

a destination IP address, a source port for TCP/UDP, a destination port for TCP/UDP, and an IP protocol. While flow pattern 130 may be represented by any data structure, an example of a data structure includes the following 5-tuple: <sourceip_address, destination_ip_address, source port, destination port, protocol_type>.

[0040] Flow pattern 130 may also include rules or other operations or tests to be performed on the data packets when evaluating whether the data packets match flow pattern 130. For example, as an alternative or addition to analyzing the source address and destination address of each of the data packets, a hash of each data packet may be calculated and compared to a hash included in flow pattern 130. In this manner, particular data packets may be selected, as further described below.

[0041] Upon receipt of flow pattern 130, virtual datacenter 302 stores flow pattern 130 in a memory area associated with virtual datacenter 302. At 406, virtual datacenter 302 distributes flow pattern 130 to one or more devices, applications, and/or other entities within virtual datacenter 302 that handle (e.g., route) data packets. In some embodiments, one or more of these entities manage a plurality of VMs as part of at least one virtual network. For example, virtual datacenter 302 may distribute flow pattern 130 to one or more of host components 308 illustrated in FIG. 3.

[0042] In some embodiments, virtual datacenter 302 transmits flow pattern 130 to each of host computing devices 100 and virtual gateways within virtual datacenter 302. In another example, only particular entities within virtual datacenter 302 are selectively targeted (e.g., a subset of the available host computing devices 100, applications, and/or virtual gateways). In some embodiments, a plurality of applications operates as, or otherwise implements some functions of, host computing device 100 and/or a virtual gateway.

[0043] At 408, over time, virtual datacenter 302 aggregates context data describing data packets that match the distributed flow pattern 130 from entities (e.g., host components 308) within virtual datacenter 302. For example, the context data is received from the entities in virtual datacenter 302 that handle the data packets, have received flow pattern 130, and have found data packets that match flow pattern 130. The context data may be aggregated across virtual networks within virtual datacenter 302 and/or across virtual datacenters 302. For example, the context data may be aggregated across one or more logical network boundaries. The context data relating to a single packet may be referred to as a footprint record.

[0044] At 410, virtual datacenter 302 determines a role of user 108. The role may be specified in, for example, the request received from user 108. The role may also be derived from past requests from user 108, feedback from user 108 resulting from past requests, default values, and/or other factors.

[0045] At 412, virtual datacenter 302 filters the aggregated context data based on the role of user 108. For example, the context data may be reduced to a level of detail associated with the determined role of user 108. At 414, virtual datacenter 302 provides the filtered context data to user 108. In some embodiments, virtual datacenter 302 displays the filtered context data to user 108.

[0046] FIG. 5 is a flowchart of an exemplary method 500 performed by host computing device 100, or other entity managing virtual machines (e.g., an edge appliance such as a virtual gateway at the edge of a virtual network to monitor

data packets and collect context data for selected data packets. For example, one or more of the operations illustrated in FIG. 5 may be performed by one or more of host components 308 illustrated in FIG. 3. However, while method 500 is described with reference to execution by host computing device 100 (shown in FIG. 2), it is contemplated that method 500 may be performed by any computing device, application, and/or other entity. For example, method 500 may be performed by a virtual gateway, hypervisor 210, VMs 235, applications 270, and/or guest operating system 265.

[0047] At 502, host computing device 100 checks whether a command or request to initiate flow monitoring has been received. For example, the command is received from management component 306 or other entity within virtual datacenter 302. Upon receipt of the command to initiate flow monitoring, host computing device 100 receives, obtains, or otherwise accesses flow pattern 130 associated with the command at 504. Host computing device 100 stores flow pattern 130 in memory area 104. For example, host computing device 100 caches flow pattern 130 into a hash table at the kernel level.

[0048] At 506, for each data packet handled by host computing device 100 (e.g., routed) within one more virtual networks, host computing device 100 compares the data packet to flow pattern 130. For example, host computing device 100 compares the source and destination addresses of each data packet to the source and destination addresses specified in flow pattern 130. At 508, for each of the data packets, host computing device 100 determines whether the data packet matches flow pattern 130, or otherwise satisfies criteria associated with flow pattern 130.

[0049] If host computing device 100 determines that the data packet matches flow pattern 130, host computing device 100 increases a pattern reference counter (e.g., at the kernel level) and collects context data describing the matching data packet at 510 (e.g., at the user level). For example, host computing device 100 collects the context data by identifying at least one of a forwarding port, any filters applied to the data packet, results of any applied filters, a virtual network encapsulating the data packet, and whether the received data packet has been copied, cloned, and/or dropped.

[0050] In some embodiments, host computing device 100 collects the context data for only those data packets matching flow pattern 130, out of a plurality of data packets handled by host computing device 100. As such, host computing device 100 does not collect the context data for each and every data packet received, but collects data for only those data packets matching flow pattern 130. In this manner, the cost of flow monitoring, in terms of latency and/or resource consumption (e.g., processing, memory, etc.), is reduced relative to merely sending a report on every single routed data packet whether or not the data packets are of interest to user 108.

[0051] At 512, host computing device 100 transmits the collected context data, for each matching data packet, to a reporting device, server, or other entity. In some embodiments, each of host computing devices 100 within virtual datacenter 302 has access to the reporting device to enable the aggregate of the context data on the reporting device.

[0052] If the data packet does not match flow pattern 130 at 508, context data is not collected and transmitted to the reporting device. Rather, processing continues with the next handled data packet at 506.

[0053] In embodiments, the operations illustrated in FIG. 5 may be performed by a virtual gateway application, operating

as a virtual gateway from a first virtual network to a second virtual network. In such embodiments, the virtual gateway application derives flow table information from flow pattern 130. As such, the virtual gateway application compares each routed data packet to flow pattern 130 by comparing each routed data packet to the derived flow table information.

[0054] FIG. 6 is a block diagram illustrating exemplary flow monitoring or tracing in a stack in host computing device 100. In the example of FIG. 6, the stack includes a netfilter front layer (e.g., an IOChain/DVFilter layer for distributed virtual filtering), a virtual switch layer (e.g., a PortGroup layer), a netfilter back layer (e.g., another IOChain/DVFilter layer for distributed virtual filtering), and a device driver layer (e.g., a NicDriver layer). The virtual extensible LAN, such as the VXLAN brand software from VMware, Inc., represents an encapsulation mechanism that runs between virtual switches to enable VMs 235 to be deployed on, and move between, any servers within virtual datacenter 302.

[0055] Exemplary operations associated with flow monitoring are next described. While the operations are described as being performed by host computing device 100 in some embodiments, the operations may be performed by any module, device, application, etc. that acts as a bi-directional packet filter between a virtual switch and a device driver to filter network traffic entering the virtual switch from VMs 235 or exiting the virtual switch to VMs 235. In some embodiments, the bi-directional packet filter is embodied in a software Ethernet switch providing network connectivity within a virtual network environment.

[0056] User 108 defines flow pattern 130 at Operation 1 and flow tracing is enabled at Operation 2. The packet footprint is traced according to flow pattern 130 at Operation 3, and the network flow is analyzed via the data packets. For each matching data packet, host computing device 100 collects the context data described in Table 1 below.

TABLE 1

Exemplary Context Data.	
Packet Metadata	Layer 3 5-tuple Metadata
portIDs	List of portIDs that the packet accessed
Distributed virtual filters	List of filters that the packet passed (e.g. invoked)
IOChains	List of function addresses that the packet passed (e.g., called)
<func, result>	Function calls and the results therefrom
<clone/copy>	Clone or copy records of the packet

[0057] Each of the elements of the context data shown in Table 1 may include additional data, such as a timestamp.

[0058] At Operation 4, the collected context data is sent to a remote server. For example, the footprint record for the data packet is sent to the remote server when the data packet is about to be freed. Alternatively or in addition, the collected context data is batched for a plurality of matching data packets and flushed to the remote server periodically and/or intermittently. For example, host computing device 100 may send 32 footprint records at a time. Upon collecting enough footprint records (e.g., satisfying a pre-defined threshold), the remote server determines the flow record.

[0059] In an exemplary implementation, host computing device 100 reserves extra memory space in each data packet when creating, copying and/or cloning the packet. The extra

memory space is used to capture context data, and may be referred to as the footprint record. When the data packet is copied or cloned, the footprint record is also copied. In some embodiments, host computing device 100 allocates 64 kilobytes for the footprint record, which includes 8 bytes for the functional address, 4 bytes for the result, 4 bytes for the portID, 16 bytes for the timestamp, and 32 bytes for other metadata.

[0060] FIG. 7 is a block diagram illustrating exemplary end-to-end flow monitoring between two organizations. In this example, a user in one organization (Org1) is accessing web services provided by another organization (Org2). The exemplary network path between Org1 and Org2 is illustrated in FIG. 7 and shown to pass through several zones or edges:

[0061] Source→vAppNetwork1→OrgNetwork→ExternalNetwork→OrgNetwork2→vAppNetwork2→Destination

[0062] In the example of FIG. 7, “edge” represents an edge gateway and “app” represents an application firewall. The virtual gateways between Org1 and Org2 may be implemented using the vShield Edge and/or the vShield App, both from VMware, Inc. Aspects of the disclosure access virtual gateway network profiles for the virtual gateways to understand the virtual network flow patterns. For example, the virtual gateway network profiles may be obtained from a cloud operating system and/or cloud application executed by virtual datacenter 302. In the example of FIG. 7, five virtual network flow patterns are illustrated. Flow pattern 130 is given by the application in Org1 with <SRC: 192.168.1.100, DST: 172.30.4.12, DST Port: 80, Protocol: TCP>. Host computing device 100 communicates with the other virtual network appliances shown in FIG. 7 to understand the routing policies on the virtual gateways and then deduce the remaining virtual network flow patterns. The remaining virtual network flow patterns represent the flow patterns (e.g., source address, destination address, port number) along the entire network path. In the example of FIG. 7, the virtual network flow patterns change several times as the traffic passes through several routers in the network path (e.g., from 192.168.1.100 to 172.25.33.22).

[0063] FIG. 8 is a block diagram illustrating exemplary end-to-end flow monitoring with network encapsulation. In the example of FIG. 8, network traffic with flow pattern 130 represented by <SRC: 192.168.1.100, DST: 192.168.1.200, DST Port: 8000, Protocol: TCP> is encapsulated and decapsulated within the same virtual network. The virtual extensible LAN application performing the encapsulation and the virtual extensible LAN application performing the decapsulation execute the operations illustrated in FIG. 5 to enable identification and tracing of data packets matching flow pattern 130. Aspects of the disclosure are operable with any encapsulation/decapsulation operations, such as those associated with the VXLAN brand software from VMware, Inc.

[0064] FIG. 9 is a block diagram illustrating exemplary flow monitoring at virtual gateways. An application implementing the operations illustrated in FIG. 4 (e.g., executing on host computing device 100) enables flow monitoring by communicating with virtual datacenter 302 (e.g., a cloud operating system) to obtain a description of the infrastructure within virtual datacenter 302 at Operation 1. For example, the application establishes an authenticated channel with the cloud operating system managing virtual datacenter 302 to obtain a network profile describing the virtual network infrastructure. The network profile represents the virtual network

topology and configuration information that identifies, among other data, host computing devices 100 sitting in the network flow path.

[0065] At Operation 2, the application identifies the hosts involved in a target network flow path, and distributes flow pattern 130 to each of the hosts to initiate flow monitoring. At Operation 3, flow pattern 130 is distributed to virtual gateways and other virtual network edge appliances to initiate flow monitoring. At Operation 4, the virtual network appliances (e.g., hosts, virtual gateways, etc.) send the context data for data packets matching flow pattern 130 to a remote server for analysis and reporting.

[0066] FIGS. 10A, 10B, and 10C are block diagrams illustrating exemplary filtering of flow information based on user role. The flow information represents the aggregated context data for data packets matching flow pattern 130. In this example, the flow information indicates a network issue and is filtered based on three different user roles: a virtual application owner (e.g., a customer), a network administrator, and a developer. FIG. 10A illustrates the virtual application owner view, FIG. 10B illustrates the network administrator view, and FIG. 10C illustrates the developer view. The checkmarks in the figures indicate successful passing of the data packets, whereas the "X" in the figures indicates where the network issue is occurring.

[0067] The virtual application owner view represents a top level view (e.g., high-level logical view) presenting the network flow from a source virtual network to a destination virtual network. The virtual application owner view indicates to the virtual application owner that the network issue is occurring at the edge of the destination virtual network. The network administrator view represents a middle level view (e.g., infrastructure or stack view) presenting the network flow across the overall virtual network infrastructure. The network administrator view indicates to the network administrator that the network issue is occurring at the edge of the destination organization network. The developer view represents a low level view (e.g., component level) presenting the functional level packet footprints and network flow across the hosts. The developer view indicates to the developer that the network issue is occurring because of a destination edge gateway configuration.

[0068] The different views illustrated in FIGS. 10A, 10B, and 10C may be presented to users 108 by an application executing on the remote server, on a device within virtual datacenter 302, by the cloud operating system managing virtual datacenter 302, and/or by another entity.

[0069] In this manner, aspects of the disclosure present differentiated troubleshooting views according to user role. For example, the functional packet tracing path is presented to technical support and developers, while the logical network path is presented to network administrators to enable efficient identification of network issues in a top-down fashion.

Additional Examples

[0070] The following scenarios are merely exemplary and not intended to be limiting in any way.

[0071] In one scenario, the remote server aggregating the context data for data packets matching flow pattern 130 allows third party vendors to develop filtering criteria to create customized views of the aggregated context data. For example, rather than presenting the three different views in FIGS. 10A, 10B, and 10C, other views may be created.

Exemplary Operating Environment

[0072] The operations described herein may be performed by a computer or computing device. The computing devices communicate with each other through an exchange of messages and/or stored data. Communication may occur using any protocol or mechanism over any wired or wireless connection. A computing device may transmit a message as a broadcast message (e.g., to an entire network and/or data bus), a multicast message (e.g., addressed to a plurality of other computing devices), and/or as a plurality of unicast messages, each of which is addressed to an individual computing device. Further, in some embodiments, messages are transmitted using a network protocol that does not guarantee delivery, such as User Datagram Protocol (UDP). Accordingly, when transmitting a message, a computing device may transmit multiple copies of the message, enabling the computing device to reduce the risk of non-delivery.

[0073] Exemplary computer readable media include flash memory drives, digital versatile discs (DVDs), compact discs (CDs), floppy disks, and tape cassettes. By way of example and not limitation, computer readable media comprise computer storage media and communication media. Computer storage media include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media are tangible and are mutually exclusive to communication media. In some embodiments, computer storage media are implemented in hardware. Exemplary computer storage media include hard disks, flash drives, and other solid-state memory. In contrast, communication media typically embody 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 include any information delivery media.

[0074] Although described in connection with an exemplary computing system environment, embodiments of the disclosure are operative with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with aspects of the disclosure include, but are not limited to, mobile computing devices, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, gaming consoles, microprocessor-based systems, set top boxes, programmable consumer electronics, mobile telephones, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0075] Embodiments of the disclosure may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. The computer-executable instructions may be organized into one or more computer-executable components or modules. Generally, program modules include, but are not limited to, routines, programs, objects, components, and data structures that perform particular tasks or implement particular abstract data types. Aspects of the disclosure may be implemented with any number and organization of such components or modules. For example, aspects of the disclosure are not limited to the specific computer-executable instructions or the specific components or modules illustrated in the figures and described herein. Other embodiments of the disclosure may include different computer-executable instruc-

tions or components having more or less functionality than illustrated and described herein.

[0076] Aspects of the disclosure transform a general-purpose computer into a special-purpose computing device when programmed to execute the instructions described herein.

[0077] The embodiments illustrated and described herein as well as embodiments not specifically described herein but within the scope of aspects of the invention constitute exemplary means for end-to-end tracing of the one or more data packets through virtual datacenter 302 based on a footprint for each of the one or more data packets.

[0078] At least a portion of the functionality of the various elements illustrated in the figures may be performed by other elements in the figures, or an entity (e.g., processor, web service, server, application program, computing device, etc.) not shown in the figures.

[0079] In some embodiments, the operations illustrated in the figures may be implemented as software instructions encoded on a computer readable medium, in hardware programmed or designed to perform the operations, or both. For example, aspects of the disclosure may be implemented as a system on a chip or other circuitry including a plurality of interconnected, electrically conductive elements.

[0080] The order of execution or performance of the operations in embodiments of the disclosure illustrated and described herein is not essential, unless otherwise specified. That is, the operations may be performed in any order, unless otherwise specified, and embodiments of the disclosure may include additional or fewer operations than those disclosed herein. For example, it is contemplated that executing or performing a particular operation before, contemporaneously with, or after another operation is within the scope of aspects of the disclosure.

[0081] When introducing elements of aspects of the disclosure or the embodiments thereof, the articles “a,” “an,” “the,” and “said” are intended to mean that there are one or more of the elements. The terms “comprising,” “including,” and “having” are intended to be inclusive and mean that there may be additional elements other than the listed elements. The term “exemplary” is intended to mean “an example of”

[0082] Having described aspects of the disclosure in detail, it will be apparent that modifications and variations are possible without departing from the scope of aspects of the disclosure as defined in the appended claims. As various changes could be made in the above constructions, products, and methods without departing from the scope of aspects of the disclosure, it is intended that all matter contained in the above description and shown in the accompanying drawings shall be interpreted as illustrative and not in a limiting sense.

We claim:

1. A system for end-to-end virtual network flow monitoring in a virtual datacenter having a plurality of virtual machines (VMs), said system comprising:

a memory area associated with the virtual datacenter, said memory area storing a flow pattern including at least a source address and a destination address, the flow pattern being received from a user in a request to perform virtual network flow monitoring; and

a processor programmed to:

distribute the flow pattern stored in the memory area to a plurality of applications in the virtual datacenter, each of the plurality of applications managing a plurality of VMs as part of at least one virtual network;

aggregate, by the virtual datacenter from the plurality of applications, context data for one or more data packets routed by the plurality of applications and matching the flow pattern; and

determine a role associated with the user; and

filter, by the virtual datacenter based on the determined role, the aggregated context data for presentation to the user.

2. The system of claim 1, wherein the plurality of applications operate as at least one of a host computing device or a virtual gateway.

3. The system of claim 1, wherein the processor is further programmed to display the aggregated context data to the user.

4. The system of claim 1, wherein the flow pattern further comprises at least one of a transmission control protocol (TCP) or a user datagram protocol (UDP).

5. The system of claim 1, wherein the user comprises at least one of a virtual application owner, an administrator, or a developer.

6. The system of claim 1, wherein the processor aggregates the context data across a logical network boundary.

7. The system of claim 1, further comprising means for end-to-end tracing of the one or more data packets through the virtual datacenter based on a footprint for each of the one or more data packets.

8. A method comprising:

for each of a plurality of computing devices associated with one or more virtual networks in a virtual datacenter, the virtual datacenter comprising a plurality of virtual machines (VMs):

accessing, by the computing device, a flow pattern providing a source address and a destination address;

receiving a data packet from the one or more virtual networks;

comparing, by the computing device, the received data packet to the accessed flow pattern; and

based on the comparison, collecting data describing the received data packet and transmitting the collected data for analysis to a reporting device accessible by each of the plurality of computing devices associated with the one or more virtual networks in the virtual datacenter.

9. The method of claim 8, wherein collecting the data comprises identifying at least one of a forwarding port, a filter applied to the data packet, a virtual network encapsulating the data packet, whether the received data packet has been copied, or whether the received data packet has been cloned.

10. The method of claim 8, wherein collecting the data comprises collecting the data on determining that the received data packet matches the accessed flow pattern.

11. The method of claim 8, wherein the computing device receives a plurality of the data packets, and wherein collecting the data comprises collecting the data for only those data packets, from the plurality of data packets, that match the accessed flow pattern.

12. The method of claim 8, wherein receiving the data packet comprises receiving, by a host computing device, a data packet.

13. The method of claim 8, wherein the computing device comprises a virtual gateway application, and further comprising deriving, by the virtual gateway application, flow table information from the accessed flow pattern.

14. The method of claim 13, wherein comparing the received data packet to the accessed flow pattern comprises comparing the received data packet to the derived flow table information.

15. The method of claim 8, wherein the computing device comprises a host, and wherein collecting the data comprises collecting the data across a network stack within the host.

16. One or more computer-readable storage media storing computer-executable components for performing end-to-end virtual network flow monitoring in a virtual datacenter having a plurality of virtual machine (VMs), the computer-executable components comprising:

- a plurality of host components that each, when executed, cause at least one processor to:
 - receive a data packet;
 - compare the received data packet to a flow pattern defined by a user, the flow pattern including at least a source address and a destination address; and
 - based on the comparison, collect context data describing the received data packet; and

a management component in communication with each of the plurality of host components, the management component, when executed, causing at least one processor to:

- aggregate the collected context data from one or more of the plurality of host components;
- determine a role associated with the user; and
- filter, based on the determined role, the aggregated, collected context data for presentation to the user.

17. The computer storage media of claim 16, wherein the management component further receives, from the user, a request to perform the virtual network flow monitoring, the request including the flow pattern.

18. The computer storage media of claim 16, wherein the management component further distributes the flow pattern to each of the plurality of host components.

19. The computer storage media of claim 16, wherein each of the plurality of host components manages a plurality of VMs.

20. The computer storage media of claim 16, wherein at least one of the plurality of host components represents a virtual gateway from a first virtual network to a second virtual network.


* * * * *




(19) **United States**(12) **Patent Application Publication****Kan et al.**(10) **Pub. No.: US 2004/0073655 A1**(43) **Pub. Date: Apr. 15, 2004**(54) **PACKET SEQUENCE NUMBER NETWORK MONITORING SYSTEM**(76) Inventors: **Chao Kan**, Plano, TX (US); **Pierrick Guingo**, Dallas, TX (US); **Emanuele Jones**, Ottawa (CA); **Hubert Ogier**, Paris (FR); **Thierry Labbe**, Dallas, TX (US)



Correspondence Address:

ALCATEL USA**INTELLECTUAL PROPERTY DEPARTMENT
3400 W. PLANO PARKWAY, MS LEGL2
PLANO, TX 75075 (US)**(21) Appl. No.: **10/267,813**(22) Filed: **Oct. 9, 2002****Publication Classification**(51) **Int. Cl.⁷ G06F 15/173**(52) **U.S. Cl. 709/224; 709/230****ABSTRACT**

A network monitoring system (60). The system comprises storage circuitry (32) for storing network packet information, wherein the network packet information includes a predicted identifier. The network monitoring system also comprises at least one monitoring circuit (36) coupled to a network (70) along which network traffic flows in a form of packets. The at least one monitoring circuit programmed to perform the steps (44) of receiving a packet communicated along the network and determining whether the received packet is communicated between a source and destination in a first set of network nodes. Each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination. The at least one monitoring circuit programmed to perform the step of, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.












FIG. 7



PACKET SEQUENCE NUMBER NETWORK MONITORING SYSTEM

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] Not Applicable.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

[0002] Not Applicable.

BACKGROUND OF THE INVENTION

[0003] The present embodiments relate to computer networks and are more particularly directed to a real-time system for monitoring network performance by evaluating packet identifiers such as packet sequence numbers.

[0004] As the number of users and traffic volume continue to grow on the internet and other networks, an essential need has arisen for both the users and network operators to have a set of mechanisms to monitor network performance, detect router availability, and troubleshoot network congestion without introducing additional traffic on the network. This is especially true for Internet Service Providers in order to comply with the Service Level Agreements ("SLAs") that they are now providing to customers. Additionally, such a need is prevalent for the underlying internet protocol ("IP") networks in the global internet. With these mechanisms, the users and network operators also can monitor and query the reliability and availability of the network nodes (e.g., IP routers) and the given internet paths, and the mechanisms also can identify the location of a segment of a given internet path, if any, that has a limitation. Additionally, the internet is evolving towards an advanced architecture that seeks to guarantee the quality of service ("QoS") for real-time applications, such as by putting limits on the upper bound on certain QoS parameters including jitter, throughput, one-way packet delay, and packet loss ratio. Accordingly, tracking QoS performance is also desirable. Given the preceding, various mechanisms for network monitoring have been developed in the art, some of which are described below.

