the central network monitoring device.

23. The medium of claim **21**, wherein each subset of assigned nodes includes a root node, each of the nodes of the subset being hierarchically related to the root node, and the priority associated with each link is based on a hierarchical distance of each node of the link from its root node.

24. The medium of claim **21**, wherein each subset of assigned nodes includes leaf nodes and branch nodes arranged in a hierarchy, and the priority associated with each link is based on a hierarchical distance of each node of the link from a hierarchically-closest leaf node, the hierarchical distance of a leaf node from a hierarchically-closest leaf node being zero.

* * * * *