[0005] A first type of network performance analysis mechanism is known in the art as an active mechanism. The active network performance monitoring system sends a series of special test packets or query packets to the underlying networks or routers, and it analyzes the response with a specific purpose. Currently, most tools for active end-to-end QoS monitoring in IP networks are based on the traditional "ping" (i.e., ICMP echo and echo response messages) to measure the roundtrip delay between two hosts. Variations of ping include "Nikhef ping," "fping," "pingplotter," "gnuplotping," "Imeter," and "traceping". Several other tools are variations based on the traditional "traceroute" (which exploits the time-to-live field in the IP packet header), for example, "Nikhef traceroute," "pathchar," "whatroutet," and "network probe daemon." A few large-scale research projects are using "ping" to continually and actively monitor the network performance between various points in the global internet. Examples of these types of projects are: (i) the PingER project at Stanford Linear Accelerator Center ("SLAC") uses repeated pings around various Energy Sciences Network ("ESnet") sites; (ii) the Active Measurement Program ("AMP") project by National

Laboratory for Applied Network Research ("NLNR") performs pings and traceroutes between various NSF-approved high-performance connection sites; (iii) the Surveyor project attempts one-way packet delay and loss measurements between various sites using the global positioning system ("GPS") for time synchronization; and (iv) the national Internet measurement infrastructure ("NIMI") project measures the performance between various sites using traceroute or transmission control protocol ("TCP") bulk transfer.

[0006] A second type of network performance analysis mechanism is known in the art as a passive mechanism and may be contrasted to the active mechanism discussed above. The passive network performance monitoring system performs its traffic analysis in a non-invasive way with respect to the observed networking environment, namely, the system does not introduce additional test packets to the network. As a result, the system does not affect the performance of the network while doing the measurements and querying. The traditional approach usually involves the following three steps: (i) collection of the entire TCP/IP packet (e.g. traffic sniffers, etc.) or the entire packet header data into a database or databases; (ii) hardware and software for analyzing the collected database(s); and (iii) off-line traffic characterization and modeling. Various examples of passive mechanisms have been implemented by certain entities. For example, the NLNR has been using OCXmon monitors to tap into the light of a fiber interconnection by means of optical splitters, and from the connection the system collects a portion of the packet, namely, the packet header. Specifically, traffic data is collected in an abstract format by extracting and storing the packet header in the database within a preset traffic aggregation period. The packet header traffic data in the database is later analyzed (i.e., an off-line traffic analysis). As another example, Cisco offers a NetFlow capability in its large routers. NetFlow identifies traffic flows based on IP source/destination addresses, protocol ID field, type of service ("TOS") field, and router port. Once identified, statistics can be collected for a traffic flow, and exported via user datagram protocol ("UDP") when the flow expires. Flow statistics may include the flow start/stop times, number of bytes/packets, and all IP header fields. As a final example, Lucent Bell Labs has various research projects in traffic analysis, which are mainly concentrated in collection of TCP/UDP/IP packet headers and off-line traffic analysis, modeling and visualization.

[0007] A third type of network performance analysis mechanism is also a type of passive mechanism in that it does not add traffic to the network, but rather than collecting information directly from a packet it instead collects statistics pertaining to the packet; this system is sometimes referred to as a network or element management system. In this system, instead of separately setting up the packet collection process, IP routers periodically (e.g., in 30 minute intervals) store and collect a set of traffic statistics into a built-in Management Information Base ("MIB") using a simple network management protocol ("SNMP") interface, and those statistics may be retrieved by a network or element manager. Typical traffic statistics might be the number of received/forwarded packets, discarded packets, error packets, port utilization, CPU utilization, buffer utilization, and so forth. As in other passive systems, these statistics are collected for later analysis. When a network congestion or event occurs, the SNMP agent embedded in the IP router

sends a trap message to the SNMP manager, which then indicates an alarm in its graphic user interface.

[0008] While the preceding systems have provided certain levels of information about the behavior of traffic in IP networks, the present inventors have observed that such systems also provide various drawbacks. As an example of a drawback with respect to the active systems, they are required to send special test or PING packets and, thus, these additional packets themselves introduce additional traffic onto the route of the regular data packets. Accordingly, at the time of detecting the network congestion or server availability, the network performance and associated QoS being measured will be negatively impacted by this additional burden. Additionally, the QoS is measured for the test packets instead of the actual data packets, and the QoS for the data packets must be inferred indirectly from the QoS measured for the test packets. As an example of a drawback with respect to the passive system, although many traffic studies have attempted to understand the random behavior or composition of internet traffic, these traffic studies have been off-line analysis of historical data. There are no known prominent research projects attempting traffic analysis and control based on real-time traffic measurement or comprehensive traffic profiling. For example, Lucent's projects reflect the traditional approach of collecting large traffic measurement datasets for lengthy periods of time such as on the order of many days and then performing subsequent off-line statistical analysis. Cisco NetFlow essentially measures the volume and duration of each traffic flow for the purposes of accounting and off-line traffic analysis, but is not intended to be used for real-time network monitoring and querying. The OCXmon tool from NLANR is only for asynchronous transfer mode ("ATM") traffic, and is not for the purpose of traffic monitoring and control. Given the off-line use of these tools which, by definition, is delayed by some time which may be on the order of days, any conclusion that may be drawn from the collected data is stale in the sense that it applies to the network environment as it existed at some previous time rather than having a high likelihood of characterizing the network at the time the analysis is concluded; in other words, the status of the network is likely to have changed between the time these prior art systems collect their data and the time they draw any inferences or conclusions during their off-line analysis. Thus, by time the results provide meaningful indicators, the real-time status of the network may have changed. Finally, as an example of a drawback of the element management system, the router MIB is usually designed based on the specific structure and implementation of the underlying network or IP router and, therefore, it will not be the same for equipment from different vendors. For example, the Argent Guardian tool from Argent Software Inc. for performance monitoring and proactive problem detection and correction has different versions for different monitored entities. The Argent Guardian for Cisco can only be used for Cisco routers because it uses the Cisco router MIB to retrieve and query the traffic information. In addition, the MIB is commonly embedded in the IP router. As a result of these constraints, there is considerable difficulty in changing the MIB once the design is complete. Accordingly, there are considerable barriers to the use of such a system in the instance where there is a change in the parameters or statistics to be monitored and where the change is not already built into the MIB.

[0009] In view of the above, there arises a need to address the drawbacks of the prior art as is accomplished by the preferred embodiments described below.

BRIEF SUMMARY OF THE INVENTION

[0010] In the preferred embodiment, there is a network monitoring system. The system comprises storage circuitry for storing network packet information, wherein the network packet information includes a predicted identifier. The network monitoring system also comprises at least one monitoring circuit coupled to a network along which network traffic flows in a form of packets. The at least one monitoring circuit is programmed to perform the steps of receiving a packet communicated along the network and determining whether the received packet is communicated between a source and destination in a first set of network nodes. Each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination. The at least one monitoring circuit programmed to perform the step of, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

[0011] Other aspects are also described and claimed.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

[0012] FIG. 1 illustrates a prior art network layer stack **10** defined to include the transmission control protocol ("TCP").

[0013] FIG. 2 illustrates a prior art plot of the evolution of TCP sequence numbers over time when a TCP session, including the source, destination, and network, behave normally.

[0014] FIG. 3 illustrates a prior art plot of the evolution of TCP sequence numbers over time when during a TCP session the network operates normally but either the source or destination is limiting the bandwidth of communication.

[0015] FIG. 4 illustrates a prior art plot of the evolution of TCP sequence numbers over time when during a TCP session the network is limiting the bandwidth of communication between the source and destination by dropping packets.

[0016] FIG. 5 illustrates a block diagram of a network system **10** into which the preferred embodiments are implemented.

[0017] FIG. 6 illustrates a block diagram of each network monitor NM₁ through NM₅ of FIG. 5.

[0018] FIG. 7 illustrates a flow chart of the operation of each network monitor NM₁ through NM₅ of FIG. 5.

[0019] FIG. 8 illustrates a flow chart of the predictor methodology preferably implemented within the flow chart of FIG. 7.

DETAILED DESCRIPTION OF THE INVENTION

[0020] By way of introduction to one preferred inventive implementation, **FIG. 1** illustrates a known prior art network layer stack **10** defined to include the internet connection protocol ("IP") as discussed below. At the top of stack **10** is an application layer **12**, which typically represents one or more programs that are commonly used by consumers for various reasons, including personal and business computing. Beneath application layer **12** are located two protocol layers, namely, a transmission control protocol ("TCP") layer **14** and a user datagram protocol ("UDP") layer **16**. TCP layer **14** is a connection-oriented layer, that is, once a connection is established from a source to a destination and through that layer, then traffic through that connection is confirmed so that if a packet is lost before it is confirmed by its destination then the source re-transmits the packet. In contrast, UDP layer **16** is connectionless and, thus, no confirmation is made as between a packet passing from a source to its destination; accordingly, if a packet is lost in a UDP communication then there is no re-transmission of that packet. Comparing then the TCP and UDP layers **14** and **16**, TCP assures greater integrity of communications yet it requires greater overhead than UDP.

[0021] Below layers **14** and **16** is an IP layer **18**, using IP as introduced above. Thus, both TCP and UDP network packets along a given network pass through the IP layer **18** and therefore also must comply with the requirements of the IP protocol. IP layer **18** is also a connectionless layer, although for TCP packets passing through the layer the connection-oriented nature of TCP layer **14** may be relied upon to provide additional data integrity. Due to the connectionless nature of IP layer **18**, there is a trend in the networking art to provide measures of so-called quality of service ("QoS") in order to give priority or otherwise provide guidance to the flow of connectionless IP packets. In any event, and importantly for one example of the preferred embodiment as detailed in this document, each TCP (as opposed to UDP) network packet passing from a single source to a single destination and through IP layer **18** includes a unique TCP sequence number that thereby distinguishes that packet from all other TCP packets communicated from the source to the destination in a given flow session. This TCP sequence number is required as a TCP parameter and may be used for various purposes as known in the art, while the preferred embodiments further monitor the number to evaluate network performance as detailed below. Before reaching that discussion and completing the discussion of **FIG. 1**, located below IP layer **18** is a physical layer **20**. Physical layer **20** is commonly associated with the physical hardware at a source or destination, as typically represented by a network interface having a media access control ("MAC") address.

[0022] As further explored below, the preferred embodiments evaluate network performance, preferably as it is reflected in communications between specified nodes such as in specified sub-networks in the network, in response to a real-time evaluation of the progression of TCP packet sequence numbers. In this regard, in part the evaluation by the preferred embodiments includes consideration of three prior art observations made by the NLNR as relating to TCP sequence numbers. These three observations are explored below in connection with **FIGS. 2 through 4**.

[0023] **FIG. 2** illustrates a plot of the evolution of TCP sequence numbers over time when a TCP session, including the source, destination, and network, behave normally. In **FIG. 2**, the horizontal axis represents the passage of time since the beginning of the TCP session observation, while the vertical axis represents the sequence number of the last byte of each data packet. Note that the sequence numbers are shown for simplicity of illustration to start from zero, whereas in reality they typically commence from a non-zero, and likely random, number, and then increase from that point. In any event, the sequence numbers will increase monotonically unless there are reordered or retransmitted packets. While intuitively one might expect a linear relationship for normal behavior, instead note that the plot provides a number of upward curves each ending with a peak, such as shown by ways of example with a curve prior to a point P_1 at time t_1 and a point P_2 at a time t_2 . Note that these curves, followed by respective peaks, occur under ideal network operation because the source transmission rate will tend to increase toward and reach the normal upper limit of the traffic rate for the destination (or the network). In other words, under TCP protocol, transmission occurs in a TCP control window, where the source starts transmitting at a certain rate, and hence with a corresponding steady increase in TCP sequence numbers, followed by an attempt within the window over time to increase the transmission rate and hence with a change in the rate of the ascension of the TCP sequence numbers. The source's periodic attempts during the TCP control window to increase the TCP sequence number frequency continues until the limit of the destination (or network) to accurately communicate packets is reached. As this limit is reached, the source will reduce its transmission rate for a time and thereafter again continue to increase it until the limit is again reached later in time.

[0024] **FIG. 3** illustrates a plot of the evolution of TCP sequence numbers over time, when during a TCP session the network operates normally but either the source or destination is limiting the bandwidth on communication. Here again, the horizontal axis represents the passage of time while the vertical axis represents the increase in sequence numbers. However, the plot of TCP sequence numbers reflects a straight line. Typically such a result is anticipated because it represents that the source is not reaching the maximum bandwidth limitation of the network or destination. Thus, while the TCP sequence numbers increase, the increase of the sequence numbers from one to the next is in a constant fashion across the entire illustrated plot, as opposed to an acceleration of those numbers that increases over time as shown in **FIG. 2** and that also peaks from time to time as a normal limit is reached.

[0025] **FIG. 4** illustrates a plot of the evolution of TCP sequence numbers over time, when during a TCP session the network is limiting the bandwidth of communication between the source and destination and, thus, it is presumed during such time that the source and destination are otherwise capable of communicating at their respective maximum bandwidths. However, the plot of TCP sequence numbers illustrates curves with discontinuities between the curves. For example, at time t_i , the TCP sequence numbers are at a point P_{10} , but transmission is interrupted until a time t_2 . At time t_2 , however, note that a slight drop occurs in the TCP sequence number as shown at point P_{11} . This type of pattern repeats later at time t_3 where a sequence number shown at point P_{12} is reached, but an interruption then occurs until

communication through to the destination from the source is achieved at time t_4 and having a sequence number at point P_{13} . The behavior illustrated by these discontinuities occurs due to packet losses along the network as resulting from network irregularities. In other words, recalling that TCP communications are connection-oriented, when a network irregularity occurs one or more TCP packets will be lost, and those packets will be required to be re-transmitted by the source. As a result, the source, which receives the indication of the failure at a time when it was prepared to transmit a new packet having a certain TCP sequence number, is instead required to backtrack to re-transmit a previous packet with a lower TCP sequence number, followed by all packets following the re-transmitted packet. The beginning of such a re-transmission, therefore, is shown by way of example at point P_{11} (and various other points in FIG. 4), where a drop is shown in TCP sequence number as compared to point P_{10} , where the drop occurs due to the backtrack and re-transmission by the source of a previously-transmitted packet(s).

[0026] FIG. 5 illustrates a block diagram of a system 60 into which the preferred embodiments are implemented. System 60 generally includes a number of stations ST_1 through ST_4 , each coupled to a network 70 via a router, and each operable to send packets as a source or receive packets as a destination. By way of example, network 70 is an internet protocol ("IP") network such as the global internet or other IP-using network, where each station and IP networks in general are well known in the art. However, one skilled in the art should appreciate that the use of the IP protocol is by way of illustration, and many of the various inventive teachings herein may apply to numerous other protocols, including by way of examples asynchronous transfer mode ("ATM"), token ring, Novell, Apple Talk, and still others. In any event, returning to network 70 as in IP network, and also by way of an example, each station ST_x may be constructed and function as one of various different types of computing devices, all capable of communicating according to the IP and to include the TCP protocol. Lastly and also by way of example, only four stations ST_x are shown so as to simplify the illustration and example, where in reality each such station may be proximate other stations (not shown) and at a geography located at a considerable distance from the other illustrated stations. For example, station ST_1 might be located in Montana while station ST_2 is located in California, and station ST_3 might be located in New York while station ST_4 is located in Florida.

[0027] Continuing with FIG. 5, along the outer periphery of network 70 are shown a number of edge routers ER_1 through ER_{11} , while within network 70 are shown a number of core routers CR_1 through CR_4 . The terms edge router and core router are known in the art and generally relate to the function and relative network location of a router. Typically, edge routers connect to remotely located networks and handle considerably less traffic than core routers. In addition and due in part to the relative amount of traffic handled by core routers, they tend to perform less complex operations on data and instead serve primarily a switching function; in other words, because of the tremendous amount of throughput expected of the core routers, they are typically hardware bound as switching machines and not given the capability to provide operations based on the specific data passing through the router. Indeed, core routers typically do not include much in the way of control mechanism as there

could be 10,000 or more connections in a single trunk. Further, typically core routers do not involve their operations with TCP related items and instead deal at the IP level and below. In contrast, edge routers are able to monitor various parameters within data packets encountered by the respective router. In any event, the various routers in FIG. 5 are shown merely by way of example, where one skilled in the art will recognize that a typical network may include quite a different number of both types of routers. Finally, note that each core router CR_x also generally may be constructed and function according to the art, with the exception that preferably the core routers may include additional functionality for purposes of traffic irregularity detection based on TCP sequence number evaluation as described below. In addition, each edge router ER_x also generally may be constructed and function according to the art, with the exception that preferably the edge routers are responsive to the traffic irregularity detection of the core routers, as also described below.

[0028] Completing the discussion of FIG. 5, note that the various stations, edge routers, and core routers therein are shown connected to one another in various fashions and also by way of example. Such connections are intended to illustrate an example for later discussion of the preferred operation and also to reflect a general depiction of how networks are generally established. Thus, each station ST_x is shown connected to a single edge router ER_x , where that edge router ER_x is connected to one or more core routers CR_x . The core routers CR_x , also by way of example, are shown connected to multiple ones of the other core routers CR_x . By way of reference, the following Table 1 identifies each station and router shown in FIG. 5 as well as the other device(s) to which each is connected.

TABLE 1

station or router	connected nodes
ST_1	ER_1
ST_2	ER_{10}
ST_3	ER_5
ST_4	ER_7
ER_1	ST_1 ; CR_1
ER_2	CR_3 ; CR_2
ER_3	CR_2
ER_4	CR_2
ER_5	ST_3 ; CR_2 ; CR_3
ER_6	CR_3 ; CR_4
ER_7	ST_4 ; CR_4
ER_8	CR_4
ER_9	CR_4
ER_{10}	ST_2 ; CR_1
ER_{11}	CR_4
CR_1	ER_1 ; ER_{11} ; ER_{10} ; ER_2 ; CR_2 ; CR_3 ; CR_4
CR_2	ER_2 ; ER_3 ; ER_4 ; CR_1 ; CR_3 ; CR_4 ; ER_5
CR_3	ER_5 ; ER_6 ; CR_2 ; CR_1 ; CR_4
CR_4	ER_7 ; ER_8 ; ER_9 ; CR_1 ; CR_2 ; CR_3 ; ER_6

[0029] Given the various illustrated connections as also set forth in Table 1, in general IP packets flow along the various illustrated paths of network 70, and in groups or in their entirety such packets are often referred to as network traffic. In this regard and as developed below, the preferred embodiments operate to identify any irregularities in such network traffic. Finally, note that FIG. 5 may represent a simplified

version of a network or the internet in that only a few stations and routers are shown, while one skilled in the art will readily appreciate that the inventive concepts in this document may be applied to a larger number of stations, routers, and the network interconnectivity between those devices.

[0030] FIG. 5 also illustrates a number of network monitors NM_1 through NM_5 , where the choice of five such network monitors is only by way of example given the limited amount of other hardware that is shown for network 70. As detailed below, each network monitor NM_x is operable to sample, and possibly respond to, irregularities in the TCP sequence number in an IP packet that passes along the conductor(s) to which the network monitor is connected. Further, the components of each network monitor NM_x are described below, but at this point the connections of each such monitor are noted in the following Table 2:

TABLE 2

network monitor	connected nodes
NM_1	CR_1 ; CR_2
NM_2	CR_2 ; CR_3
NM_3	CR_1 ; CR_3
NM_4	CR_3 ; CR_4
NM_5	CR_1 ; CR_2 ; CR_3 ; ER_7 ; ER_8 ; ER_9

[0031] FIG. 5 and Table 2 demonstrate that each of network monitors NM_1 through NM_4 is connected to sample packets passing along the conductor(s) between a pair of core routers. However, network monitor NM_5 is by way of an alternative example incorporated into core router CR_4 . As a result, network monitor NM_5 is able to sample packets communicated with any of the nodes to which core router CR_4 is connected, namely, core routers CR_1 , CR_2 , CR_3 , and edge routers ER_7 , ER_8 , and ER_9 . Thus, the contrast of network monitor NM_5 to the other illustrated network monitors NM_1 through NM_4 is shown to demonstrate that in the preferred embodiment each network monitor NM_x may sample packets as a stand alone entity or may be combined with the hardware and software of an existing router; indeed, in the preferred embodiments a network monitor NM_x also may be combined with network or element management systems. In any event and by way of introduction to details provided later, in the preferred embodiments the sampling functionality of each network monitor NM_x permits real-time monitoring of the TCP sequence numbers in network packets, and in response determinations may be made, and actions may be taken, based on detecting network traffic irregularities based on the behavior of the TCP sequence numbers.

[0032] FIG. 6 illustrates a block diagram of each network monitor NM_1 through NM_4 of FIG. 1, with the further understanding that functionally the following description also may be applied to network monitor NM_5 with the addition that certain functionality may be provided by the hardware and software already available from core router CR_4 . Turning then to FIG. 6, first note that a console 30 is associated with network monitor NM_x , where in the preferred embodiment a single such console 30 communicates with multiple network monitors NM_x . For example, returning briefly to FIG. 5, preferably each of network monitors

NM_1 through NM_5 communicates with a single console 30, where such communications also may be by way of packets between console 30 and the network monitors NM_x . Console 30 may be constructed by one skilled in the art using various forms of hardware and software, where the selection is a matter of implementation choice in order to achieve the functionality described in this document. Turning to that functionality, console 30 preferably provides an administration (configuration) function and a characterization (reporting) function. To permit a user to perform these functions, various interfaces may be used such as a graphical Web-based configuration tool or a command line tool, where one preferable approach implements console 30, by way of example, using the Apache Web server with the PHP server-side scripting language to provide a dynamic interface to a flow store 32. In any event, the user interface preferably provides different screens for each of the administration and characterization functions. The actual interface and screens, however, may be implemented in various forms, where it is desirable for any form that a user can properly control console 30 to perform its administration and characterization function with respect to its flow store 32. Preferably, a network administrator or the like uses the administration functionality of console 30 to create a set of rules and to provide those rules along with a predictor methodology to an RTFM meter 36 described later, where those rules cause meter 36 to retrieve TCP sequence numbers and the predictor methodology is later used to evaluate the history of the TCP sequence numbers for a flow, as also discussed later. Also, once the system is configured and left to collect TCP sequence number information from or about the monitored packets, the network administrator can use the characterization screen to query the information so as to generate network status reports based on the collected information and to thereby detect any network irregularities.

[0033] As introduced above, console 30 is connected to a flow store 32, which preferably represents a storage medium that stores a flow database relating to monitored TCP sequence numbers. In the preferred embodiment, each network monitor NM_x includes its own flow database, although alternative embodiments may be created where more than one network monitor NM_x shares a common flow database. In the preferred embodiment the flow database in flow store 32 is an SQL-compatible database using the PostgreSQL relational database management system, and console 30 communicates with this database through the Web server's PHP link to the SQL database. Thus, any administration and configuration changes made via console 30 are passed directly to flow store 32. Given the preceding, one skilled in the art should appreciate that access to flow store 32 can be achieved by SQL queries, enabling network administrators to automate the configuration process or integrate report generation. In the preferred embodiment, flow store 32 stores what is referred to in this document as a "rule set" (or "rule sets" when plural), which is initially provided to the flow store 32 from console 30 as part of the administration function and which is also conveyed to meter 36. Additionally, flow store 32 stores the progression of TCP sequence numbers from a monitored flow between a source and destination node, as further detailed below. Moreover and as also discussed below, flow store 32 may store numerous different progressions of TCP sequence numbers, each corresponding to a monitored flow between sets of corresponding source and corresponding destination nodes. Still fur-

ther, flow store 32 preferably stores packet time of arrival data, as also detailed later. Lastly, recall that an RTFM meter 36, described below, stores a predictor methodology. From this predictor methodology, analyses are preferably made on the TCP sequence number progression(s) and times of arrival as stored in flow store 32, and those analyses may be reported to console 30 and are also used as a basis to respond to network irregularities.

[0034] Continuing with FIG. 6, recorder layer 34 provides an interface between flow store 32 and a meter 36 (or meters) in use and coupled to the network, and meter 36 copies TCP sequence numbers from respective selected network packets passing through the meter. Specifically, as each network packet is identified as between a particular same source and destination node, meter 36, in response to a predictor methodology, provides the TCP sequence number from the packet to recorder layer 34, which in turn communicates that information to flow store 32. In the preferred embodiment, the monitored source and destination node can be determined randomly, or it can be identified from one or more rules in a rule set in flow store 32. Thus, the applications of recorder layer 34 can be separated into two categories: manager applications and meter reading applications. Manager applications configure meter 36, based on information, including rules in one or more rule sets, in flow store 32. Meter reading applications permit the TCP sequence number data collected by meter 36 to be provided in a data format usable by flow store 32 and, indeed, recorder layer 34 facilitates the passage of that data into the flow database of flow store 32. Recorder layer 34 applications may pass information between flow store 32 and the network probes of meter 36 either synchronously or asynchronously. This gives network administrators the flexibility of either using real-time network flow meters (i.e. NeTraMet) or importing data from other network analysis tools that are not able to provide real-time data (e.g. Cisco NetFlow data).

[0035] In the preferred embodiment, meter 36 is a Real-Time Traffic Flow Measurement ("RTFM") meter which is a concept from the Internet Engineering Task Force ("IETF"). As known in the RTFM art, RTFM meters are previously known to be used in systems for determining the service requested by IP packets that are passing through a network for purposes of collecting revenue, where such a service is identifiable by the transport port number specified in each IP packet. For example, RTFM meters are currently being considered for use in systems whereby an internet user is charged based on the service he or she is using on the internet; for example, a different fee may be charged to the user for each different internet service, including mail, video, phone calls, and web browsing. However, as detailed in this document, the preferred embodiment implements the RTFM meter instead to analyze each IP packet and to determine if the packet is a TCP packet between a given source and destination and, if so, to copy its TCP sequence number for storage so that the copied numbers can be interpreted for purposes of monitoring network activity so as to evaluate network performance and, hence, identify any irregularities with such performance. More specifically, meter 36 physically probes the underlying network traffic and each time meter 36 detects an IP packet on the network, it determines whether the packet is a TCP packet and is between a given (i.e., either randomly or specifically identified) source and destination. Accordingly, during the real-

time passage of numerous IP packets by the evaluating meter 36, that meter 36 does not always copy a fixed portion of each such packet into a database such as the entire packet or the entire packet header; instead, each meter 36 evaluates the source and destination address of each TCP packet and, if a given source and destination are matched, then meter 36 communicates the TCP sequence number for the packet to recorder layer 34. In response, recorder layer 34 inserts the captured TCP sequence number into flow store 32. Accordingly, that stored information is then available for analysis and reporting to console 30, as detailed below.

[0036] FIG. 7 illustrates a flow chart of an overall method 40 of operation for each network monitor NM_x of FIGS. 5 and 6, where method 40 is now described and additional preferred details are provided later with respect to FIG. 8. Additionally, while method 40 is described below with respect to a single network monitor NM_x , preferably each network monitor NM_x performs a comparable method at its respective location in network 70. As a result, distributed operations occur at different network locations, where by way of example the results of each individual network monitor NM_x may be accumulated by a common resource, such as at a single console 30. From this accumulated information, greater accuracy may be achieved in identifying network irregularities and their causes.

[0037] Turning to method 40, in step 42 a network monitor NM_x according to the preferred embodiment first detects and stores the TCP sequence numbers for each of a number n TCP packets between a first set of source and destination nodes. In the preferred embodiment, the first set of source and destination nodes can be selected randomly or may be specified by a rule in flow store 32 that is conveyed to meter 36. By way of one example, preferably such a first set includes only a single source and a single destination. Also by way of example, each of the single source and single destination is likely in a different physical sub-network, where each such sub-network is identifiable in that nodes in each sub-network share a same portion of their address as directed to the sub-network, where typically such an address portion is that following the subnet mask in the packet address. By way of an illustrative example, assume that the source is the FIG. 5 station ST_1 in a first physical sub-network in Montana and the destination is station ST_3 in a second physical sub-network in New York. Thus, for step 42, the RTFM meter 36 of the given network monitor NM_x evaluates each packet passing through the meter to determine if the packet is a TCP packet from the certain source to the certain destination in the two respective sub-networks. If so, then the TCP sequence number for that packet is stored. Once n such TCP sequence numbers are stored (i.e., corresponding to n packets between the same source and destination along the logical sub-network of interest), then step 42 also processes those sequence numbers so that the regularity of subsequently received packets may be predicted, but a discussion of such processing is deferred until later. In any event, note that at the initial completion of step 42, a first window of n packets will have been detected and processed; this window is updated to maintain a total of n packets by later processing subsequent selected packets, where for each such subsequent packet its information is processed and added to the overall information while at the same time the information for the then-oldest packet is removed from the processed information. These details are

better understood below. In any event, once step 42 is complete, method 40 continues to step 44.

[0038] In step 44, the network monitor NM_x continues to monitor IP packets until it detects another TCP packet between the step 42 source and destination, and then step 44 determines whether an irregularity exists in the TCP sequence number of the additionally-detected packet. Specifically, from the plots of FIGS. 2 through 4, once a sufficient number n of TCP sequence numbers are collected in step 42, a plot of those numbers and the TCP sequence number from the additional packet in step 44 may be compared as against the trends represented by the plots of FIGS. 2 through 4, where the case of FIG. 4 is of particular interest in the preferred embodiments. However, in practice, such an approach may require a considerably large value set for the value n of step 42 or it may involve a considerable delay in either collection or analysis time. Accordingly, in the preferred embodiment, step 44 collects one TCP sequence number from a next TCP packet in the first set. Additionally and as detailed below, through the use of the predictor methodology in meter 36, as also introduced with respect to the processing of n packets in step 42 so that the regularity of future packets may be predicted, step 44 determines whether the actual TCP sequence number of the newly-collected TCP packet is within a range as predicted from the n previously-collected TCP sequence numbers of step 42; further, step 44 also preferably determines whether the newly-collected TCP packet arrives within a time as expected based on the time of reception of one or more of the n previously-collected sequence numbers of step 42. Additional details for this predictor methodology are described below. For sake of continuing with FIG. 7 at this point, however, step 44 generally determines, such as through the use of the predictor methodology, whether an irregularity in the TCP sequence number of the newly-detected packet exists, such as may be indicated if the TCP sequence number of the newly-detected packet, and/or the time of arrival of the newly-detected packet, deviates beyond a threshold. Such an irregularity may reflect the conditions of FIG. 4. If no such irregularity is detected, then the flow returns to step 44 to detect another TCP packet in the first set and to repeat the same analysis with respect to its time of arrival and TCP sequence number. If an irregularity is detected, then flow continues from step 44 to step 46.

[0039] Before discussing step 46, note also in step 44 that regardless of whether an irregularity is detected, step 44 also indicates that the newly-detected packet is processed to update the data based on n packets. Specifically, recall that step 42 detects a first group of n packets in the first set and processes the sequence numbers of packets so that the regularity of future packets may be predicted. Looking now to step 44, the latter indication of processing in that step 44 is intended to demonstrate that step 44 also updates the previous data based on the n packets of the preceding step 42. In the preferred embodiment, this updating is such that the packet detected in step 44 becomes one of the n packets, whereas the oldest first set packet received by step 42 is now removed from the n packets. In other words, preferably the data relating to n packets is updated in a sliding window type manner, where at any given time after step 42, the data and the predictions derived from that data as detailed below, relate only to the most recently received n packets in the first

set. Once updated in this manner, the flow continues from step 44 either to step 46 or to return to step 44, as discussed earlier.

[0040] In step 46, having been reached due to a detected TCP sequence number irregularity between a source and destination in the first set from step 44, the network monitor NM_x detects a TCP sequence number from one or more additional TCP packets communicated in the first set (e.g., as between the single source and single destination mentioned in the example of step 42). For each of these additionally-detected packets in the first set, step 46 determines whether the TCP sequence number (and arrival time) also indicates an irregularity as compared to previous TCP sequence numbers and times of arrival of packets between the same source and destination. If an irregularity is not detected, then the preferred embodiment operates under the assumption that the irregularity detected in the preceding step 44 was either an isolated incident, a relatively minor irregularity, or otherwise not reason to make a network response or adjustment, and the flow of method 40 returns from step 46 to step 44. On the other hand, if an irregularity is detected in step 46, then the preferred method flow continues to step 48. As detailed below, preferably an irregularity is considered as detected if the TCP sequence number and/or time of arrival of the TCP sequence packet deviates beyond a threshold, where the deviation relative to the threshold is measured based on an predicted TCP sequence number and time of arrival as provided by the predictor methodology. This threshold may be selected by one skilled in the art, and by way of example may be on the order of five percent from the predicted value.

[0041] Before discussing step 48, note that step 46 also indicates at its conclusion, and in the same manner as discussed above with respect to step 44, that regardless of whether an irregularity is detected, step 46 also processes the newly-detected packet to update the data based on n packets. Once more, therefore, this indication is intended to represent the update operation achieved by the sliding window, to again encompass information based on a total of n packets. Accordingly, in step 46 the updating is such that the packet (or packets) detected in step 46 becomes one of the n packets, whereas the oldest first set packet received by step 42 or step 44 is now removed from the n packets. Again, therefore, the data after step 46, and the predictions derived from that data as detailed below, relate only to the most recently received n packets in the first set. Once updated in this manner, the flow continues from step 46 either to step 48 or to return to step 44, as discussed earlier.

[0042] In step 48, the network monitor NM_x continues to monitor IP packets and now collects IP sequence numbers and time of arrival data for packets in a second set of source and destination nodes, where the second set is preferably related in some rational sense to the first set (e.g., the single source and single destination) of steps 42, 44, and 46 and the second set is also defined by a rule or rules conveyed from flow store 32 to meter 36. For example, recall earlier that step 44 was demonstrated by monitoring packets to detect those between source station ST_1 in a first sub-network in Montana and destination station ST_3 in a second sub-network in New York. Further, for step 48, and since the second set is preferably related to the first set, then as a first example of step 48 consider that the second set includes any source node in the same physical sub-network as station ST_1 in

Montana and any destination node in the same physical sub-network as station ST₃ in New York. As a second example of step 48, the second set may include any source node in a different physical Montana sub-network (i.e., not including the sub-network of station ST₁) and any destination node in a different physical New York sub-network (i.e., not including the sub-network of station ST₃). As a third example of step 48, the second set may include any source node in the same physical sub-network as station ST₁ in Montana, but any destination node in a specified physical sub-network as station ST₄ in Florida; this latter example may prove beneficial where again this second set relates in some sense to the first set, such as may be the case where the destination sub-network of the first set (i.e., in New York) communicates packages received from Montana onward to Florida. In any event, the relationship between the first and second set is preferably stated or implied in the rule set(s) stored in flow store 32. In this manner, one skilled in the art may establish different rules to thereby encompass different relationships between the first and second monitored steps, and to thereby draw corresponding conclusions based on any irregularities in the traffic flow between the nodes of the two different sets. Also in this regard, in one preferred embodiment the second set may include the first set, and it also preferably includes a greater number of nodes than the first set. For example, step 44 (and 46) may be directed to a single source node and a single destination node, while step 48 is directed to multiple source nodes and multiple destination nodes. In any event, therefore, in step 48 and for each packet passing through the network monitor NM_x, if it is a TCP packet of the second set, then the network monitor NM_x preferably copies the TCP sequence number, source address, destination address, and time of arrival for each of n packets between a same source and destination in the second set. In other words, if the second set has three sources, each communicating with three respective destinations (in addition to the first set source and destination), then step 48 endeavors to collect packet data for each these three pairs of source/destination nodes, thereby accumulating data for a total of three times n packets. Thereafter, method 40 continues from step 48 to step 50.

[0043] In step 50, the network monitor NM_x generally operates in the same fashion as described above with respect to step 46, that is, another packet is detected and processed; however, step 50 performs its analysis with respect to the packet information obtained from step 48 in view of a next received second set packet. In other words, while step 46 analyzed packets from the first set, step 50 analyses the TCP sequence numbers and times of arrival for packets from the second set, again determining whether those two parameters deviate beyond a certain percent relative to a predicted value provided by a predictor methodology. In the preferred embodiment, the threshold of step 50 may be the same or different than the threshold used by step 46. In any event, if the deviation, if any, is less than or equal to the step 50 threshold, then method 40 proceeds under the assumption that the earlier-detected irregularity from steps 44 and 46 were limited to the nodes in the first set, and the method returns to step 44; optionally, however, as also shown in FIG. 7, before returning to step 44, the flow can pass through a step 50a wherein an alarm is issued relating solely to the first set of packets. In other words, if no irregularity is found in the second set packet as examined by step 50, then there is still the observation from the affirmative

findings in steps 44 and 46 that there was an irregularity in at least two first set packets. Thus, if the level of granularity at interest is such that irregularities solely within the first set are of interest, then step 50a represents an indication, such as an alarm, relating to the first set so that future examination or action may be taken with respect to the cause of the first set irregularity. Continuing with step 50, however, if the second set packet deviation exceeds the step 50 threshold, then method 40 continues to step 51.

[0044] Before discussing step 51, note that step 50, in the same manner as steps 44 and 46 discussed above, also indicates at its conclusion the operation of processing the newly-detected packet to update the data based on n packets. Here, note that the packet at issue is a second set packet (as monitored from step 48). Accordingly, the data based on a sliding window of n packets continues to be updated, but now that data is reflective of packets in the second set, which recall preferably reflects a larger number of source and/or destinations as compared to the first set. Further, the updated data is only that relative to the specific source and destination between which the step 50 packet was communicated. Lastly, as this update occurs for a packet, the data from the oldest packet of the previously-processed n packets is preferably no longer maintained.

[0045] In step 51, having been reached due to a detected TCP sequence number irregularity between a source and destination in the second set from step 50, the network monitor NM_x detects one or more additional TCP packets communicated in the second set. For each of these additionally-detected packets in the second set, step 51 determines whether the TCP sequence number (and arrival time) also indicates an irregularity. If an irregularity is not detected, then the preferred embodiment operates under the assumption that the irregularity detected in the preceding step 50 was either an isolated incident, a relatively minor irregularity, or otherwise not reason to make a network response or adjustment, and the flow of method 40 returns from step 51 to step 44, again with the option as shown in step 50a of indicating an alarm relating to the irregularities detected earlier with respect to the first set packets. On the other hand, if an irregularity is detected in step 51, then the preferred method flow continues to step 52. Lastly, as with steps 44, 46, and 50, for a given packet between a specific source and destination, step 51 also concludes by updating the data from the sliding window of the most recent n packets in the second set as between that specific source and destination, as updated for each packet (or packets) as detected by step 51.

[0046] From the preceding, one skilled in the art will appreciate that step 52 is reached when method 40 has detected irregularities in the TCP sequence numbers and/or TCP packet times of arrival in both the first and second sets of source and destination nodes. Thus, irregularities have been detected in communications between a plural number of source and destination nodes having a relationship to one another, where that relationship is implied by the rule(s) that define the first and second sets. Note that the actual irregularity may be for various reasons. Moreover, the specific reason(s) causing the irregularity may be inferred based on the amount of the deviation between the actual and predicted values. For example, if the predicted and real values differ by a figure of approximately 50 percent, then the preferred embodiment concludes that a packet (or more than one packet) was lost. This might occur because the source TCP

control window was halved or the destination may have dropped a packet. Still other examples may be ascertained by one skilled in the art. In any event, in step 52, method 40 attempts a corrective action to improve traffic flow in view of the irregularity. Preferably, such a corrective action is taken by the network meter issuing an event report, such as via a network packet, to the edge routers ER_x on the logical sub-network having the detected irregularity(ies). The edge router, in response, may take various actions. For example, one or more alarms may be set off to indicate to a user or network manager that an irregularity or irregularities have been detected. As another example, traffic queuing at the edge routers ER_x relating to the logical sub-network at issue may be adjusted. As still another example which is likely a rare case, there may be an observation that 90% of the flows are losing packets independently of their source and destination; in this case there is likely a considerable issue with the link where the abnormal behavior has been observed.

[0047] Two additional observations are also noteworthy with respect to method 40. As a first observation, note that method 40 only generally illustrates the use of two sets of packets, where by way of example the first set is reflected in steps 42 through 46 and in one instance may be packets between only a single source and single destination of two respective physical sub-networks, where the second set is reflected in steps 48 through 51 and in one instance may be all packets between any source and any destination of the same two respective physical sub-networks. However, the present inventive teachings contemplate different divisions for each set, where there may be additional iterations for number, where preferably each additional set encompasses all packets in a lesser analyzed set. In this manner, the analysis of method 40 effectively continues to expand its scope to consider a greater amount of network traffic, each time including the traffic as considered during the lesser scope of analysis. In operating in this fashion, a greater amount of data may be collected to determine the magnitude of a potential irregularity in network traffic. As a second observation, note that each of steps 46 and 51 states that it evaluates one or more additional packets. This is intended to demonstrate that the irregularity detection for a given set can be based on any number of packets; thus, by looking only at FIG. 7, steps 44 and 46 may appear to represent the analysis of only two first set packets, but the earlier statement that step 46 may include multiple packets is intended to indicate that the inventive scope includes a variation where more than two first set packets are processed, that is, for each of those packets, they are analyzed to determine if an irregularity exists, and the data from those packets is also used to update the sliding window of n packets. Similarly, from FIG. 7 alone, steps 50 and 51 may appear to represent the analysis of only two second set packets, but the earlier statement that step 51 may include multiple packets is intended to indicate that the inventive scope includes a variation where more than two second set packets are processed, that is, for each of those packets, they are analyzed to determine if an irregularity exists, and the data from those packets is also used to update the sliding window of n packets corresponding to each source and destination between which each packet is communicated.

[0048] FIG. 8 illustrates a flow chart of the predictor methodology introduced above and which is therefore preferably used in connection with method 40 of FIG. 7. Indeed, for sake of cross-reference between FIGS. 8 and 7, the

predictor steps of FIG. 8 may be considered to occur as part of various steps from method 40 of FIG. 7. Thus, by way of example, FIG. 8 illustrates steps 40 and 42 from FIG. 7 and which are shown generally with an elaboration of various sub-steps as described below. Further, while the illustration of FIG. 8 is directed to steps 42 and 44, one skilled in the art should appreciate that the discussion below with respect to newly-received and processed as described above, also receives an additional packet(s) and updates the corresponding sliding window of n packet data, where from the following it will become apparent that from each update additional parameters are predicted and those predictions form the basis for detecting irregularities in each of steps 44, 46, 50, and 51. This processing is now further explored below in the various sub-steps of step 42.

[0049] In sub-step 42a, the network monitor NM_x detects a network packet of the first set, that is, a packet having a source and destination address matching the source and destination (or multiple sources and destinations) that are specified, by a rule(s), to be monitored. In sub-step 42b, the TCP sequence number ("SN") and the time of arrival ("TOA") of the detected TCP packet are stored, such as in a buffer located in flow store 32. In sub-step 42c, a difference, or offset, is determined and stored for the TCP SN parameter of sub-step 42b relative to the immediately-preceding packet that was received for the same source and destination. In other words, a TCP SN offset is determined by subtracting from the TCP SN of the presently-detected packet the TCP SN of the immediately-preceding packet that was received for the same source and destination. For the sake of reference, this offset value is hereafter referred to as $\Delta(\text{TCP SN})$, and in the preferred embodiment it is stored in flow store 32. Further, because sub-step 42c is repeated for each of the n packets, then in the preferred embodiment the TCP SN is stored and the value $\Delta(\text{TCP SN})$ is computed for each of the n packets (although the first packet in the n packets may not provide a value of $\Delta(\text{TCP SN})$ because there is no immediately-preceding packet from which to make the determination of the then-present packet relative to the immediately-preceding packet).

[0050] In sub-step 42d, a difference, or offset, is determined for the TOA parameter of sub-step 42b relative to the immediately-preceding packet that was received for the same source and destination. However, unlike the operation with respect to $\Delta(\text{TCP SN})$, preferably there is no storage of multiple offset values of TOA for the multiple ones of the n packets. Instead, in the preferred embodiment and in order to reduce complexity and storage requirements, for each packet received, first the difference between the actual TOA of the received packet and the stored value of the immediately-previous packet (of the same set) is determined, and that difference, hereafter referred to as $\Delta(\text{TOA})$, is stored to a single location. Once that value of $\Delta(\text{TOA})$ is stored, then the actual TOA of the present packet is also stored to single storage location, thereby overwriting the actual TOA of any previous packet. Thus, when a next packet in the n packets is received, sub-step 42d again operates in this manner, thereby determining the $\Delta(\text{TOA})$ as between that packet and the immediately-preceding packet, and then overwriting the previous value of $\Delta(\text{TOA})$ with a new value of $\Delta(\text{TOA})$ and also overwriting the previously-stored actual TOA of the immediately-preceding packet with the actual TOA value of the then-processed packet.

[0051] Sub-steps 42a through 42d continue for a total of n packets in the first set, where such a repetition is shown by the condition sub-step 42e. In other words, until n packets are processed through sub-steps 42a through 42d, then those sub-steps are repeated, thereby storing the value $\Delta(\text{TCP SN})$ for each of the n packets relative to the packet that preceded it. However, with respect to the TOA, only the TOA of the immediately-preceding packet is saved and it is overwritten once another packet is detected and the value of $\Delta(\text{TOA})$ with respect to that packet is determined and stored. Once n packets are processed in this manner, step 42e causes the flow to continue to sub-step 42f.

[0052] Sub-step 42f predicts an offset to the TCP SN for a future packet having the same source and destination as the n packets already detected. In the preferred embodiment, in order to provide the prediction of the offset with respect to the next TCP SN, first step 42f reviews the n stored values of $\Delta(\text{TCP SN})$ and chooses the most common stored value as the predicted TCP SN offset value for future packets. For sake of reference, the predicted offset value is hereafter referred to as $\text{PO}_{\Delta}(\text{TCP SN})$. Note also that one skilled in the art may implement alternative manners of processing the n stored values of $\Delta(\text{TCP SN})$ in order to provide the predicted value of $\text{PO}_{\Delta}(\text{TCP SN})$. In any event, once $\text{PO}_{\Delta}(\text{TCP SN})$ is determined, method 40 continues to step 44.

[0053] Looking to step 44 in FIG. 8, recall first from FIG. 7 that step 44 detects another TCP packet between the step 42 source and destination, and then step 44 determines whether an irregularity exists in the TCP sequence number of the new packet as compared to past TCP sequence number history. According to the preferred predictor methodology, the determination of an irregularity is made using the predicted offset value $\text{PO}_{\Delta}(\text{TCP SN})$ and the last-determined offset time of arrival value, $\Delta(\text{TOA})$, from sub-step 42d by determining predicted values with respect to the newly-detected packet. Specifically, the offset value of $\text{PO}_{\Delta}(\text{TCP SN})$ is added to the TCP SN of the immediately-preceding received packet in the first set, that is, the previous TCP SN is increased by the offset value it is predicted to change based on step 42f. Also, the value of $\Delta(\text{TOA})$ is added to the TOA of the immediately-preceding received packet in the first set, that is, the TOA of the immediately-previous first set packet is increased by the amount of time that occurred between that previous packet and the packet that preceded it. These two totals, therefore, represent what are a predicted values for the TCP SN and TOA of the packet detected in step 44. Given these expected values, the actual detected values of the packet detected in step 44 are compared to the predicted values, and a measure of deviation may be determined between the two. If the deviation, for either TCP SN or TOA, between the actual value and the predicted value is beyond a threshold, then an irregularity is deemed to occur and, as described above with respect to FIG. 7, method 40 continues to step 46. To the contrary, if the deviation is equal to or less than the threshold, then no irregularity is deemed to occur and, as described above with respect to FIG. 7, method 40 returns to step 44 to detect a next TCP packet in the first set and to process that step as just described. Note also now that the threshold may be adjusted by one skilled in the art based on various considerations, including the observations made in FIGS. 2 through 4.

[0054] Before proceeding recall also that step 44 in FIG. 7 concludes with an update of the data based on n packets.

Given the sub-steps shown in FIG. 8, one skilled in the art should now appreciate that this update operates in the same general manner as for the first n packets in step 42, but for step 44 the packet at issue is the step 44 detected packet. Thus, for step 44, the offset value $\Delta(\text{TCP SN})$ is determined with respect to the step 44 packet and in view of the immediately-preceding packet from step 42; particularly, $\Delta(\text{TCP SN})$ is determined by subtracting from the TCP SN of the presently-detected step 44 packet the TCP SN of the immediately-preceding packet that was received for the same source and destination in step 42. Also, the step 44 value $\Delta(\text{TCP SN})$ is stored while the oldest value of the n stored values of $\Delta(\text{TCP SN})$, as corresponding to the oldest received first set packet in a group of n packets, is discarded (e.g., overwritten). As a result, the determination of the predicted offset value, $\text{PO}_{\Delta}(\text{TCP SN})$, also may be updated given the newest value of $\Delta(\text{TCP SN})$ and the removal of the oldest value of $\Delta(\text{TCP SN})$. Also as part of the step 44 update, a new value of $\Delta(\text{TOA})$ is determined, here by comparing the TOA of the step 44 packet with the last received packet in step 42. Again, the new value of $\Delta(\text{TOA})$ is stored, thereby overwriting the actual TOA of any previous packet. Lastly, note that these same update principles, as applying to $\Delta(\text{TCP SN})$ and $\Delta(\text{TOA})$, also apply to steps 46, 50, and 51, that is, each indicated update is preferably achieved in the same manner with respect to these two parameters, where in each instance the most recently detected packet, for either the corresponding first or second set, is used to adjust $\Delta(\text{TCP SN})$ so that is reflective of the last n received packets as between the same source and destination as the previous packets, and similarly $\Delta(\text{TOA})$ is adjusted as a comparison between the presently received packet and its immediate predecessor of the same set.

[0055] Having described the application of the preferred predictor methodology with respect to step 44, as mentioned above note that a comparable analysis is also preferably performed with respect to steps 46, 50, and 51 in FIG. 7. For example, recall from FIG. 7 that step 46 detects whether an irregularity occurs for an additional packet in the first set only once step 44 detects a first TCP packet in the first set with an irregular deviation in its throughput (e.g., as detected based on TCP SN and TOA). In the case where step 46 detects multiple packets, then it updates and uses the most current value of $\text{PO}_{\Delta}(\text{TCP SN})$ as relating to the immediately-preceding n packets in the first set. Thus, for each received packet its actual TCP SN is compared to the sum of the TCP SN of the immediately-preceding packet plus the then-updated value of $\text{PO}_{\Delta}(\text{TCP SN})$. Additionally, a time of arrival analysis is made for each packet, where the time of arrival of the then-detected packet is compared to the predicted time of arrival based on the $\Delta(\text{TOA})$ value from the immediately-preceding packet. Thus, the value of $\Delta(\text{TOA})$ is also changed for each successive packet. In other words, each time a first set packet is detected, the TOA of the immediately-preceding first set packet is subtracted from the TOA of the then-received first set packet, and the difference is stored as $\Delta(\text{TOA})$; that difference is then added to the TOA of the then-received first set packet to predict the TOA of the next received first set packet, and when that next received first set packet is received, its TOA is compared to the immediately-previous predicted TOA. Again, a deviation beyond a threshold gives rise to some type of corrective/

warning response. Further, the preceding as applied to step 46 will also be readily applied by one skilled in the art to steps 50 and 51.

[0056] An additional observation is also contemplated with respect to the TOA and $\Delta(\text{TOA})$ values included within the present inventive scope. Particularly, the preceding describes the use of only two successive packets to determine those two values with respect to any of steps 42, 44, 46, 50, and 51. However, note that the changing of $\Delta(\text{TOA})$ immediately after an irregular packet is received may cause difficulty in detecting further TOA irregularities. Specifically, if the TOA of the irregular packet is used to overwrite a previous TOA and also to determine the next $\Delta(\text{TOA})$, then that next value of $\Delta(\text{TOA})$ may itself be irregular because it is based on an irregular packet. As such, the next packet may arrive at a regular timing, but its timing may be misconstrued as irregular because it will be compared to $\Delta(\text{TOA})$ as derived from the previously-received irregular packet. Accordingly, further contemplated within the present inventive scope are other manners of determining $\Delta(\text{TOA})$. As one example, rather than using only the TOA from two successive packets, the TOA of a number of packets greater than two may be compared against one another to determine a $\Delta(\text{TOA})$, such as via the average of the difference between the TOAs of each of the previously-received packets.

[0057] Having described the application of the preferred predictor methodology with respect to steps 44 and 46, note that a comparable analysis is also preferably performed with respect to steps 50 and 51, but as relating to the second set. Therefore, when a new second set packet is detected, the value of $\text{PO}_{\Delta}(\text{TCP SN})$ comparable to that from sub-step 42f is added to the TCP SN of the immediately-preceding second set packet and the result is compared to the actual TCP SN of the new second set packet. Similarly, when the new second set packet is detected, its actual TOA is compared to the sum of the TOA of the immediately-preceding second set packet with the difference in TOA between that immediately-preceding second packet and the second set packet which preceded it. The comparisons provide an indication of a deviation, if any, between the predicted TCP SN and TOA values and the actual values, and for the sake of step 50 or 51 those deviations are compared to the step 50 or 51 threshold as described above, with a response taken based on how the deviations, if any, compare to the step 50 threshold.

[0058] From the above illustrations and description, one skilled in the art should appreciate that the preferred embodiments provide a manner of monitoring the flow of network packets, with a particular benefit having been shown for monitoring TCP packets based on their TCP sequence numbers as well as their times of arrival. By the preferential analysis of flows in terms of TCP sequence numbers and times of arrival, a measure of throughput is realized, which as evidenced in FIGS. 2 through 4 depicts the change of TCP sequence numbers as it relates to the times of arrival. The embodiments provide numerous benefits over the prior art. As one example of a benefit, unlike the prior art active systems, the preferred embodiments do not send special test packets (e.g., PING) and do not introduce additional test traffic onto the route of the regular data packets. Thus, in this sense the preferred embodiments do not impact the network performance. As another example of a benefit, as compared to the traditional passive measure-

ment mechanisms in which the IP packets (or portions thereof) are routinely stored and then later studied using an off-line analysis of historical data, the preferred embodiments use real-time RTFM meters to retrieve packet information from packets as they are incurred during actual real-time traffic flow and to store portions of those packets; this stored data is also preferably used in a real-time or near real-time manner, such as on the order of less than one second of when the packet sequence numbers and times of arrival are collected, to evaluate the regularity of the underlying network and to take corrective action(s), if necessary. As still another example of a benefit, prior art MIBs provide single point analysis directed to the traffic flow at the location of the MIB. In contrast, the preferred embodiments contemplate analyzing real-time collected packet information from multiple points in the network and they are not constrained to the hardware type or manufacturer of each router. Still further, the preferred embodiments also provide numerous benefits inherited from the following intrinsic advantages of RTFM meter architecture: (1) all data on the IP network is recordable; (2) functional even when the given network elements are incapable of flow monitoring; (3) independence from physical and data layer technologies; and (4) good performance even during fault conditions. In all events, the preceding as well as other benefits should be appreciated by one skilled in the art. As another benefit, while the preferred embodiments are implemented with respect to TCP sequence numbers, the inventive teachings likewise may apply to another selected field(s) in a packet that readily distinguishes that packet from past and future packets in the sequence, where based on past packets the future packets can be monitored given an expectation of the value of the selected field for each of the future packets. As a final benefit, while the preferred embodiments are particularly advantageous in IP networks, they also may be applied to numerous other networks as well. In addition, while the present embodiments have been described in detail, various substitutions, modifications or alterations could be made to the descriptions set forth above without departing from the inventive scope which is defined by the following claims.

1. A network monitoring system, comprising:

storage circuitry for storing network packet information, wherein the network packet information includes a predicted identifier; and

at least one monitoring circuit coupled to a network along which network traffic flows in a form of packets, the at least one monitoring circuit programmed to perform the steps of:

receiving a packet communicated along the network;

determining whether the received packet is communicated between a source and destination in a first set of network nodes, wherein each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation

between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

2. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the identifier deviation to a threshold; and

responsive to the identifier deviation exceeding a threshold, indicating an alarm.

3. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the identifier deviation to a threshold; and

responsive to the identifier deviation exceeding a threshold, adjusting packet communications between the source and the destination.

4. The system of claim 3:

wherein the at least one monitoring circuit monitors packets between core routers; and

wherein the adjusting step adjusts packet communications at an edge router.

5. The system of claim 1 wherein the packet identifier comprises a sequence number.

6. The system of claim 1:

wherein the packet comprises a TCP packet; and

wherein the packet identifier comprises a TCP sequence number.

7. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform, prior to the comparing step, the steps of:

receiving a plurality of packets communicated along the network;

for each packet in the plurality of packets:

determining whether the received packet is communicated between a source and destination in the first set of network nodes; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, storing in the storage circuitry the packet identifier of the received packet.

8. The system of claim 7 wherein the at least one monitoring circuit is programmed to further perform the step of determining the predicted identifier in response to a plurality of packet identifiers stored by the storing step corresponding to a plurality of packets in the first set.

9. The system of claim 8 wherein the at least one monitoring circuit is programmed to further perform the step of determining the predicted identifier as a most common one of the plurality of packet identifiers.

10. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the step of, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing a time of arrival of the received packet to a predicted time of arrival to determine a time deviation between the time of arrival and the predicted time of arrival.

11. The system of claim 10 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the time deviation to a threshold; and

responsive to the time deviation exceeding a threshold, indicating an alarm.

12. The system of claim 10 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the time deviation to a threshold; and

responsive to the time deviation exceeding a threshold, adjusting packet communications between the source and the destination.

13. The system of claim 10 wherein the at least one monitoring circuit is programmed to further perform, prior to the step of comparing a time of arrival, the step of determining the predicted time of arrival in response to a difference in time between two successively-received packets between the source and the destination in the first set of network nodes.

14. The system of claim 1 wherein the first set consists of a single source and a single destination node.

15. The system of claim 1 wherein the first set comprises a plurality of source nodes in a first physical sub-network and a plurality of destination nodes in a second physical sub-network that differs from the first physical sub-network.

16. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the steps of, responsive to determining that the identifier deviation exceeds a threshold:

receiving an additional packet communicated along the network;

determining whether the additional packet is communicated between the source and the destination in the first set of network nodes; and

responsive to determining the additional packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the additional packet to a predicted identifier to determine an identifier deviation between the packet identifier of the additional packet and a predicted identifier.

17. The system of claim 16 wherein the at least one monitoring circuit is programmed to further perform, responsive to determining that the identifier deviation between the packet identifier of the additional packet to a predicted identifier exceeds a threshold, the steps of:

receiving a plurality of packets communicated along the network; and

determining whether each packet in the plurality of packets is communicated between a source and destination in a second set of network nodes, wherein the second set of network nodes comprises the first set of network nodes.

18. The system of claim 17:

wherein the first set consists of a first number of nodes; and

wherein the second set consists of a second number of nodes greater than the first set of nodes.

19. The system of claim 18 wherein the first set consists of a single source and a single destination node.

20. The system of claim 18:

wherein the single source node corresponds to a first physical sub-network and the single destination node corresponds to a second physical sub-network; and

wherein the second set comprises a plurality of source nodes corresponding to the first physical sub-network and a plurality of destination nodes corresponding to the second physical network.

21. The system of claim 17 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in the second set of network nodes, the step of comparing the packet identifier of each respective received packet to a predicted identifier to determine a respective identifier deviation between the packet identifier of the respective received packet and the predicted identifier.

22. The system of claim 21 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in a second set of network nodes, the steps of:

comparing the respective identifier deviation to a threshold; and

responsive to the respective identifier deviation exceeding a threshold, adjusting packet communications along the second set of network nodes.

23. The system of claim 22 wherein the step of adjusting packet communications along the second set of network nodes comprises adjusting packet communications at least one an edge router.

24. The system of claim 21 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in a second set of network nodes, the steps of:

comparing the respective identifier deviation to a threshold; and

responsive to the respective deviation exceeding a threshold, indicating an alarm.

25. The system of claim 16 wherein the at least one monitoring circuit is programmed to further perform the step of, responsive to determining the additional packet is communicated between a source and destination in the first set of network nodes, comparing a time of arrival of the additional packet to a first predicted time of arrival to determine a time deviation between the time of arrival of the additional packet and the first predicted time of arrival.

26. The system of claim 25 wherein the at least one monitoring circuit is programmed to further perform, responsive to determining that the time deviation exceeds a threshold, the steps of:

receiving a plurality of packets communicated along the network; and

determining whether each packet in the plurality of packets is communicated between a source and destination

in a second set of network nodes, wherein the second set of network nodes comprises the first set of network nodes.

27. The system of claim 26 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in a second set of network nodes, the step of comparing a time of arrival of the packet to a corresponding predicted time of arrival to determine a corresponding time deviation between the time of arrival of the packet and the corresponding predicted time of arrival.

28. The system of claim 27 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the corresponding time deviation to a threshold; and

responsive to the corresponding time deviation exceeding a threshold, indicating an alarm.

29. The system of claim 27 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the corresponding time deviation to a threshold; and

responsive to the corresponding time deviation exceeding a threshold, adjusting packet communications between at least some of the second set of nodes.

30. The system of claim 1 wherein the network comprises an internet protocol network.

31. The system of claim 1 wherein the internet protocol network comprises the global internet.

32. The system of claim 1 and further comprising a plurality of monitoring circuit, each coupled to the network, and each programmed to perform the steps of:

receiving a packet communicated along the network;

determining whether the received packet is communicated between a source and destination in a first set of network nodes, wherein each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

33. A method of monitoring a network along which network traffic flows in a form of packets, the method comprising:

storing network packet information in storing circuitry, wherein the network packet information includes a predicted identifier and wherein the storing circuitry is accessible to the network; and

operating at least one monitoring circuit coupled to the network along which network traffic flows in a form of packets, the operating step comprising the steps of:

receiving a packet communicated along the network;

determining whether the received packet is communicated between a source and destination in a first set of network nodes, wherein each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

34. The method of claim 33, wherein the operating step further comprises:

comparing the identifier deviation to a threshold; and

responsive to the identifier deviation exceeding a threshold, one of the steps from the group consisting of indicating an alarm and adjusting packet communications between the source and the destination.

35. The method of claim 33 wherein the packet identifier comprises a sequence number.

36. The method of claim 33, wherein the operating step further comprises, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing a time of arrival of the received packet to a predicted time of arrival to determine a time deviation between the time of arrival and the predicted time of arrival.

37. The method of claim 36 wherein the operating step further comprises:

comparing the time deviation to a threshold; and

responsive to the time deviation exceeding a threshold, one of the steps from the group consisting of indicating an alarm and adjusting packet communications between the source and the destination.

38. The method of claim 33 wherein the operating step further comprises, after determining that the identifier deviation between the packet identifier of the packet to a predicted identifier exceeds a threshold, the steps of:

receiving a plurality of packets communicated along the network; and

determining whether each packet in the plurality of packets is communicated between a source and destination in a second set of network nodes.

39. The method of claim 38 wherein the second set of network nodes comprises the first set of network nodes and wherein the second set consists of a number of nodes greater a number of nodes in the first set of nodes.

40. The method of claim 38 wherein the operating step further comprises, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in the second set of network nodes, the steps of:

comparing the packet identifier of each respective received packet to a predicted identifier to determine a respective identifier deviation between the packet identifier of the respective received packet and the predicted identifier;

comparing the respective identifier deviation to a threshold; and

responsive to the respective identifier deviation exceeding a threshold, one of the steps from the group consisting of indicating an alarm and adjusting packet communications between the network.

* * * * *



(51) International Patent Classification:

H04L 29/06 (2006.01) H04L 12/26 (2006.01)
H04L 12/24 (2006.01)

(21) International Application Number:

PCT/EP2011/002145

(22) International Filing Date:

28 April 2011 (28.04.2011)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant (for all designated States except US): VOIPFUTURE GMBH [DE/DE]; Wendenstrasse 4, 20097 Hamburg (DE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): SCHOLZ, Hendrik [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). WALLBAUM, Michael [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). KRUEGER, Michael [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). BASTIAN, Jan [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE).

(74) Agent: BOEHMERT & BOEHMERT; Felix Hermann, Hollerallee 32, 28209 Bremen (DE).

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

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

Published:

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

(54) Title: CORRELATION OF MEDIA PLANE AND SIGNALING PLANE OF MEDIA SERVICES IN A PACKET-SWITCHED NETWORK




Fig. 4

(57) Abstract: This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided. To correlate media streams and signaling sessions of services, the invention proposes to independently generate correlation keys in a media plane probe for monitored media streams and correlation keys for signaling sessions that are monitored by a signaling plane probe in a fashion that matching correlation keys are generated for a respective service. By identifying reports on the media streams and reports on the signaling session that contain matching correlation keys respectively, the media streams and the signaling streams are correlated in a correlation unit.

Correlation of Media Plane and Signaling Plane of Media Services in a Packet-Switched Network

FIELD OF THE INVENTION

This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided.

10 TECHNICAL BACKGROUND

VoIP Services

VoIP ("Voice over IP" – IP denoting the Internet Protocol) services may be considered to consist of a signaling plane and a media plane. On the signaling plane various protocols describe the session (call) flow in terms of involved parties, intermediary VoIP entities (i.e. VoIP proxies, routers) and the characteristics of the VoIP service (call). The media plane typically carries the media information (i.e. audio and/or video data) between the involved parties. Neither the media plane nor the signaling plane alone is sufficient to carry a VoIP service. On the signaling plane protocols like SIP (see IETF RFC 3261, "SIP: Session Initiation Protocol", available at <http://www.ietf.org>) or ITU-T recommendation H.323 (see H.323, "Packet-based multimedia communications systems", Edition 7, 2009, available at <http://www.itu.int>) are commonly used, whereas protocols like RTP (Real-time Transport Protocol, see IETF RFC 3550, "RTP: A Transport Protocol for Real-Time Applications", available at <http://www.ietf.org>), MSRP (see IETF RFC 4975, "The Message Session Relay Protocol (MSRP)", available at <http://www.ietf.org>) or ITU recommendation T.38 (see T.38, "Procedures for real-time Group 3 facsimile communication over IP networks", Edition 5 (2007) or Edition 6 (2010), available at <http://www.itu.int>) may be present on the media plane.

In contrast to the traditional PSTN (Public Switched Telephone Network) network both planes may be on different infrastructure using different protocols and even take different routes through a network.

Correlation Information

- 5 The property of disjunct signaling stream(s) and media stream(s) of a VoIP service within packet-switched networks may necessitate a correlation mechanism. This correlation may only be based on the common parameters carried on both planes of the session. Monitoring equipment used in VoIP networks may be limited to either the signaling or media plane due to the network structure. A correlation of signaling stream(s) and media stream(s) is thus
10 needed outside the VoIP network in between the monitoring components.

In RTP-based VoIP networks the data basis is rather limited as defined by the RTP protocol and the underlying transport and Internet layer protocols. An RTP stream which typically carries voice or video information from one party to another of a VoIP service is made up of plural RTP packets. The header of each RTP packet does not contain an identifier which can
15 be used to distinctively associate single RTP packets or the entire RTP stream to a VoIP service (defined by the VoIP signaling). The same way the VoIP signaling protocol does not carry information to relate a signaling session (signaling plane of the VoIP service) to RTP protocol header information.

The parties originating or receiving a VoIP service actively dictate the transport and Internet
20 layer address (IP address and port) on which they want to receive media streams. This configuration is done on the signaling plane. Using SIP as a signaling protocol, the address tuple(s) to be used on the media plane is/are defined in a so-called SDP body (see IETF RFC 3264, "An Offer/Answer Model with the Session Description Protocol (SDP)", available at <http://www.ietf.org>). The address tuple(s) - usually consisting of an IP address and port –
25 is/are the sole information to correlate media plane of a VoIP service with the signaling plane of the VoIP service, and vice versa.

Challenges in Correlation of Signaling Plane and Media Plane

Beside the limited information available for correlation, the correlation of the signaling plane and media plane imposes also numerous further challenges, especially in passive mid-point
30 monitoring scenarios. While an address tuple (IP address and port) may be enough information to identify a single RTP stream at a given point in time, multiple streams

identified by different address tuples may exist in a regular VoIP call. In fact at least two media streams are needed to allow two-way communication between the caller and callee.

It will be outlined in further detail in the following how multiple media streams may occur in a single VoIP service and how they are communicated between the parties involved (which may not always be the case).

Mid-point monitoring equipment (i.e. the signaling and media probes) is introduced into the network path which active components (e.g. VoIP phones and servers) use to communicate with each other. This passive equipment is not involved in the communication itself and has to derive all information from the network packets it observes on said network path. The correlation challenge increases as the passive equipment lacks information the active components internally have. A VoIP phone may internally assign a media port to a signaling session but never send out this information combined so that a mid-point monitoring probe could benefit from this information.

Time Constraints due to Port Reuse

The IP address and port combination transmitted inside the signaling plane of a VoIP service as well as the IP address and port observed on the media plane of the VoIP service may be reused. For exemplary purposes it can be assumed that each VoIP device has a single IP address. The port range from which the client can draw is limited to roughly 65000 ports by the underlying transport layer protocol UDP (see RFC 768, "User Datagram Protocol", available at <http://www.ietf.org>). In practice additional specifications and configuration parameters significantly further limit this port range. RFC 3550 which defines the Real-Time Transport Protocol mandates that only even port numbers may be used to transport media streams. The odd ports are reserved for Real-Time Transport Control Protocol (RTCP) packets. This standardization alone effectively halves the number of ports available for VoIP communication.

As a result a VoIP device has to re-use a certain IP address and port combination, thus eliminating a unique identification criterion. Essentially the IP address and port combination can only be assumed to be unique for a limited amount of time. The time-frame depends on the utilization of the VoIP device as well as the number of IP addresses required in an Internet Service Provider's network. If a correlation of the signaling and media plane takes place without a time criteria, it would therefore produce so-called false positives: Multiple media streams would match the requested IP address and port combination, since the signaling plane may have caused re-use of the very same IP address and/or port for another VoIP service.

Media Streams

Multiple streams may exist per direction of each VoIP service. The caller sends data to the callee and vice versa resulting in at least one media stream per direction. Desirable features, such as for example codec changes, may trigger additional streams. On the network multiple streams may use the same IP address and port tuple depending on the media protocol. For example, each RTP stream is additionally identified by a unique identifier - the Sender Synchronization Source (SSRC) – which is used to distinguish between multiple streams.

Fig. 1 shows exemplary situations where multiple media streams may be triggered. The focus is on the media plane traffic and its timing. The x-axis denotes the time basis, starting on the left side. The vertical lines in Fig. 1 denote signaling events and the bold horizontal lines are RTP streams present in the respective exemplified VoIP service example (please note that the different VoIP service examples are separated by broken horizontal lines). Only the RTP stream direction from the callee to the caller is shown per VoIP service example. The RTP stream from the caller to the callee can be assumed to be less complex (e.g. as shown in the “normal” VoIP service example).

The three vertical lines denote important events in the phase of a call. The leftmost line denotes the start timestamp. This is the time at which the caller submitted the call setup request to the network. The connected point is the timestamp at which a callee or his endpoint/handset picked up and moved the service into a fully established state. The service ends when one of the parties sends a termination request.

The most common VoIP services will be the trivial session examples where there is only one RTP stream per direction (see “normal” or “early media”). In a “normal” VoIP service flow both parties only start sending once the session is fully established on the signaling plane. More often the callee would start sending early media, i.e. when sending ring-tones or (pre-call) announcements. In “session forking” VoIP service example there may be multiple concurrent or sequential RTP streams at the beginning of a service e.g. due ring-tones being sent to multiple VoIP devices of the callee, while one is answered at the connect event

During a service media codecs may be changed (see “code change” VOIP service example) which is commonly resulting in a new media stream. Another service scenario may be the VoIP call being put on hold and picked up at a later time essentially generating multiple media streams (see “call on hold” VoIP service example). Silence suppression and other RTP features may cause multiple streams sporadically during a service (see “silence suppression” VoIP service example). Also retransmissions on the signaling plane may

desynchronize the signaling plane and media plane (see "SIP retransmissions" VoIP service examples).

Geographic Distribution of Media Plane and Signaling Plane

As indicated above, signaling plane data and media plane data of VoIP services may take different routes. In examples, where SIP is used for session setup, this is indeed the case as those sessions usually contain signaling-only components (i.e. SIP proxies, REGISTRARS). As a result passive mid-point monitoring solutions may be split geographically along the routes of the VoIP traffic.

An exemplary passive mid-point monitoring scenario for this case is shown in Fig. 2. Fig. 2 shows three geographically distributed POPs (Point of Presence) which are linked over the carrier's internal network. The signaling data flows on the signaling plane between POP A and B via the signaling-only POP C. POP C would typically host a centralized routing entity which reduces the complexity in the POPs A and B. To prevent back-hauling traffic from POP A to POP C before sending it to POP B the VoIP system is configured in a way that would allow the media to flow directly on the media plane between POPs A and B.

Typical mid-point monitoring locations would be inside each POP. Due to the layout of the network no media traffic would be visible in POP C. To allow for successful correlation of signaling and media streams by the correlation mechanism it is therefore desirable to support distributed networks. If the same session traverses multiple POPs it would be desirable for the correlation mechanism to identify this and support correlation of the session legs across multiple sites.

Limitations/Challenges in Correlation of Signaling Plane and Media Plane

The correlation may be solely based on the IP address and port used on the transport and Internet layer. On the media plane each RTP stream is destined to exactly one IP address and port on the transport layer. On the signaling plane one IP address and port is typically exchanged inside the Session Description Body of the SDP protocol during the VoIP service setup. However this IP address and port combination may not be the combination used throughout the VoIP service. In this case the above mentioned correlation mechanism fails: It fails upon searching a media stream for a specific signaling session of the VoIP service. The same way it fails to locate a signaling session of the VoIP service for a specific media stream on the media plane. The reasons for using another IP address port combination than advertised upon session setup are manifold.

Network Address Translation (NAT)

Network Address Translation (NAT, RFC 1631) was developed to solve Internet address space depletion and routing problems. In today's Internet one may distinguish between internal addresses which are often only unique locally (i.e. typically used inside company or home networks), and external addresses of global validity and uniqueness as used on the public Internet.

Fig. 3 shows an exemplary network using internal and external addresses. The internal network on the left side uses IP addresses from the so-called private space, as for example defined in IETF RFC 1918, "Address Allocation for Private Internets", available at <http://www.ietf.org>. The external network on the right side is the public Internet with globally reachable addresses. A router with embedded Network Address Translation (NAT) functionality is located in between the two networks. The router does not directly route packets in between the two connected networks, but "translates" (i.e. replaces) IP source and destination addresses within the routed IP packets on demand.

Essentially, the NAT functionality of the router hides the entire internal network (10.0.0.0/8 in this example) behind one public IP address (192.0.2.1). Every IP connection from within the internal network to the Internet will seem to originate from the sole public IP address (192.0.2.1) of the entire network.

For each node inside the internal network this imposes some limitations. Although Internet access works in general, a node in the internal network may not be reachable from the outside. This is due to the fact that IP traffic from an outside node (i.e. a node in the external network) would be destined to the public IP address (192.0.2.1), but the router with this IP address does not know at this point in time where to forward the traffic to on the internal network, given that there was no prior communication from an node inside of the internal network and the external node.

Additionally the internal node may not be aware that it is behind a router which performs the NAT. As required by the transport protocols, the internal node will use the internal IP address for communication purposes. On the VoIP signaling and media layer (i.e. session layer or application layer), the application will also use the internal IP address (i.e. 10.0.0.123) instead of the public IP address. VoIP signaling protocols (e.g. SIP or H.323) contain various IP addresses on the application layer. An intermediary NAT device (i.e. the router in Fig. 3) will not perform translation of the addresses on the session layer or application layer, but only translates IP addresses in the network layer. The internal IP addresses inside the VoIP

signaling (i.e. SIP or H.323 messages) impose a problem to the application as internal IP addresses are most likely not reachable from the public Internet.

VoIP devices involved in the processing of a certain call may detect this problem by comparing the IP addresses seen on the Internet and transport layers with the IP addresses set on the application layer. The following message is an exemplary SIP message as seen on the network and transport layer, which means that both the Internet and transport layer information (IP addresses and ports) are visible:

```
U 2010/09/09 13:57:31.925226 192.0.2.99:37682 -> 10.0.0.1:5060
```

The Internet and transport layers show that the UDP packet (U) was received from the external address 192.0.2.99 (port 37682).

The following message represents the same exemplary SIP message as seen on the application layer, which means that both the Internet and transport layer information (IP addresses and ports) as well as the application layer (SIP message contents) are visible:

```
INVITE sip:echo@example.com SIP/2.0
Via: SIP/2.0/UDP
192.168.1.216;branch=z9hG4bK3BB64874;rport=37682;received=192.0.2.99
CSeq: 7726 INVITE
To: <sip:echo@example.com>
Content-Type: application/sdp
From: "John Doe" <sip:jdoe@example.com>;tag=693C7725
Call-ID: 1341723561@192.168.1.216
Subject: sip:jdoe@example.com
Content-Length: 230
User-Agent: kphone/4.2
Contact: "John Doe" <sip:jdoe@192.168.1.216;transport=udp>
v=0
o=username 0 0 IN IP4 192.168.1.216
s=The Funky Flow
c=IN IP4 192.168.1.216
t=0 0
m=audio 60942 RTP/AVP 0 97 8 3
a=rtpmap:0 PCMU/8000
a=rtpmap:3 GSM/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 iLBC/8000
a=fmtp:97 mode=30
```

In contrast to the information available at the network and transport layers claiming that the message was received from the external address 192.0.2.99 (port 37682), the application claims that the source was the internal address 192.168.1.216 (port 5060 - the port 5060 is not viewed for the source address as it is the default port).

- 5 In general the translation of the network (IP) address also applies to the port numbers which are used on the transport layer. This feature is named Port Address Translation (PAT). NAT and PAT commonly happen at the same time in the same device.

PAT becomes necessary if multiple devices attempt to use the same port. Since different devices have different IP addresses their IP address and port tuple is unique. If packet streams from both devices traverse the same NAT device, the NAT device has to translate the port of one of the two packet streams. Otherwise there would be two distinct packet streams sharing the same external IP address of the NAT device and the same port. While this may not be a problem if the destination IP address and port tuple of the destination device are different for the two packet streams, it adds overhead which the NAT device wants to prevent. By rewriting the source port of one of the packet streams the address tuples of the external IP address and ports become unique again.

However, in networks where NAT (and PAT) is performed to for example hide private networks from the outside Internet, the correlation of signaling sessions of VoIP services and the associated media plane streams is becoming complex.

20 *Multiple Communication Paths*

Numerous solutions have been developed to allow traffic of VoIP services to operate in networks with Network Address Translation (NAT). These mechanisms require the involved caller and callee to exchange information, for instance on additional transport protocols or addresses which can be used to reach the remote party. These include – but are not limited to - IPv6 (Internet Protocol Version 6, see IETF RFC 2460, “Internet Protocol, Version 6 (IPv6), Specification”), TURN (Traversal Using Relays around NAT, see IETF RFC 5766, “Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN)”) addresses or ICE (Interactive Connection Establishment, see IETF RFC 5245, “Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols”) candidates (all RFCs available at <http://www.ietf.org>).

These additional tuples of IP addresses and ports communicated in these solutions establish different data paths in media plane and signaling plane of a VoIP service and thereby break

up the previous one-to-one relation of signaling session to media stream to a given VoIP device. Previously one media stream to the given VoIP device could be distinctively identified by exactly one IP address and port. In the new scenario on the signaling plane a list of IP address and port tuples is transported on which the VoIP device may receive media thereby increasing the complexity of correlation between media plane and signaling plane of a VoIP service.

SUMMARY OF THE INVENTION

One object of the invention is to suggest correlation mechanisms that allow overcoming one or more of the limitations and challenges on the correlation of signaling plane and media plane of services.

The object is solved by the subject matter of the independent claims. Advantageous embodiments of the invention are subject to the dependent claims.

One of the various aspects of the invention is to suggest a definition of correlation keys that allow the correlation of media streams of services and the associated signaling session of the respective services based on data records (reports) that provide information on the media plane and signaling plane of the respective services. These correlation keys may be for example provided in form of hash keys. A correlation key may be associated to a given direction of the media data paths of the media streams belonging to a given service (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call). Hence, assuming that the service comprises at least two media streams of different directions, at least two correlation keys are provided on the signaling plane for later correlation of the media plane and signaling plane of the service. On the media plane one correlation key is provided per media stream.

The correlation keys for use in correlation are generated by a media plane probe and a signaling plane probe independently from each other, but the two probes generate matching correlation keys for a given media service (i.e. the signaling plane probe and the media plane probe independently each generate the same correlation key for a given direction/media data path of the media service, given that there are not any exceptional situations preventing the match of the correlation keys).

In response to the detection of the end of a signaling session, the signaling plane probe generates for this signaling session at least one correlation key for each "direction" of the service to which the signaling session belongs. Each correlation key (e.g. hash key) is generated by the signaling plane probe based on a timestamp indicative of the point in time

at which the end of the respective signaling session has been detected by the signaling plane probe and using information on respective direction of the media plane data flow (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call). The correlation keys generated for the different directions of the service corresponding to the signaling session the end of which has been detected are then collated in a signaling session report, which may be for example stored in a database or a flat file.

For the media plane, the media plane probe generates a correlation key (e.g. hash key) for each media stream. The correlation keys for the media streams are generated upon the media plane probe detecting the end of the respective media stream. The correlation key for a respective terminated media stream is generated by the media plane probe based on the direction of the respective media stream, e.g. based on the media data path of a media stream, and on a timestamp indicative of the point in time at which the respective media stream has ended. The timestamp at which the respective media stream has ended is for example the time at which the last packet of the media stream has been received by the media plane probe. The media plane probe stores the correlation key (e.g. hash key) for the media stream in a media stream report, which may be for example stored in a database or exported as a flat file.

Another aspect of the invention is the use of the correlation keys comprised in the signaling session reports and the media stream reports to identify the media streams and signaling session belonging to a single media service. As indicated above, the correlation keys generated by the signaling plane probe and the media plane probe should be matching each other (i.e. the correlation keys of a given signaling session report should be found in the media stream reports, respectively, the correlation key of a media stream report should be found in one of the signaling session reports), provided that signaling session and related media streams are terminated nearly at the same point in time, i.e. within a given timeframe. Hence, a correlation unit can identify which media streams belong to which signaling session (or vice versa) and may therefore generate a service record that summarizes a service based on information contained in the media stream reports and information contained in the signaling session report of the respective service.

According to one exemplary embodiment of the invention, a method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network is provided. It is assumed for exemplary purposes that one service has one signaling session.

In this method a signaling plane probe generates at least one correlation key for each direction of a service based on the information of the signaling session of the respective

service monitored at the signaling plane probe. As noted above the direction may be equivalent to a respective media data path of a media stream belonging to a given service. The correlation keys for a respective service may be generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service. Furthermore, the signaling plane probe stores a signaling session report for each signaling session in a database. Each of the signaling session reports comprises at least two correlation keys generated for the respective directions of the service.

- 5
- 10 Moreover, in this method a media plane probe also generates a correlation key for each media stream monitored by the media plane probe. The media plane probe generates the correlation key using a timestamp indicative of the point in time of the end of the respective media stream as detected by the media plane probe and using information on the media data path of the respective media stream. The media plane probe stores a media stream report
- 15 for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

A correlation unit further correlates the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

- 20 In a more detailed exemplary embodiment of the invention, the generation of the respective correlation keys for the respective media streams by the media plane probe is based on the media data path information of the respective media stream and the timestamp indicative of the point in time at which the respective media stream ended.

- 25 According to another exemplary embodiment of the invention, information on the direction of the respective service – based on which the correlation keys are generated by signaling plane probe for a respective service – is information on one or more media data paths of media streams of the service learned from messages of the signaling session of the service. In other words, the direction of the service may be for example given by the media data path which is yielding a given direction. For example in a simple VoIP call between two peer devices, there are two directions: caller-to-callee and callee-to-caller as noted above.

- 30 In a further embodiment of the invention, the signaling plane probe monitors signaling messages of the respective signaling sessions of the services, and detects media streams belonging to each service associated to a respective monitored signaling session based on the signaling messages of the respective signaling sessions. The correlation keys for a respective one of the signaling sessions may be generated by the signaling plane probe

based on information on the media data path of the media streams belonging to the service associated to the respective one of the signaling sessions and the timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe. For example, in one possible implementation information on a
5 media data path of a media stream is a destination IP address and port number of the media stream.

In another embodiment, the timestamps used for the generation of the correlation keys by the signaling plane probe and the media plane probe could be for example rounded (e.g. in a configurable or fixed interval) before the generation of the correlation keys, in order to ensure
10 that timing differences in the events monitored in the media plane and signaling plane of the services do not lead to generation of non-matching correlation keys for a given service. Alternatively, in another exemplary implementation, the timestamps – by default – have a granularity that is higher than one second, e.g. a granularity in the range of one to ten seconds.

15 In a further embodiment, instead of storing a single correlation key per direction of a signaling session or per media stream, additional correlation keys may be generated and stored. The input for these additional correlation keys are the same IP address and port information, but a different timestamp. For example, the signaling plane probe could calculate an additional correlation key using a timestamp earlier than the timestamp
20 indicative of the end of the respective signaling session.

In case of using a rounded timestamp this earlier timestamp may be the next earlier rounded timestamp before the rounded timestamp indicative of the end of the respective signaling session. In one exemplary implementation, the signaling plane probe could generate two timestamps from a (non-rounded) timestamp indicative of the end of the respective signaling
25 session which are each used together with the same IP address and port information to generate two respective correlation keys for a given direction: One of these two generated timestamps used for calculation of the correlation keys may be the timestamp indicative of the end of the respective signaling session rounded to the nearest lower value according to the rounding interval, while the other of these two generated timestamps may be the
30 timestamp indicative of the end of the respective signaling session rounded to the nearest higher value according to the rounding interval.

All correlation keys that are generated by the signaling plane probe for a given direction may be stored as part of the signaling session records.

In a similar fashion as described above for the signaling plane probe and the directions indicated in the signaling, also the media plane probe may optionally generate more than one correlation key for a given media data path/media stream. In case of using rounded timestamps also the media plane probe could generate two timestamps from a (non-
5 rounded) timestamp indicative of the end of the respective media stream which are each used together with the same IP address and port information of the media stream to generate two respective correlation keys. One of the two timestamps generated by the media plane probe may be the timestamp indicative of the end of the respective media stream rounded to the nearest lower value according to the rounding interval, while the other of
10 these two generated timestamps may be the timestamp indicative of the end of the respective media stream rounded to the nearest higher value according to the rounding interval. All correlation keys that are generated by the media plane probe for a given media stream may be stored as part of the media stream records.

In case of rounding timestamps in the media plane probe and the signaling plane probe, the
15 rounding interval in both probes should be identical. The gain of having both probes generating additional correlation keys and correlating a multitude of correlation keys may not reflect in the increased processing needs.

In a further exemplary embodiment of the invention, the correlation of the media streams and the signaling session of the respective services is performed by a correlation unit periodically
20 or in response to a query.

Another exemplary embodiment of the invention is related to a method for correlating media streams and signaling sessions of services. According to this method a correlation unit correlates the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and
25 signaling session reports of the signaling sessions to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services. For correlating media streams and the signaling sessions, the correlation unit identifies one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media
30 stream reports and signaling session reports. Further, the correlation unit generates a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained in the signaling session report of said at least one service.

In one exemplary implementation, these steps are performed by a correlation unit periodically for all signaling session reports and media stream reports generated within a configurable time interval, or in response to a request.

5 In case the correlation unit is unable to match one or more media streams to said signaling session report, the correlation unit may optionally generate a service record without quality information of the media streams of the service or may optionally generate a corresponding indication of the user/operator of the correlation unit.

10 A service report generated for a service by the correlation unit may comprise – according to one exemplary implementation – information on the service and the quality information of the media streams of the service.

15 In one exemplary implementation, the correlation keys provided in a respective signaling session report may have been generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the corresponding service has been detected and using information on the media streams of the service. Similarly, the correlation keys provided in a respective media stream report may have been generated using a timestamp indicative of the point in time at which the respective media stream ended and using information on the media data path of the media stream.

20 In one further exemplary implementation, the correlation unit first attempts a correlation of media stream reports and signaling session reports using the correlation keys that have been generated based on the timestamps indicative of the end of the signaling session, respectively, the end of the media stream(s). In case the correlation is unsuccessful, the correlation unit may further use optionally available additional correlation keys contained in the signaling session reports and/or media plane reports which are based on other timestamps, e.g. before and after the actual end of the respective signaling session or media stream(s), respectively. This layered approach allows to first conduct the correlation using the correlation keys which are more likely to yield a result and as a fallback conduct the correlation using a wider timeframe.

30 According to a more detailed exemplary embodiment of the invention, the correlation of the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions by the correlation unit comprises matching the respective correlation keys of the media stream reports of media streams to the respective correlation keys of the signaling session reports, or vice versa. A match of the correlation keys of a one of the media stream reports and one of the signaling session reports (i.e. the

two correlation keys are identical) yields that the media stream of said one media stream report belongs to the same (media) service to which the signaling session of said one signaling session report belongs.

5 In another exemplary embodiment of the invention, for correlating the media streams and the signaling sessions of one or more services by the correlation unit, same may obtain from a signaling session report correlation keys identifying the media streams belonging to the same service as the signaling session. Further the correlation unit could then identify those media stream reports which include one of the correlation keys obtained from the signaling session report.

10 Alternatively, in another exemplary embodiment, the correlation of the media streams and the signaling sessions of one or more services could also be implemented by the correlation unit obtaining from a media stream report a correlation key identifying the media stream, and further the correlation unit identifying a signaling session report among the signaling session reports which includes the same correlation key obtained from the signaling session report.

15 According to one exemplary embodiment of the invention, a method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network is provided. It is assumed for exemplary purposes that one service has one signaling session.

20 Another embodiment of the invention is related to a passive monitoring system for correlating media streams and signaling sessions of services in a packet-switched network. The passive monitoring system comprises a signaling plane probe that generates at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe. As noted above, there may be one or more correlation keys for each direction of the service. The correlation keys for a
25 respective service are generated by the signaling plane probe using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and further using information on the direction of the respective service (e.g. information on the media data path of the given direction, such as a destination IP address and UDP port number). The signaling plane probe
30 stores a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

Furthermore, the passive monitoring system also comprises a media plane probe for generating a correlation key for each media stream monitored by the media plane probe. The

correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream. The media plane probe stores a media stream report for each media stream in a database. Each media stream report comprises at least the correlation key
5 generated for the respective media.

The passive monitoring system further comprises a correlation unit for correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports. The correlation unit may for example use one of the different correlation methods described herein for this purpose.

- 10 Another embodiment of the invention is related to the design of the signaling plane probe. According to this embodiment of the invention, the signaling plane probe comprises a network interface card for receiving packets of signaling sessions of multiple media services and a monitoring unit for monitoring the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a
15 respective one of the signaling sessions. The monitoring unit of the signaling plane probe is further capable of detecting the point in time at which a respective signaling session of the respective service terminates.

- The signaling plane probe also comprises a processing unit for generating at least one correlation key for each direction of a service based on a timestamp indicative of the point in
20 time at which the end of the respective signaling session of the respective service has been detected by the monitoring unit and based on information on the direction of the respective service. Moreover, the monitoring unit stores a signaling session report for each signaling session in a database, wherein each signaling session report comprises at least two correlation keys generated for the respective directions of the service.

- 25 Another embodiment of the invention is related to a media plane probe which comprises a network interface card for receiving packets of media streams of multiple media services, and a monitoring unit for monitoring the media streams based on the packets of the media streams received by the media plane probe. Thereby the media plane probe obtains information on a media data path a respective one of the media stream. Further, the
30 monitoring unit detects the point in time at which a respective media stream of the respective service terminates, and a processing unit of the media plane probe generates a correlation key for each media stream monitored by the monitoring unit. This correlation key may be for example generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective
35 media stream.

The media plane probe's monitoring unit further stores a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

5 The invention according to another embodiment is related to a correlation unit capable of correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network. The correlation unit comprises a database access unit for accessing a database comprising signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services. The signaling session reports each comprise at least one correlation key for each direction of a
10 respective service, and the media stream reports each comprise one correlation key indicating the direction of the media stream. A processing unit of the correlation unit is used to correlate the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a
15 respective one of said one or more services. The processing unit identifies one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports. The correlation unit may further comprise a service record generation unit (e.g. implemented by a processing unit) for generating a service record for a
20 respective service based on information comprised in the media stream reports and information comprised the signaling session report of said at least one service.

In a further embodiment of the invention, the correlation unit is further adapted, e.g. by comprised respective configured means, to perform the steps of the correlation methods according to one of the different embodiments described herein. Furthermore, also the
25 signaling plane probe and the media plane probe may be adapted, e.g. by comprised respective configured means, to implement the different functions thereof as described herein.

Further embodiments of the invention relate to the implementation of the different functions of media plane probe, signaling plane probe and correlation unit in software. In this context,
30 another embodiment of the invention provides a computer-readable medium storing instructions that, when executed by a processor of a signaling plane probe, cause the signaling plane probe to receive packets of signaling sessions of multiple media services, and to monitor the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of
35 the signaling sessions. The execution of the instructions further cause the signaling plane probe to detect the point in time at which a respective signaling session of the respective

service terminates, and to generate at least one correlation key for each direction of a service based a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected and based on information on the direction of the respective service. Moreover, the executed instructions may further cause
5 the signaling plane probe to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

Another exemplary embodiment of the invention is providing a computer-readable medium storing instructions that, when executed by a processor of a media plane probe, cause the
10 media plane probe to receive packets of media streams of multiple media services, and to monitor the media streams based on the packets of the media streams received by the media plane probe to thereby obtain information on a media data path a respective one of the media streams. Moreover, the execution of the instructions cause the media plane probe to detect the point in time at which a respective media stream of the respective service
15 terminates, and to generate a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream. Furthermore, the execution of the instructions may further cause the media plane probe to store a media stream report for each
20 media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

Another embodiment of the invention is providing a computer-readable medium storing instructions that, when executed by a correlation unit, cause the correlation unit to correlate media streams and signaling sessions of services in a passive monitoring system of a
25 packet-switched network. The execution of the instructions cause the correlation unit to access a database comprising signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services. The signaling session reports each comprise at least one correlation key for each direction of a respective service, and the media stream reports each comprise one correlation key
30 indication the direction of the media stream. The execution of the instructions further cause the correlation unit to correlate the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services and to identify one or more media streams of a
35 respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports.

The execution of the instructions also cause the correlation unit to generate a service record for a respective one of said at least one service based on information comprised in the media stream reports and information comprised the signaling session report of said at least one service.

- 5 Another embodiment is providing a computer readable medium storing instruction that, when executed by the processor of the correlation unit, cause the correlation unit to perform the steps of the method for correlating media streams and signaling sessions of services according to one of the different embodiments described herein.

BRIEF DESCRIPTION OF THE FIGURES

- 10 In the following the invention is described in more detail in reference to the attached figures and drawings. Similar or corresponding details in the figures are marked with the same reference numerals.

Fig. 1 shows exemplary situations where multiple media streams may be triggered in a media service,

- 15 **Fig. 2** shows three geographically distributed POPs (Point of Presence) which are linked over the carrier's internal network,

Fig. 3 shows an exemplary network using internal and external addresses,

- Fig. 4** shows an exemplary overview on a system according to an embodiment of the invention comprising signaling plane probe, a correlation unit and a media plane probe in a packet-switched network,
20

Fig. 5 shows a flow chart of the operation of the correlation unit according to an exemplary embodiment of the invention,

Fig. 6 shows the structure of a correlation unit according to an exemplary embodiment of the invention,

- 25 **Figs. 7 & 8** shows two exemplary flow charts of the operation of a signaling plane probe according to different exemplary embodiments of the invention,

Fig. 9 shows the structure of a signaling plane probe according to an embodiment of the invention,

- Fig 10** shows the structure of a media plane probe according to an embodiment of the invention,
- Fig. 11** shows a flow chart of the operation of a media plane probe according to an embodiment of the invention,
- 5 **Fig. 12** shows a network in which locations are indicated, where signaling and media plane probes of the invention may be advantageously located according to an embodiment of the invention,
- Fig. 13** illustrates the use of rounded timestamps within a signaling plane probe and a media plane probe to highlight potential sources of error in correlation, and
- 10 **Fig. 14** illustrates the use of two rounded timestamps for generating respective correlation keys according to an exemplary embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The following paragraphs will describe various embodiments of the invention. For exemplary purposes only, most of the embodiments are outlined in relation mainly to VoIP services as
15 an example for a media service, but the invention is not limited to VoIP services. For example, the invention may also be used for other types of media services provided in packet-switched networks. The invention may be advantageously used in media services which have signaling plane and media plane separated and/or taking different routes through the packet-switched network.

20 The explanations given in the Technical Background section above are intended to better understand the often VoIP call-related exemplary embodiments described herein and should not be understood as limiting the invention to the described specific implementations of processes and functions in a packet-switched communication network.

As indicated above, the invention generally relates to methods for correlating at least one
25 media data path and a signaling session of a service in a passive monitoring system that is monitoring a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit, as well as passive monitoring system comprising same. Based on the assumption of disjunctive signaling and media components
30 of a service, the correlation mechanism suggested herein supports correlation across

multiple hosts connected via a network (e.g. a Local Area Network (LAN), Wide Area Network (WAN) or the Internet).

One of the various aspects of the invention is to suggest a definition of correlation keys that allow the correlation of media streams of services and the associated signaling session of the respective services based on data records (reports) that provide information on the media plane and signaling plane of the respective services. In the following these correlation keys are implemented as in form of hash keys for exemplary purposes. The correlation keys may be used to query database(s) to correlate media stream reports for media streams monitored by a media plane probe and a signaling session report on the signaling session provided by a signaling plane probe that belongs to the same service – it should be noted that for the proper functioning of the invention, it is not decisive whether media stream reports and signaling session reports are stored in the same or different databases, or whether there are even multiple database partitions provided. The media stream reports and signaling session reports may even be stored in files.

The correlation keys are generated by one or more media plane probes and one or more signaling plane probes that monitor a packet-switched network independently from each other. This independent calculation of the correlation keys facilitates distributed correlation spanning across multiple geographically distributed components or even across multiple vendors of software components. Furthermore, as will become more apparent from the following, for a given media data path – respectively a corresponding given direction –, the parameters that are used for generating correlation keys at the signaling plane probe(s) and the media plane probe(s) ensure that the identical correlation keys are generated for the given media data path by the signaling plane probe(s) and the media plane probe(s) to facilitate correlation. A respective correlation key may be associated to a given direction of the media data paths of the media streams belonging to a given service. For example, assuming for exemplary purposes a simple bi-directional service, e.g. VoIP call, there are two directions: caller-to-callee and callee-to-caller.

The signaling plane probe may generate a correlation key for each direction of the service. For example, the signaling plane probe generates one correlation key for the caller-to-callee direction and another correlation key for the callee-to-caller direction. However, the signaling plane probe may also generate multiple correlation keys per direction, as will be outlined below in further detail. Consequently, in case it is assumed that a service comprises at least two media streams of different directions, then at least two correlation keys are calculated by the signaling plane probe for later correlation of the media plane and signaling plane of the service. The correlation key for a given direction or media data path may be generated based on information on the direction or media data path, such as for example a tuple comprising

the destination IP address and UDP port for a media stream in the service's media plane corresponding to the given direction or media data path. In addition, the generation of the correlation key is further based on a timestamp. With respect to the generation of the correlation key by the signaling plane probe, the timestamp may for example indicate the point in time at which the signaling plane probe has detected the end of the signaling session (and thus the service). In another exemplary implementation, the timestamp used in the generation of a correlation key for a media data path of a media stream in the media plane of the service could also be the point in time of the termination (end) of the media stream in the media plane detected by the signaling plane probe.

10 The signaling plane probe may generate the correlation keys for the different endpoints (the endpoint indicate a direction of the related media data flow) of a service upon detection of the end of a signaling session and additionally or alternatively, upon detection of the end of a media stream in the service's media plane. Typically but not limited thereto, there is one endpoint per direction, e.g. the caller and callee. The callee may be represented by multiple endpoints in case of more complex media service scenarios, e.g. when session forking is used (see Fig. 1).

The end of a signaling session may be for example detected in the signaling plane probe by detecting a signaling protocol specific hang-up message, a terminal failure to connect the involved endpoints or a session timeout. Each correlation key (e.g. hash key) is then generated by the signaling plane probe based on a timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe and using information on respective direction of the media plane data flow (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call).

The direction may be for example given by a media data path of a media stream of the respective service. A media data path could be for example indicated by a tuple comprising the destination IP address and UDP port number for transmitting the media packets of a media stream of the service. The signaling plane probe may thereby generate only one single correlation key per direction or multiple correlation keys per direction. In one exemplary embodiment of the invention, the signaling plane probe generates one correlation key per direction, where the correlation key for a respective direction is based on a media data path of a media stream. This media stream of which the media data path information is used to generate the correlation key should be existent at the end of the signaling session (according to the session signaling monitored by the signaling plane probe), in order to be able to match the correlation key to a corresponding correlation key in a media stream report generated by a media plane probe. The media data path(s) of a service and optionally their

existence (e.g. start time and termination time) could be for example learned by the signaling plane probe from the messages of the signaling session of the service.

The correlation keys generated for the different directions of the service corresponding to the signaling session the end of which has been detected are then comprised in a signaling session report, which may be for example stored in a database or (flat) file.

Alternatively, the signaling plane probe may also generate a correlation key for each media stream in the media plane of the service. In this case it may be assumed that the signaling plane traffic, i.e. the packets of the signaling protocol of the signaling plane, allow the signaling plane probe to detect the termination (i.e. end) of respective media streams in the media plane. In this case, the detection of the end of a media stream could trigger the generation of a correlation key for the respective media stream based on information on its media data path and a timestamp indicating the point in time at which the signaling plane probe detected the termination of the media stream.

Further, in another alternative implementation, the detection of the end of a signaling session, respectively, service as well as the end of a media stream as detected by a signaling plane probe could be used to trigger the generation of correlation keys in the signaling plane probe. For example, during the ongoing service the detection of the termination of a media stream in the media plane as learned from the signaling plane traffic could trigger the signaling plane probe to generate and include respective correlation keys for the media streams terminated during the ongoing session into the signaling session report of the service, while the detection of the end of the signaling session, and thus the service, could be used as a trigger for generating correlation keys for the respective media streams being "alive" in the media plane of the service at the signaling plane probe and their inclusion into the signaling session report of the service.

In the media plane, the media plane probe generates a correlation key (e.g. hash key) for each media stream. The correlation keys for the media streams are generated upon the media plane probe detecting the end of the respective media stream. The correlation key for a respective terminated media stream is calculated by the media plane probe based on the direction of the respective media stream, e.g. based on the media data path of a media stream, and on a timestamp indicative of the point in time at which the respective media stream ended. This point in time is for example the point in time at which the last packet of the media stream has been received at the media plane probe. In this latter case, the media plane probe may thus maintain a timestamp indicating the time of reception of the last packet of the media stream at the media plane probe.

The timestamp used for the generation of the correlation key may indicate another point in time than the point in time at which the media plane probe detects the end of the media stream. For example, in a packet switched network, there may be interruption in the packet stream of the media stream so that there may be a timer that measures the time since the last reception of a packet of the media stream. In case the time since the last reception of a packet of the media stream is exceeding a threshold (e.g. the timer "expires") – for example after 15 or 30 seconds, or the like – then the media plane probe detects the end of the media stream. The timestamp used in the correlation key generation may nevertheless be the point in time at which the last packet of the media stream has been received. The media plane probe stores the correlation key (e.g. hash key) for the media stream in a media stream report, which may be for example stored in a database or exported as a flat file.

Furthermore the timestamp used for the generation of a correlation key may be rounded. This counters differences in the time basis (local time) of multiple signaling and media probes as slightly different timestamps still yield the same rounded timestamp. Additional correlation keys may be generated with different timestamps, e.g. indicative of the – optionally rounded – end of a respective signaling session and/or media stream or a timeframe before or thereafter. These additional correlation keys may be stored in the respective signaling session reports and/or media session reports.

A media stream may also be interrupted due to codec features such as Voice Activity Detection (VAD). If a sender of a media stream is mute the codec may not encode any data resulting in an interruption of the media stream. This is typically indicated in the packet stream transporting the media and should not be considered by the media plane probe as an indication of the end of the media stream itself. In case a timer mechanism is used to detect terminated media streams, the timer should be configured such that codec features that lead to a temporary interruption of the media stream are not triggering the termination of the media stream (according to the detection mechanism for detecting terminated media streams in the media plane probe).

Furthermore, in another exemplary embodiment, the correlation of media streams of services and the signaling sessions of the services may also be based on a compound keys used as a correlation keys. This exemplary alternative implementation may however require a more complex database structure of the media stream reports and signaling session reports, as will be outlined below.

Another further aspect of the invention is the use of the correlation keys comprised in the signaling session reports and the media stream reports to identify the media streams and signaling session belonging to the respective services. This correlation of media streams and

signaling session could be done for example periodically or on request. The correlation keys generated by the signaling plane probe and the media plane probe should be matching each other (i.e. the correlation keys of a given signaling session report should be found in the media stream reports, respectively, the correlation key of a media stream report should be found in one of the signaling session reports). Hence, a correlation unit can identify which media streams belong to which signaling session (or vice versa) based on the correlation keys contained in the media stream reports and the signaling session reports and may therefore generate a service record for a respective service based on information contained in the media stream reports and information contained the signaling session report of the service.

The correlation methods according to the different embodiments of the invention proposed herein are based on a probe monitoring the signaling plane traffic of services (i.e. the signaling plane probe) identifying new services and associated media data paths of the media plane by detecting their corresponding signaling sessions and determining for each service correlation keys allowing to identify at least the media streams of the media plane existent upon termination of the service. It should be noted that "at least the media streams" means that it is optionally possible to generate a correlation key for each media data path of the media plane learned from the signaling plane traffic of the service, in case it is possible to detect the termination of the media data paths of the media plane in the signaling plane traffic. In this case the correlation key for a respective media stream may be based on the information on the media data path of the respective media stream visible in the signaling plane of the service and the point in time at which the signaling plane probe detects the termination of the media stream in the signaling plane.

Regarding the use of the point in time of the termination of the signaling session (e.g. a signaling message in the signaling plane indicating "hang up") for the generation of correlation keys it should be noted that the termination of the signaling session also means that all existent media data paths, respectively, the corresponding media streams of the media plane are terminated. Hence, in case it is not possible to detect the termination of the media data paths of the media plane in the signaling plane traffic, it is assumed to be possible to detect at least the end of the signaling session and to thereby conclude on the termination of the respective media data paths of the media plane and to thereby facilitate correlation of the media streams that have been "alive" at the end of the service and the corresponding signaling session of the service.

As noted above, another possibility may be to use compound keys as correlation keys. A compound key is defining a set of information comprised e.g. in a given set of fields of a database entry of a signaling session report and media stream report that – together with

additional criteria on the point in time of the termination of a media stream or the service – allows to correlate media streams and signaling sessions of the services. In one exemplary embodiment, the one or more databases that are used to store the signaling session reports and the media stream reports are partitioned into partitions covering only a given amount of time (i.e. time interval) of signaling session reports and/or media stream reports, respectively. The “size”, i.e. the time interval covered by the partitions may be for example determined based on the average time interval in which port reuse occurs in the target monitored packet-switched network (e.g. the partition could cover $\frac{1}{2}$ of the average time interval in which port reuse occurs). In the given partitions it may thus be assumed that the media data path information, such as destination IP address and UDP port number are unique. In this case, the signaling session reports and media stream reports may not include separate correlation keys, but the correlation key can be generated on-demand as a compound key from the information on the media data paths contained in the signaling session reports and media stream reports.

Hence, it would be sufficient that the signaling session reports and media stream reports store information on the media data paths of the service (e.g. destination IP address and UDP port number) and a timestamp that indicates the point in time of termination of the respective media data path in the media plane, as detected by the signaling plane probe from the monitored signaling sessions and the media plane probe from the monitored media streams. As noted above, the signaling plane probe can also detect the point in time of termination of the media data path(s) in the media plane by detecting the point in time of termination of the signaling session, which is – in most practical cases – coinciding with the termination of all media streams of the service in the media plane as well.

In this case, the correlation could be realized by first selecting the database partition of the media plane reports according to a timestamp of the termination of a media data path/signaling session within a given signaling session report, and by then determining the media stream reports that comprise media stream information matching those indicated in the given signaling session report.

Fig. 4 shows a packet switched network 400 according to an exemplary embodiment of the invention that comprises a signaling plane probe 401, a correlation unit 402 and a media plane probe 403. Please note that more than one signaling plane probe 401 and media plane probe 403 may be provided. Optionally, there may be multiple correlation units 402 present, working independent of each other and not impacting the performance or correlation accuracy of any other component.

The signaling monitoring probe 401 and the media plane probe 402 are monitoring the traffic on a packet-switched network 400, which may be for example an IP based packet switched network. It is assumed for exemplary purposes that the signaling plane probe 401 and the media plane probe 403 monitor the traffic of the packet-switched network 400 at different
5 locations of the packet-switched network 400, but observe the same network traffic.

In one exemplary implementation, the signaling plane probe 401 and the media plane probe 403 are probes of a minimal non-intrusive, passive monitoring system. The respective monitoring probes 401, 403 are connected to the packet-switched network 400 via respective test access ports (TAPs). A TAP is a passive network device that mirrors network traffic of
10 the packet-switched network without interference of the original network traffic. It provides a copy of every packet sent or received on the network, by duplicating the data packets observed on the physical layer of the network. Instead of a TAP a Switch Port Analyzer (SPAN) port may be used to monitor the traffic of the packet-switched network.

In packet-switched environments, the signaling data of the signaling plane and media data of
15 the media plane is packaged into packets on the transmitter side. Packets are transported over the packet-switched network independent of each other. The receiver of the packets – identified by the destination IP address and port – may have to regroup and sometimes reorder packets in order to compile streams from single packets.

On the signaling plane multiple packets belong to one signaling session. The first packet of a
20 signaling session – usually containing a “call setup” message – starts a new signaling session (respectively a new service). At this point, the signaling plane probe 401 generates a session entry along with auxiliary data, i.e. a unique correlation key identifying the session internally (signaling session identifier). Packets following the initial signaling packet are matched to one of the currently ongoing signaling sessions (respectively services) and
25 signaling sessions managed by the signaling plane probe. Packets which do not start a new signaling session and do not belong to existing sessions may be discarded.

Matching single signaling packets to existing signaling sessions is specified in the standards which define the signaling protocol itself. Transport layer information like IP addresses and protocol-specific information (e.g. SIP Call-IDs and tags) may be used.

30 A signaling session may be torn down upon reception of a special packet – a “call hang-up” message. This message again is specific to the signaling protocol used. Upon reception of this message the signaling plane probe 401 can conclude on the termination (end) of the signaling session and generates the correlation keys for use in correlation of media plane and signaling plane of the services. Furthermore, the signaling session may be removed

from its internal storage. Additionally the signaling monitoring component may remove sessions from its internal storage, if it identifies stale sessions.

The media plane 403 of the service may for example identify the media plane data in form of RTP streams. RTP streams are – analog to the signaling sessions – made up from individual RTP packets. The transmitter of an RTP stream generates RTP packets by splitting the media data into smaller pieces. In general these packets are relatively small to reduce the delay between sender and receiver. As a result many packets are created and form a single RTP stream – typically 50 packets per second. Single RTP packets can be grouped into an RTP stream by media plane probe 403. When using RTP, the media plane probe 403 has transport layer information as well as the RTP protocol header which contains a so called Synchronization source (SSRC). The SSRC is a 32-bit numeric identifier which is contained in every single packet and deemed mandatory by the RTP. When media plane probe 403 first receives a packet which does not match an existing RTP stream in its internal storage it creates an RTP stream. Consecutive RTP packets can then be matched to the existing streams by comparing the transport layer information and SSRC value. If no packets are received for a single RTP stream for a certain time the stream is considered ended and can be removed from the internal storage. A short interruption in the packet stream should not trigger the media plane probe 403 to mark an RTP stream as ended, thus the time before a media plane probe 403 can consider an RTP stream as ended is multiple seconds and has to take features like voice activity detection into account.

The signaling plane probe 401 and the media plane probe 403 generate information about the signaling session and media streams, respectively. This generated information is comprised in signaling session reports and media stream reports, respectively. In order to track running sessions and media streams the entities 401, 403 store information internally while the signaling sessions and media streams are in progress. When a signaling session or media stream is terminated the respective monitoring component 401, 403 may generate a signaling session report with signaling information, respectively media stream report with information on the media stream. As will be outlined below in further detail, the signaling session report provides a summary of the signaling session (service) and includes the correlation keys (e.g. hash keys) for correlation. Similarly, the media stream reports summarize the media stream and also include a respective correlation key (e.g. hash key) to facilitate correlation.

Each signaling session report for a signaling session may be for example written by the signaling plane probe 401 upon completion of a signaling session and contains data which describes the signaling session it monitored. For example, in one embodiment of the invention, the data stored in the signaling session report and describing the signaling session

include – but are not limited to – a set of caller and callee specific properties which will be available per endpoint (caller or callee). These endpoint-specific properties may for example include at least one of the following:

- endpoint address/name (e.g. SIP URI),
- 5 – phone number,
- User-Agent (endpoint software and version),
- capabilities/features,
- supported media types,
- IP addresses and ports for signaling and media streams (from SDP body), and
- 10 – hash key of the media path (for example described by the destination IP address and UDP port of the expected media stream)
- signaling routing information.

Additionally, one or more of the following signaling session properties which are not endpoint-specific may also be included in the signaling plane report:

- 15 – timestamps: e.g. session start, Post-Dial-Delay (PDD), connect time, session end;
- information on the session flow (e.g. signaling message(s) with timestamps),
- session state (e.g. successful, unsuccessful, timed out, reason for failure), and
- events during session (media on hold, codec renegotiation, session forking, authentication)
- 20 In one exemplary embodiment of the invention, the signaling session report comprises all of the above noted endpoint-specific properties, and – optionally – all signaling session properties listed above, as far as they are available from the signaling protocol messages.

- Depending on the specific implementation and capabilities of the signaling plane probe, the signaling session reports may contain information about multiple endpoints and/or multiple
- 25 media streams per endpoint. In this case the session end timestamp may be available per involved endpoint, and not per session.

For example, the most trivial VoIP service is a phone call between two parties (including two media streams, one in each direction). In this case, the signaling session report could store each party as one endpoint with the above mentioned information and the media stream's destination IP address and port information. If the VoIP service would contain a video call (i.e. one audio and video stream per direction) the endpoint information may be extended to contain the destination IP address(es) and port numbers of the two media streams per direction (one IP address & port tuple for audio, one IP address & port tuple pair for video). Considering a VoIP service, the respective tuples of destination IP address and UDP port number of the media streams of the session may be for example learned by the signaling plane probe from the SDP body of SIP signaling.

Likewise, the signaling plane probe may extend the signaling session report to include additional endpoints, if multiple endpoints were part of the VoIP service. Multiple endpoints could be for example involved in a VoIP service, in case a caller tries to reach a callee which has call forwarding configured. The caller may then hear a ring-tone from the callee's landline, followed by a ring-tone from the callee's mobile phone. Finally a voicemail box picks up the session generating a third endpoint on the callee's side.

The media plane probe 403 monitors the media traffic, i.e. the individual media streams that convey the media information of the different services. The media plane probe 403 identifies a new media stream of a media service based on the network traffic it observes on the monitored packet switched network 400. The identification of a new media stream for a given service may happen shortly after this service is set up on the signaling plane. Throughout the duration of the media stream information on the media stream may be stored in an internal data record on the media stream. The media plane probe 403 may note for each stream the last RTP packet received for the respective media stream in an internal data record on the media stream in order to allow detecting the media stream's termination.

The media plane probe 403 may optionally perform a real-time analysis of the media streams monitored on the packet switched network and may generate a fine-grained media stream quality report with media stream related information. Fine-grained media stream quality reports could be for example written in predetermined time intervals, e.g. every five seconds, and could provide a summary of the media stream in the respective interval as well as quality information of the media stream in the respective interval. The information of the fine-grained media stream quality reports of a media stream could be used by the media plane probe 403 to generate a single media stream report on the media stream upon the end of the media stream is detected by the media plane probe 403. The media stream report and optionally also the fine-grained media stream quality reports of a media stream may be stored in individual databases.

The media stream report generated by the media plane probe 403 upon detecting the end of the media stream may contain data which describes the unidirectional media stream it monitored as well as correlation key that can be for example generated based on the destination IP address and UDP port number of the media stream and a timestamp that indicates the point in time at which the media stream ended/terminated. The media plane probe 403 could for example learn the destination IP address and UDP port number from the IP header encapsulating the RTP packets on the network layer. The characteristics of each media stream can be grouped into two categories. The transport and RTP header information characterize the media stream, e.g. for correlation purposes, whereas the quality information enable quality related processing of the media stream. In general the transport information is the same for all media packets processed as part of a single media stream. The stream quality information includes information covering all media packets and may be aggregated or stored on a per-packet basis.

In one exemplary embodiment of the invention, the media stream report generated by the media plane probe 403 for a given media stream may for example include one or more of the following transport media information:

- source and destination IP address,
- source and destination UDP port,
- RTP header Synchronization source (SSRC) value,
- RTP payload type, and
- correlation key(s) (e.g. hash value(s)) calculated based on destination IP address, destination UDP port and a timestamp indicating the point in time at which the of the media stream terminated (e.g. the timestamp of the last packet received for the media stream) or timestamp indicative of the timeframe before or after the termination timestamp.

The stream quality information that may be added to the media stream report by the media plane probe 403 may for example include/indicate one or more of the following:

- timestamp information (e.g. start and/or end time),
- RTP format violations,
- protocol conformance (e.g. sequence errors and/or DSCP errors),

- transport information (e.g. packet interval and jitter statistics, information on packet loss, information on duplicate packets, etc.), and
 - MOS values (e.g. per time-slice (e.g. every 5 seconds), minimum/average/maximum value for entire stream duration, etc.),
- 5 – events (e.g. comfort noise, silence suppression, DTMF tones, codec changes, marker bits).

10 In one exemplary embodiment of the invention, the media stream report comprises all of the above noted transport media information and optionally further all stream quality information listed above, as far as they are available to the media plane probe from the packets of the media streams.

In addition to “measured values” (e.g. jitter between multiple media packets) the media stream report may also contain calculated information like MOS (Mean Opinion Score) values and information about impairments experienced by the receiver of the stream.

15 As further shown in Fig. 4, a correlation unit 402 correlates the media stream reports and signaling session reports generated by the signaling plane probe 401 and the media plane probe 403 in order to identify the media streams and the signaling session that belong to a respective service. It should be noted that each service is assumed to have one signaling session, but may have one or more media streams.

Generation of Correlation Key

20 The correlation key that is used for correlation may be built from:

- the timestamp that indicates the end of the signaling session, respectively the end of a given media stream, and
- information on a media stream, such as for example the destination IP address and UDP port of the media stream.

25 The correlation of signaling plane and media plane by the correlation unit 402 is based on this compound correlation key.

30 As noted earlier, the correlation key may be for example a hash key generated based on the before-mentioned parameters. The hash function to calculate the hash key (sometimes also called hash value) from the given data should be known to all parties involved (i.e. signaling plane probe 401, and media plane probe 403 and optionally also correlation unit 402) so that

the result for a given input is the same for all parties. Therefore, the hash keys are determined by the signaling plane probe 401 and media plane probe 403 using media data path information in addition to a timestamp that indicates point in time when the signaling session, respectively, the media stream ended/terminated. The correlation unit 402 may generate additional correlation keys based on the timestamps indicative of the timeframe before or after the timestamp that indicates the point in time when the respective session ended.

The hash function to determine the hash key may be a simple concatenation of one or more of the input parameters. For example, a concatenation of a destination IP address, destination UDP port number and a timestamp may (optionally including a delimiter character in between the parameters) be used for defining the hash keys. Of course more complex mechanisms may also be used to calculate a hash key representing the signaling session identifiers and/or the media path identifiers. For example the MD5 message digest function (see IETF RFC 1321, "The MD5 Message-Digest Algorithm", 1992, available at <http://www.ietf.org>) could be used to create a valid hash key. The MD5 message digest function may be fed with a concatenation of a destination IP address, destination UDP port number and the timestamp indicating the time of termination of the signaling session, respectively, the media stream and the resulting string can be used as a correlation key. Another hash function that could be used may be selected from the SHA cryptographic hash function family (see FIPS PUB-180, "Secure Hash Standard (SHS)", 2008 available at <http://www.itl.nist.gov/fipspubs/>). Furthermore, in another embodiment, a linear hash function may be used.

As outlined with respect to Fig. 1 earlier herein, ephemeral media streams may exist at any point during a service, such as a VoIP service. The ephemeral media streams may be generated by early media streams, service forking and codec changes. As a result the start and end timestamps of media streams may not necessarily concur with the start and end time of the VoIP service itself. Yet the timestamps of the start or end of a signaling session or media stream are an important piece of information for a successful correlation, as it significantly reduces the search window and amount of data to search. In typical scenarios, it can be expected that the start time of signaling and media streams deviate significantly and are thus not ideal candidates for a correlation mechanism.

At the end of a VoIP service both caller and callee have an established signaling session and send each other media on the media plane. Once a party hangs up the session a hang-up signal is send on the signaling plane and at the same time stops sending media packets. Once the remote party receives the hang-up signal it also stops transmitting media. The timestamps of the end of the signaling session of a service and the end of the corresponding

media streams in the media plane of the service are commonly very close to each other compared to the start of a media service in signaling plane and media plane. All these events will be subject to the same network- and application-imposed delays. Still the time-frame in which these happen is typically in the few milliseconds range or in the sub-second-range on long-distance or low bandwidth networks. This makes the timestamps of the end of the signaling session and the end of the media streams a good choice for calculating a correlation key.

Nevertheless, since signaling plane probe 401 and media plane probe 403 independently determine the correlation keys based on timestamps (and even minor deviations of the timestamps in the millisecond-range could lead non-matching correlation keys), the granularity of the timestamps used for the generation of the correlation keys may be chosen to account for those deviations of events in signaling plane and media plane events (i.e. deviations in the end of the signaling session in the signaling plane and the end of the corresponding media streams in the media plane of a service). Therefore the timestamps could have a granularity for example in the range of one or more seconds or at least 250 milliseconds or 500 milliseconds. In one exemplary embodiment, UNIX timestamps (which is indicating the seconds passed since January 1, 1970, 00:00:00) are used for the calculation of the correlation keys. Alternatively, also timestamps with a high resolution (e.g. resolution of one or ten or hundred milliseconds) could be used, but the timestamps may be rounded for the calculation of the correlation keys. It should be noted that the choice of the granularity or rounding interval may also impact the performance of a database or databases storing the media stream reports and signaling session reports for correlation. The larger the granularity, respectively the rounding interval and thus the search window the more data have to be processed to match signaling sessions and media streams.

Rounding timestamps could be for example implemented by rounding the timestamp to the nearest n second interval. Assuming for example that the interval is 5 seconds the following UNIX timestamps would be rounded as exemplified in Table 1.

Unix timestamp value	rounded value (5 sec)	description
1283936131	1283936130	rounded down
1283936133	1283936135	rounded up
1283936135	1283936135	unchanged
1283936136	1283936135	rounded down
1283936139	1283936140	rounded up

Table 1

The rounding interval can be any time interval in seconds or even milliseconds whereas the specific value is a trade-off between speed and accuracy. As noted above, the larger the interval the more data may have to be processed to match signaling sessions and media streams. The smaller the search window is the more likely it is to miss a match due to misaligned timestamps or derivations due to the delays outlined above.

The best rounding interval depends on the network monitored. Reasons for increased time differences are possible derivations from a shared time source (NTP) or differences seen in signaling and media events, i.e. the delay between a call being terminated and the endpoints stopping to send media.

- 10 For a smaller search window of three seconds the rounded values are different, as for example shown in Table 2.

Unix timestamp value	rounded value (3 sec)	description
1283936131	1283936130	rounded down
1283936133	1283936133	unchanged
1283936135	1283936136	rounded up
1283936136	1283936136	unchanged
1283936139	1283936139	unchanged

Table 2

Using normalized "session end" timestamp (i.e. timestamps of low granularity or timestamps that have been rounded) may be advantageous to counter the effects (e.g. false positive hash key matches) of the re-used port numbers. Since the end of a service (and/or end of media streams) becomes part of the correlation criteria only the port numbers that are used and valid at this point in time are considered in the correlation.

Another issue that could complicate the correlation processing is that deviations in detecting the end of a media stream, respectively a signaling session at the media plane probe 403 and the signaling plane probe 401 leads to the two entities rounding the timestamps to different values. For example, the signaling plane probe 401 could round the timestamp up, while the media plane probe 403 rounds the timestamp down to the next rounding interval value due to (minor) deviations in the timestamps at the signaling plane probe 401 and media plane probe 403. In another exemplary embodiment, additional correlation keys may be therefore generated by either one or both, the signaling plane probe 401 and media plane probe 403.

The signaling plane probe 401 and/or media plane probe 403 may use timestamps indicative of the rounding intervals before or after the timestamp indicative of the point in time where a

particular signaling session or media stream ended. For example a session ended at a point in time identified by the Unix timestamp value 1283936135. The correlation unit applies the rounding mechanism based on 3 second intervals. The rounded value thus is 1283936136. The correlation unit may generate additional keys with the timeframes before or after the mentioned rounded value. In this case the timestamps used for generation of the additional keys are 1283936133 (1283936136 – 3 seconds) and 1283936139 (1283936136 + 3 seconds). Of course, especially in cases where the rounding interval is short, it may also be desirable to more rounded timestamps rounded to the next and subsequent rounding interval(s) based on which correlation keys are calculated.

As noted above, the inclusion of multiple correlation keys calculated based on different timestamps for a given direction of a signaling session and/or media stream may for example be useful to counter correlation failures due to applying different rounding of timestamps in the signaling plane probe 401 and media plane probe 403 as will be outlined in further detail based on Fig. 13 and Fig. 14 in the following. The end of the signaling session $t_{\text{signaling}}$ of a service as observed by the signaling plane probe 401 and the end of the one or more media streams t_{stream} of the service as observed in the media plane probe 403 may not perfectly match each other as shown in Fig. 13. Even in case signaling plane probe 401 and media plane probe 403 calculate correlation keys based on the same IP address and UDP port tuple, the use of different timestamps $t_{\text{signaling}}$ and t_{stream} for the calculation of the correlation key will result in the correlation keys generated by signaling plane probe 401 and media plane probe 403 not matching each other. Additionally the use of rounded timestamps may not necessarily resolve this problem, as it cannot always be assured that the timestamps $t_{\text{signaling}}$ and t_{stream} are rounded to the same value. As also highlighted in Fig. 13 it could well happen that signaling plane probe 401 rounds the timestamp $t_{\text{signaling}}$ to the next lower rounded value T_1 in accordance with the rounding interval R , while the media plane probe 403 rounds the timestamp t_{stream} to the next higher rounded value T_2 in accordance with the rounding interval R . Again, correlation keys calculated on the different rounded timestamps T_1 and T_2 would not match each other.

As noted above, one possible solution to this problem is to include additional correlation keys at least to one of the signaling session reports and media stream reports. For example, in one exemplary embodiment, the media plane probe 403 could include two correlation keys per direction to the signaling session report. For a given media data path of a media stream (e.g. IP address and UDP port tuple) observed by the signaling plane probe, the signaling

plane probe could calculate one correlation key $C_{stream,1}$ on the media data path information and a rounded timestamp T_1 , i.e. using the timestamp indicative of the end of the signaling session $t_{signaling}$ rounded to the next lower value according to the rounding interval, and the signaling plane probe could further calculate one correlation key $C_{stream,2}$ on the media data path information and a rounded timestamp T_2 , i.e. using the timestamp indicative of the end of the signaling session $t_{signaling}$ rounded to the next higher value according to the rounding interval. Both correlation keys could then be included for the direction/media data path to the signaling session report.

A further implementation may include even further correlation keys, e.g. a third correlation key that is based on the rounded timestamp T_0 , indicative of the next lower timestamp to T_1 according to the rounding interval.

The media plane probe 403 may calculate only one correlation key C_{stream} based on a rounded timestamp T_1 or T_2 , i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next lower or higher value according to the rounding interval (and rounding rules). This correlation key for the media stream would then be included to the media stream report.

The correlation of the signaling session to the media stream(s) of the service would be successful in this embodiment, since either the correlation key $C_{stream,1}$ or the correlation key $C_{stream,2}$ within the signaling session report for the given direction/media stream would correspond to the correlation key C_{stream} included in the media stream's media stream report.

Optionally, in another embodiment of the invention, also the media plane probe 403 may calculate two or more correlation keys for each media stream, one correlation key based on the rounded timestamp T_1 , i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next lower value according to the rounding interval, and the other correlation key based on the rounded timestamp T_2 , i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next higher value according to the rounding interval. Both timestamps would then be included in the media stream report for the media stream by the media plane probe. In this exemplary implementation, the correlation may check whether the two correlation keys of the service's signaling session report for the direction/media stream match the two correlation keys within

the media stream report of the service's media stream to confirm the signaling session report and media stream report belonging to the same service.

Please note that instead of or in addition to the media plane probe 403 including two or more correlation keys to the media stream reports, also the signaling plane probe 401 could
5 calculate two or more correlation keys per direction/media data path that is to be indicated in the signaling session report. Including multiple correlation keys per direction/media data path in the signaling session reports may be advantageous as it reduces the overhead information, since there are commonly less signaling session reports than media stream reports generated in the passive monitoring system.

10 Furthermore, in another exemplary implementation, the database(s) for signaling session reports and media stream reports may be further partitioned to reduce the search space for correlation. For example, the (respective) database could be partitioned in chunks of multiple hours, e.g. 4 or 6 hours or the like. The signaling session reports and media stream reports may further include as a separate element the timestamps of the end of the signaling
15 session, respectively, the end of the media streams, based on which the database(s) can be partitioned and based on which the correct partition could be chosen for correlation later. Hence, in a more advanced implementation, besides the correlation keys a service end timestamp can be used as a search criteria to limit the amount of data that are to be searched in a database significantly.

20 In another exemplary implementation, the signaling plane probe may also generate a respective correlation key for all media streams of a service, in case the end of media streams is detectable within the session signaling, as noted previously. In this case the signaling session reports could further include a timestamp that indicates the end of each media stream, in order to allow choosing the correct partition of media stream reports in the
25 database for correlation.

Correlation Mechanisms

For correlation of the media streams and the signaling sessions, different possibilities exist how the correlation by the correlation unit can be implemented. For example, according to exemplary embodiments of the invention, the correlation unit may correlate the media
30 streams and the signaling session of the respective services periodically (or proactive) or in response to a query (i.e. on demand).

For example, a proactive correlation process could be utilized by a third party entity which has access to both signaling session reports and media stream reports. Their data may be

present in one or more databases to ease access to the raw data. As noted above, the correlation aims at correlating signaling and media information of the services and at merging them into combined service reports that may be then accessed by a third party, e.g. for evaluation of the services.

- 5 The proactive correlation process periodically takes all new signaling session reports which outline – but not fully describe – the media services, e.g. the monitored VoIP sessions. For each signaling session report the correlation process draws the calculated correlation keys from the signaling session report. There are at least two correlation keys in each signaling session report, at least one per direction. Additional correlation keys may optionally exist in
- 10 the signaling session report, for instance if it describes a video call or a session with multiple endpoints. For each correlation key contained in a respective signaling session report the correlation process searches the media stream report database for the media streams with the same correlation key. If the correlation keys in the signaling session report are generated based on a timestamp indicating the end of the signaling session, every media stream of the
- 15 service that existed at the end of the service in the media plane should have the same correlation key and will be in the result set of the search. If the correlation keys in the signaling session report also contain correlation keys for media streams that have been terminated during the ongoing service, the search may also further detect those media streams in the media stream report database based on the correlation key.
- 20 If a search returned two or more media stream reports, the correlation process encountered one of two scenarios. Either both media streams belong to the signaling session (i.e. service) or one of them was returned due to a correlation key collision. Both cases may be recovered by comparing the contents of the media stream report against the signaling session report. If a media stream in the result set starts before or ends after a signaling session, it was added
- 25 to the result set due to a correlation key collision. Due to time information it cannot match the signaling session and is deleted from the result set by the correlation unit. If the media stream result set still contains multiple entries after the previous correlation key collision check, all remaining media streams (i.e. their media stream reports) belong to the media service (e.g. the VoIP session).
- 30 If a search returned no usable entry, i.e. no signaling session reports and media stream reports comprising matching correlation keys have been found, the correlation unit may optionally use additional correlation keys contained in either the signaling session reports and/or media stream reports to run another search, as explained with respect to Fig. 13 and Fig. 14 above.

The correlation process may optionally select only the most significant media stream for each direction for inclusion to the service report on the media service. The parameters to define the significance of the media streams may be for example configurable by the user/administrator. One parameter for deciding the significance of a media stream in the result set could be the duration of a media stream. For example, the service report generated for a given media service by the correlation unit could include the media stream with the longest duration, in case multiple media stream reports match a correlation key of a signaling session report. Another option would be to include in the service report the last media stream that has been alive in the media service, in case multiple media stream reports match a correlation key of a signaling session report. The correlation process may optionally add information from multiple media streams to the service report, e.g. when the service consisted of an audio and video part.

For a given signaling session (i.e. service), the correlation process may run at least twice, one time for each correlation key in a signaling session report. For example, in case the media service is a VoIP service, the correlation process may be for example performed by the correlation unit once for the caller's media stream and once for the callee's media stream, both identified by their respective correlation keys in the signaling session report.

The correlation unit combines the information contained in the media stream reports that are correlated to a respective signaling session report and the information contained in the signaling session report to provide a service report, providing an outline of the service. The content of this combined service report may be for example the sum or a subset of information of the signaling and media stream reports of a given service. Furthermore, the content of the service reports for the respective services may be configurable and/or depending on the application.

In another exemplary embodiment of the invention, the correlation unit may perform the correlation based on the media stream reports as starting points which are matched to their corresponding signaling session report. Ideally, each media stream of a media stream report should match to one signaling session report based on its correlation key. This implementation may optimize the search time in the correlation processing. The above correlation mechanism starting from the signaling session reports searches the rather large media stream report database at least twice for each signaling session. If the correlation process is based on the correlation keys within the media stream reports entries as starting point, the correlation would periodically take all new media stream reports, obtain the correlation keys thereof and search the signaling session reports for matching signaling sessions. At this point the same steps are taken as outlined above.

Another option may be to correlate media stream reports and signaling session reports on demand. An on-demand correlation employs the same mechanism as the proactive correlation mechanism described above, but instead of running periodically, e.g. when introducing new signaling session reports or media stream reports into the database(s), the on-demand correlation process runs when a user requests particular data.

In an exemplary implementation, the database may contain independent sets of signaling session reports and media stream reports. An operator could request service reports of services linked to some given search criteria, e.g. a specific user, a specific source and/destination IP address, a certain service type, services using a particular software application, codec, etc. The user could be for example specified by providing a username, phone number, the user's SIP URI, etc. or another identifier of the user that can be found in a signaling session report. The correlation unit would then search for the signaling session reports for all signaling sessions matching the particular search criteria. The search criteria may be for example provided to the correlation unit through a web application running on a browser.

For each signaling session report matching the search criteria, the correlation unit subsequently searches for the matching media stream reports as described above for the proactive correlation process. The input for this on-demand correlation is a set of signaling session reports matching the search criteria and not all signaling session reports as in the proactive correlation procedure described above. The on-demand correlation process generates for each signaling session report matching the search criteria a service report, for example by enriching the information of the signaling session report with the matching media stream report information and returns the service report for each matching signaling session report. The returned data may be for example viewed by the web application as individual service reports. Optionally the on-demand correlation process may also store the generated service reports in a database. This would save the on-demand correlation processing time in case an operator requests the same information multiple times.

Correlation Unit

Next, the correlation unit 402 as shown in Fig. 4 is discussed in further detail with respect to Fig. 5 and Fig. 6. Fig. 6 shows the structure of a correlation unit 402 according to an exemplary embodiment of the invention.

The correlation unit 402 may be reachable via the packet-switched network, e.g. through a web based application. For this purpose, the correlation unit 402 may comprise one or more network interface cards 601. The correlation unit 402 has a receiving unit 602 and a

transmitter unit 603, for example for receiving correlation related messages, such as for example on-demand correlation requests, and for responding thereto. Furthermore, the receiving unit 602 and a transmitter unit 603 may transmit and receive messages for accessing remote database(s) that contain(s) the signaling session reports and the media stream reports. The receiving unit 602 and transmitter unit 603 may receive and transmit the correlation related messages from its network interface card(s) 601.

The correlation unit 402 may also have a processing unit 604. The processing unit 604 may be used to determine correlation respective signaling sessions to their corresponding media streams. Furthermore, the processing unit 604 may be used to generate services reports.

The service reports may for example comprise information on the respective services and the quality information of the media streams of the respective services. This information may be obtained from the signaling session report and media stream reports belonging to a given service. For correlation, it may be necessary to query one or more databases that comprise the signaling session reports and the media stream reports as provided by one or more signaling plane probes 401 and media plane probes 403. The correlation unit 402 may therefore comprise a database access unit 605 which provides functionality to access these one or more databases. In one exemplary implementation, the database(s) may be located remote from the correlation unit 403.

The correlation unit 402 may also have a storage unit 606, e.g. a RAM, for (temporarily) storing information and data related to the correlation processing. For example, the result set of media stream reports for correlation process, the data of the signaling session reports, and the service report generated for a respective service may be stored in the storage unit 606.

Fig. 5 shows a flow chart of the operation of the correlation unit 402 according to an exemplary embodiment of the invention. Exemplarily, the correlation processing for determining the media streams belonging to one given signaling session report is illustrated in Fig. 5. In case a correlation should be performed for more than one signaling session (report), the procedure of Fig. 5 is repeated for each signaling session report. For example, when using a proactive correlation mechanism, the correlation unit 402 could obtain all signaling session reports that have been generated since the previous correlation process from a signaling session database and process each of the signaling session reports according to Fig. 5. In another embodiment related to an on-demand correlation process, the correlation unit 402 could also receive a correlation request that is indicating one or more search criteria, and may first – not shown in Fig. 5 – query a signaling session database for all signaling session reports that match the search criteria. The signaling session reports matching the search one or more criteria may then be each processed according to Fig. 5 and the service report(s) are returned to the querying party.

The correlation unit 402 obtains 501 a signaling session report to be processed and generates 502 a service report for the service corresponding to the given signaling session. This service report may include information contained in the signaling session report and will be further supplemented with additional information as described below. Furthermore, the correlation unit 402 obtains 503 a correlation key contained in the signaling session report – as noted previously there should be at least two correlation keys contained in the signaling session report. Using the obtained correlation key, the correlation unit 402 queries 504 a media report database to receive all media stream reports that match the correlation key i.e. that contain the same correlation key.

As noted above, it may happen that the query results in the database returning more than one media stream report for the given correlation key, for example due to a hash collision of the correlation key. Therefore, the correlation unit 402 checks 505 whether the returned result set of the database query comprises more than one media stream report. If not (no), the correlation unit 402 can immediately enrich 508 the service report with information on the media stream as contained in the one media stream report returned from the media stream database.

In case the result set returned by the database comprises multiple media stream reports, the correlation unit 402 may try to detect 506 and resolve a hash collision. For example, the hash collision could be resolved by checking further information in the media stream reports of the result set for parameters that identify the media stream (report) that belongs to the correlation key that has been queried. For example, in case the media stream reports contain timestamps of the beginning and/or end of the media stream, same could be compared to respective timestamps contained in the signaling session report, which yield the beginning and/or end of the signaling session. The media stream report(s) of media stream(s) that start before or end after the signaling session may be obtained from the database due to a hash collision. Accordingly, the correlation unit 402 may select 507 only the most relevant media stream report, i.e. the media stream report that is most likely corresponding to the service of the signaling session report from the result set and may add 508 at least a part of its information to the service report. Further, the correlation unit checks 509 whether there is a further correlation key contained in the signaling session report and, if so (yes), repeats steps 503 to 508. These steps are repeated for each further correlation key(s) comprised in the signaling session report.

Once all correlation keys contained in the signaling session report have been processed, the service report for the signaling session is complete, and may be optionally stored 510 by the correlation unit 402 in a service report database.

Depending on the correlation type or application, the procedure of Fig. 5 could be supplemented with additional steps. For example, in case of a proactive/periodic correlation, the correlation unit 402 may check whether there is another signaling session report to be matched to its corresponding media stream reports, and if so the procedure of Fig. 5 would be executed for each signaling session report to be processed. In case of performing a correlation on demand, the correlation unit 402 may likewise check whether there is another signaling session report to be matched to its corresponding media stream reports, and if so the procedure of Fig. 5 would be executed for each signaling session report to be processed. Once all signaling session reports have been processed the correlation unit 402 may generate a correlation response message comprising the service reports generated by the correlation unit 402 in response to the search criteria provided in a correlation request, and returns the correlation response message to the querying entity, e.g. another device/computer.

Signaling Plane Probe

Fig. 9 shows an exemplary implementation of a signaling plane probe 401 and Figs. 7 and 8 show two exemplary flow charts of the operation of same according to different embodiments of the invention. In general, the signaling plane probe 401 observes one or more network links to obtain a copy of all packets passing these network links/interfaces. For this purpose, the signaling plane probe 401 comprises a monitoring unit 901 that receives 701 packets of a service on the packet-switched network 400. For example, this monitoring unit 901 may receive the packets from a one or more specialized packet capture cards 906 (network interface card - NIC) which is capable of receiving and processing packets on the physical interface and to provide them to the application layer, nearly without requiring CPU processing time and operation system functionality. An optional filter in each packet capture card of the signaling plane probe 401 may be used to prevent unwanted packets from being copied from the network interfaces to the processing processes. These filters may be for example implemented as stateless or stateful Internet Protocol (IP) or payload filters. For example, this filter may check 702, whether the currently processed packet is a signaling packet or not. If the packet does not belong to a signaling session, i.e. is not packet of a signaling related protocol, the next packet may be processed.

Each signaling packet received at the signaling plane probe 401 will be processed by a detection unit 902 which analyzes the contents of the packets provided by the monitoring unit 901 so as to identify and process packets as signaling packets. Specifically, the detection unit 902 detects signaling packets and matches these to signaling sessions of the media service, and is capable of identifying one or more media data paths of a media plane of the

service based on signaling information of the service comprised in the received packets. Optionally, the detection unit 902 may be provided in one or more specialized packet capture cards.

For example, the detection unit 902 identifies packets containing signaling information (signaling packets) in a parsing process based on their header information. If parsed successfully as a signaling packet, the signaling packet is forwarded to a state engine. The state engine is designed to identify 703 and form signaling sessions from individual signaling packets.

The state engine (which may be considered an internal database with correlation related information maintained in the storage unit 903) of the detection unit 902 has a processing facility and a state table which holds information about signaling sessions (signaling session records). Based on the protocol-specific identifiers in the signaling packets the detection unit 902 can identify the start and end of a signaling session, respectively media service, or "append" a respective processed signaling packet to an existing signaling session. Signaling packets may also contain information on the start and termination of media streams in the media plane of the media service, which are events that may be tracked by the detection unit 902. Depending on whether the identified signaling session is already registered in the signaling plane probe 401 (step 704), the detection unit 902 registers 705 the new signaling session in the state table, including the generation of a new signaling session record for the signaling session.

The signaling session record may be supplemented with information on the signaling session while same is running to account – for example – for endpoint address/name (e.g. SIP URI); the phone number of the user; the user-agent (endpoint software and version); capabilities/features; supported media types; IP addresses and ports for signaling and media streams (from SDP body); signaling routing information; timestamps: e.g. session start, Post-Dial-Delay (PDD), connect time, session end; information on the session flow (e.g. signaling message(s) with timestamps); session state (e.g. successful, unsuccessful, timed out, reason for failure), and/or events during session (media on hold, codec renegotiation, session forking, authentication), etc.

Furthermore, by analyzing 706 the contents of the signaling packets, e.g. a SDP body of a SIP messages comprised in signaling packets, the detection unit 902 may further learn one or more media data paths of the media plane of the service belonging to the media service that is associated to a given signaling session. The signaling plane probe 401 may further comprise a storage unit 903, e.g. a volatile memory, to store the information describing the state of the respective signaling sessions in the state table. In case the signaling packet

indicates the start of a new media stream (step 706: yes), the detection unit 902 may add 707 information on the media data path (e.g. IP address and port number of the destination, and optionally the source of the media stream) of this media stream to the signaling session record. Likewise, in case the signaling packet indicates the end/termination of a new media stream (step 708: yes), the detection unit 902 may remove the information on the respective media stream from the signaling session record maintained in storage unit 903. In case the signaling packet indicates the termination of the signaling session (step 710: yes), the detection unit 902 may either itself or trigger the processing unit 904 to generate 711 a respective correlation key for each media data path indicated in the signaling session record.

10 The correlation key for a respective media data path of a media stream may be generated for example based on information on the directivity of the media stream (e.g. as given by the tuple of destination IP address and UDP port number) and a timestamp indicating the point in time at which the signaling session was terminated. For example, this timestamp may indicate the point in time of reception of the signaling packet indicating the event of the

15 termination of the signaling session at the signaling plane probe 401. For example, the network interface card(s) 906 of the signaling plane probe could generate a timestamp indicating the time of reception of the respective signaling plane packet and could make same available to the other components of the signaling plane probe 401. Furthermore, the detection unit 902 or the processing unit 904 may further generate 912 a signaling session

20 report based on the information on the signaling session gathered in the signaling session record. Further, the signaling session report also included the correlation keys for the media data paths of the service that exist at the end of the service according to the signaling plane information. The signaling session report may then be stored in a signaling session report database, for example using the database access unit 907 of the signaling plane probe 401.

25 Moreover, the signaling session is deregistered 713 from the state table.

Once a given signaling packet has been processed, it may be discarded. It should be also noted that the order of the checks in steps 706, 708 and 710 and the respective actions to be taken is only exemplary. It is only important that the information on the media streams of the service that exist in the media plane are tracked by the signaling plane probe 401, and that

30 the signaling plane probe 401 is able to detect the termination of the signaling session.

Optionally (not shown in Fig. 7), in another exemplary embodiment, the detection unit 902 of the signaling plane probe 401 periodically (e.g. every 5, 10 or 15 seconds) traverses all entries in the state table in order to find and clean up stale signaling sessions. For this purpose, the signaling plane probe 401 may add a timestamp to each signaling session

35 record in the state table that is updated each time a new signaling packet/message of the respective signaling session is received. Based on this timestamp allows to identify stale

signaling session by simply checking based on the signaling sessions' timestamps whether the last signaling packet/message for a given signaling session has been received longer than a threshold time period ago. For example, stale sessions might be those which did not receive new signaling packets/messages for 86400 seconds or signaling sessions may be defined stale and are terminated after a certain time. This clean-up removes unnecessary information from the state.

Fig. 8 shows another flow chart of the operation of a signaling plane probe 401 according to another embodiment of the invention. In the example discussed with respect to Fig. 7, the signaling plane probe 401 keeps track of the media streams currently existing in the media plane of the service, eliminating those media streams from the signaling session record that are terminated during the ongoing signaling session/service. In the flow chart of Fig. 8, which is corresponding to Fig. 7 to a large extent, the termination of one or more media streams in the media plane, as detected in step 708, causes the detection unit 902 or the processing unit 904 (triggered by the detection unit 902) to generate 801 a correlation key for each terminated media stream. Similar to the other correlation keys, the correlation key for a terminated media stream may be generated based on its media data path information (e.g. destination IP address and UDP port number) and a timestamp that indicates the point in time the media stream was terminated. This timestamp may indicate the point in time of reception of the signaling packet indicating the event of the termination of the media stream at the signaling plane probe 401. The correlation key generated is then added 802 to the signaling session record of the signaling session by the detection unit 902 or the processing unit 904 and is later then also stored in the signaling session report in step 712 together with the other correlation keys either generated in step 801 or step 711.

Media Plane Probe

Next, the media plane probe 403 according to an exemplary embodiment of the invention is described with respect to Fig. 10 and Fig. 11. Fig 10 shows an exemplary embodiment of a media plane probe 403, while Fig. 11 shows a flow chart of its exemplary operation according to an embodiment of the invention.

Similar to the signaling plane probe 401, the media plane probe 403 observes one or more network links to obtain a copy of all packets passing these interfaces. The media plane probe 403 may observe one or more network links to obtain a copy of all packets passing these interfaces. The media plane probe 403 comprises a monitoring unit 1001 that receives 1101 packets of a service on the packet-switched network 401. For example, this monitoring unit 1001 received the packets from one or more specialized packet capture cards 1006 (network interface card - NIC) which are capable of receiving and processing packets on the physical

interface and to provide them to the application layer, nearly without requiring CPU processing time and operation system functionality. An optional filter in the packet capture card of the media plane probe 403 may be used to prevent unwanted packets from being copied from the network interfaces to the processing processes. This filter may be for example implemented as stateless or stateful Internet Protocol (IP) or payload filter. For example, this filter may check 1102, whether the currently processed packet is a media plane packet (media packet) or is some other packet. If the packet does not belong to the media plane, i.e. is not packet of a media related protocol, such as for example RTP, the next packet may be processed.

Each media packet received will be processed by the media plane probe 403. The detection unit 1002 of the media plane probe 403 is designed to detect and form media streams from single media packets. The media plane probe 403 may further comprise a processing unit 1002 that may hold a state table which stores information about the media streams monitored by the media plane probe 403 in a local storage unit 1003, e.g. a volatile memory. Based on the protocol-specific identifiers in the media streams the media plane probe 403 (e.g. based on header information in the RTP packets and/or lower layer headers) the detection unit 1002 identifies 1103 the media data path, i.e. the media stream, the media packet belongs to. If the media data path has not been yet registered (step 1104: no), i.e. the processed media packet starts a new media stream (with the given media data path) the detection unit 1002 generates a new media data record in the state table to account for the new media data path observed by the media plane probe 403. This media stream record may contain information on the media stream such as for example the destination (and optionally the source) IP address and port combination of the media stream, the SSRC in the packet header (in case RTP is used for media transport), etc. If the media data path is already known (step 1104: yes), the media stream record already exists. In any case, the detection unit 1002 updates a timestamp 1106 in the media stream record to indicate the point in time at which the last media packet of the media stream has been received. This timestamp may be used for calculating the correlation key for the media stream. Furthermore, the detection unit 1002 may (re)start a timer 1108 for the media stream record, which when timing out is indicating the termination of the media stream as will be outlined below. Alternatively, no timer is used, but the media plane probe 403 periodically checks for media streams the last media packet of which was received more than a threshold time period ago, in order to determine terminated media streams.

The media stream records of the media streams may for example indicate the respective media data paths thereof by an IP address and port tuple of the destination device, i.e. the receiving device of the media stream. The media plane probe's detection unit 1002 stores

information describing the state of the detected media streams in the media stream records maintained in the state table. Optionally, the individual media streams (i.e. their media data paths) may be identified by a media path identifier (e.g. stream hash key) as unique key. Further optionally, upon detecting a new media stream, respectively media path (step 1104:
5 no), the processing unit 1004 could generate 1105 a media path identifier (e.g. a media path identifier (e.g. stream hash key)) for the media stream, respectively media path and stores 1106 the media stream record in the state table as outlined above.

Optionally, the processing unit 1004 could analyze every received packet for quality information 1107 and – further optionally – store them in the storage unit 1003 for a given
10 time period. Periodically the processing unit 1004 may aggregate or compress the quality information to form fine-grained media stream quality reports with media stream related information. Fine-grained media stream quality reports could be for example written in predetermined time intervals, e.g. every five seconds, and could provide a summary of the media stream in the respective interval and the aggregated quality information of the media
15 stream in the respective interval. Upon generation of a respective fine-grained media stream quality report, the packets used for the generation of this fine-grained media stream quality report can be removed from the storage unit 1003.

The processing unit 1004 of the media plane probe may for example periodically check 1109 (e.g. every 5, 10 or 15 seconds); whether a timer set of a media stream (record) in step 1108
20 has expired. Alternatively, as noted above, the media plane probe 403 may also check timestamps indicating the last update (i.e. last reception of a media packet) of the respective media stream reports in the state table. A media stream may be for example considered terminated, if no new media packets have been received by the monitoring unit 1001 for the existing media stream for a given amount of time (e.g. 60 seconds). Accordingly, if using a
25 timer to monitor the termination, the timer value could be set to this given amount of time. Further alternatively, instead of periodically checking the media stream records or expiry of their timers, the expiry of a timer for a media stream (record) could cause an interrupt invoking the processing unit 1004 to perform steps 1110 to 1112 as outlined below.

Upon detecting the termination of a media stream, the processing unit 1004 processes the
30 respective information in the media stream record of this media stream the state table. The processing unit 1004 generates 1110 a correlation key for the media stream based on its media data path information (e.g. the destination IP address and UDP port number) and the timestamp in the media stream record that is indicating the point in time of receiving the last media packet of the media stream (see step 1106). Furthermore, the processing unit 1004
35 generates 1111 a media stream report for the media data path, respectively media stream, which includes the generated correlation key and further information on the media stream.

This further information may be for example transport media information and optionally stream quality information. For example, in case RTP is used for media transport, the transport media information could include the source and destination IP address of the RTP stream, the source and destination UDP port of the RTP stream, the RTP header Synchronization source (SSRC) value and the RTP payload type. Furthermore, the processing unit 1004 may also generate information on the media stream quality that is included to the media stream report as stream quality information. This stream quality information could for example include timestamp information (e.g. start and/or end time of the media stream), an indication of RTP format violations (e.g. the number thereof), information on the protocol conformance (e.g. sequence errors and/or DSCP errors), and transport information (e.g. packet interval and jitter statistics, information on packet loss, information on duplicate packets, etc.), MOS values (e.g. per time-slice (e.g. every 5 seconds), minimum/average/maximum value for entire stream duration, etc.), and/or events (e.g. comfort noise, silence suppression, DTMF tones, codec changes, marker bits). The contents of the media stream reports may be configurable by an operator and may be adapted to customers' needs.

The media stream record may be further stored in a media stream report database using the database access unit 1007 of the media plane probe 403. Finally, the processing unit 1004 of the media plane probe 403 deregisters 1112 the media data path from the state table, e.g. by deleting the corresponding media stream record.

Steps 1110 to 1112 are performed for every media stream that is considered terminated by the media plane probe.

Optimizations may be used to speed up processing of entries in the state tables inside the signaling plane probe 401, the correlation unit 402 and the media plane probe 403. A simple list would be sufficient to implement the features as outlined above. Add, modify, delete and search operations on simple lists may however be too slow. These operations may be implemented utilizing more complex but faster data structures, e.g. hash tables or data structures heavily dependent on CPU and Operating System specific behavior and functionality. The Intel Thread Building Block (TBB) library for instance can be used to implement in-memory databases used for the state tables in the signaling plane probe 401, the correlation unit 402 and the media plane probe 403. The Intel Thread Building Block (TBB) library may also be used to aid implementation of the processing queues inside and in between the components outlined above.

Please note that in case the signaling plane probe 401 and the media plane probe 403 should perform passive non-intrusive monitoring of the signaling plane traffic, respectively

media plane traffic, a separate network interface card is employed for the transmission of correlation related messages (i.e. the network interface card(s) for monitoring the traffic are only employed for packet reception but not for packet transmission).

5 Please also note that in the procedures described above the detection unit 902, 1002 and the processing unit 904, 1004 may use the database access unit 907, 1007 to access the database/state table in the storage unit 903, 1003.

Application

10 The correlation mechanism for signaling plane and media plane of a service may be used in manifold applications. For example, the correlation mechanism is useful in small networks but is more important in larger scale networks where the various signaling and media components are decomposed and present on different, sometimes geographically split, network entities. This is the case in IP Multimedia Subsystem (IMS) or Next Generation Network (NGN) setups.

15 Fig. 12 shows a simplified network in which locations are indicated where correlation may be desired and can be implemented using the correlation mechanism according to one of the various embodiments described herein. Possible monitoring locations are for example (and not limited thereto):

- The Access Session Border Controller (SBC) network interfaces with end customers (shown on the left side of Fig. 12),
- 20 - a decomposed gateways system, e.g. to the PSTN, is shown as Media Gateway (MGW) and a Media Gateway Controller (MGC) – see in the top right section of Fig. 12, and
- an Interconnection SBC interfaces to one or more interconnection or transit partners (see the lower right side of Fig. 12).

25 It should be noted that the location of the monitoring probes for signaling plane and media plane may also be subject to requirements on the monitoring solution.

In some cases, signaling plane and media plane may originate and/or terminate at the same endpoint (device). In case of an end customer this endpoint most likely uses the same IP address on the signaling and media plane. These streams may be for example observed, if the signaling plane probes 401 and the media plane probes 403 are placed close to the end customer devices. The Access Session Border Controller may be considered to be close to
30 the carrier internal network's "border" towards end customer devices and to therefore receive

- traffic from the end customers on one side and transmits this traffic on the other side towards the carrier internal network. Most commonly these Session Border Controllers (SBC) are used to separate the devices outside the own trust domain from the inside network. Although the traffic is in one geographic location the SBC may use different IP addresses to load
- 5 balance traffic. In this case a simple IP address based correlation mechanism not using the principles and ideas outlined herein would already fail whereas the correlation mechanism described herein would succeed. The signaling plane probes 401 and media plane probes 403 may be for example placed in a way that they observe the signaling streams and media streams respectively.
- 10 In the upper right corner of the diagram shown in Fig. 12, a decomposed gateway is shown which uses different physical entities as media and signaling components (MGC and MGW, respectively). In between those devices a third party protocol may be used to synchronize the devices. For the sake of a minimum size network diagram only one of each component
- 15 types is shown but in real life scenarios multiple media components are controlled by one signaling component (as it is the case in IMS, NGN). The media components may be co-located in the same location with the signaling component or spread across the internal network of the carrier. In either case the correlation task is rather complex as it involves multiple IP addresses on different physical network links and thus may require multiple
- 20 signaling plane probes 401 as well as media plane probes 403.
- 20 In a third exemplary scenario the correlation has to correlate signaling and media streams from different locations. A media plane probe 403 component on the internal side of the interconnection point (lower right corner of the diagram – Interconnection SBC) observes a media stream from the Access SBC (middle left). The corresponding signaling session is observed next to the aforementioned Access SBC by a separate signaling plane probe 401.
- 25 The location of the correlation unit 402 does not impact what signaling sessions or media streams can be correlated. The correlation unit 402 should have access to the signaling session reports and media stream reports of the signaling plane probes 401 and media plane probes 403 and can be implemented as distinct physical machine or co-located on an existing signaling plane probe 401 or media plane probe 403.
- 30 The terms probe and correlation unit used herein can be considered to refer to dedicated hardware, a software program or a combination thereof to implement the desired functionality of the respective probe, respectively, correlation unit. Accordingly, another embodiment of the invention relates to the implementation of the above described various embodiments using hardware and software. It is recognized that probes and correlation unit of the
- 35 invention may be implemented by having computing devices (processors) execute a software

program that causes the desired functionality being executed. A computing device or processor may for example be general purpose processors, digital signal processors (DSP), application specific integrated circuits (ASIC), field programmable gate arrays (FPGA) or other programmable logic devices, etc. Some of the computing devices may store the software program internally.

The various embodiments of the invention may also be performed or embodied by a combination of computing devices and software programs providing the desired functionality stored on any kind of computer readable storage media, for example RAM, EPROM, EEPROM, flash memory, registers, hard disks, CD-ROM, DVD, etc.

It should be further noted that the individual features of the different embodiments of the invention may individually or in arbitrary combination be subject matter to another invention.

It would be appreciated by a person skilled in the art that numerous variations and/or modifications may be made to the present invention as shown in the specific embodiments without departing from the spirit or scope of the invention as broadly described. The present embodiments are, therefore, to be considered in all respects to be illustrative and not restrictive.

CLAIMS

1. A method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, the method comprising:
 - a) generating at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe, wherein the correlation keys for a respective service are generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service,
 - b) storing, by the signaling plane probe, a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service,
 - c) generating a correlation key for each media stream monitored by the media plane probe, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,
 - d) storing, by the media plane probe, a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media, and
 - e) correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.
2. The method according to claim 1, wherein the generation of the respective correlation keys for the respective media streams by the media plane probe is based on the media data path information of the respective media stream and the timestamp indicative of the point in time at which the respective media stream ended.
3. The method according to claim 1 or 2, wherein information on the direction of the respective service, based on which the correlation keys are generated by

signaling plane probe for a respective service, is information on one or more media data paths of media streams of the service learned from messages of the signaling session of the service.

4. The method according to one of claims 1 to 3, further comprising the steps of:
 - f) monitoring, by the signaling plane probe, signaling messages of the respective signaling sessions of the services, and
 - g) detecting, by the signaling plane probe, media streams belonging to each service associated to a respective monitored signaling session based on the signaling messages of the respective signaling sessions,wherein the correlation keys for a respective one of the signaling sessions are generated based on information on the media data path of the media streams belonging to the service associated to the respective one of the signaling sessions and the timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe.
5. The method according to one of claims 1 to 4, wherein information on a media data path of a media stream is a destination IP address and port number of the media stream.
6. The method according to one of claims 1 to 5, wherein the timestamps used for the generation of the correlation keys by the signaling plane probe and the media plane probe are rounded before the generation of the correlation keys or have a granularity in the range of one to ten seconds.
7. The method according to one of claims 1 to 6, wherein the signaling plane probe generates for each direction of the service at least one correlation key using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and being rounded to the next lower timestamp value according to a rounding interval, and at least another correlation key using the timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and being rounded to the next higher timestamp value according to the rounding interval.

8. The method according to one of claims 1 to 7, wherein the media plane probe generates for each media stream at least one correlation key using a timestamp indicative of the point in time at which the respective media ended and being rounded to the next lower timestamp value according to a rounding interval, and at least another correlation key using the timestamp indicative of the point in time at which the respective media stream ended and being rounded to the next higher timestamp value according to the rounding interval.
9. The method according to one of claims 1 to 8, wherein the correlation of the media streams and the signaling session of the respective services is performed by a correlation unit periodically or in response to a query.
10. A method for correlating media streams and signaling sessions of services, the method comprising the steps of:
 - a) correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services, wherein correlating the media streams and the signaling sessions comprises identifying one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys contained in the media stream reports and signaling session reports, and
 - b) generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.
11. The method according to claim 10, wherein correlation keys provided in a respective signaling session report have been generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the corresponding service has been detected and using information on the media streams of the service.
12. The method according to claim 10 or 11, wherein correlation key provided in a respective media stream report has been generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the media stream.

13. The method according to one of claims 10 to 12, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:
- matching the respective correlation keys of the media stream reports of media streams are to the respective correlation keys of the signaling session reports, or vice versa,
- wherein a match of the correlation keys of a one of the media stream reports and one of the signaling session reports yields the media stream of said one media stream report belonging to the same service to which the signaling session of said one signaling session report belongs.
14. The method according to one of claims 10 to 13, wherein the service report comprises information on the service and the quality information of the media streams of the service.
15. The method according to one of claims 10 to 13, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:
- obtaining from a signaling session report correlation keys identifying the media streams belonging to the same service as the signaling session, and
- identifying those media stream reports which include one of the correlation keys obtained from the signaling session report.
16. The method according to one of claims 10 to 13, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:
- obtaining from a media stream report a correlation key identifying the media stream, and
- identifying a signaling session report among the signaling session reports which includes the same correlation key obtained from the signaling session report.

17. The method according to one of claims 10 to 16, wherein steps a) and b) are performed by a correlation unit periodically for all signaling session reports and media stream reports generated within a configurable time interval, or in response to a request.
18. The method according to one of claims 10 to 17, further comprising steps a) to d) of claim 1 or the steps of the method according to one of claims 2 to 9.
19. A passive monitoring system for correlating media streams and signaling sessions of services in a packet-switched network, the passive monitoring system comprising:
 - a signaling plane probe for generating at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe, wherein the correlation keys for a respective service are generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service,
 - wherein the signaling plane probe is further adapted to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service,
 - a media plane probe for generating a correlation key for each media stream monitored by the media plane probe, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,
 - wherein the media plane probe is adapted to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media, and
 - a correlation unit for correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.
20. A signaling plane probe comprising:

a network interface card for receiving packets of signaling sessions of multiple media services,

monitoring unit for monitoring the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions, wherein the monitoring unit is further adapted to detect the point in time at which a respective signaling session of the respective service terminates, and

a processing unit for generating at least one correlation key for each direction of a service based on a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the monitoring unit and based on information on the direction of the respective service,

wherein the monitoring unit is adapted to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises at least two correlation keys generated for the respective directions of the service.

21. A media plane probe comprising:

a network interface card for receiving packets of media streams of multiple media services,

monitoring unit for monitoring the media streams based on the packets of the media streams received by the media plane probe to thereby obtain information on a media data path a respective one of the media streams, wherein the monitoring unit is further adapted to detect the point in time at which a respective media stream of the respective service terminates, and

a processing unit for generating a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

wherein the monitoring unit is adapted to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

22. A correlation unit for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, the correlation unit comprising:

a database access unit for accessing a database containing signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services, wherein the signaling session reports each contain at least one correlation key for each direction of a respective service, and the media stream reports each contain one correlation key indication the direction of the media stream,

a processing unit for correlating the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services, wherein the processing unit is adapted to identify one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys contained in the media stream reports and signaling session reports, and

a service record generation unit for generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

23. The correlation unit according to claim 22, further comprising means to perform the steps of the method according to one of claims 10 to 17.

24. A computer-readable medium storing instructions that, when executed by a processor of a signaling plane probe, cause the signaling plane probe:

to receive packets of signaling sessions of multiple media services,

to monitor the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions,

to detect the point in time at which a respective signaling session of the respective service terminates,

to generate at least one correlation key for each direction of a service based a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected and based on information on the direction of the respective service, and

to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

25. A computer-readable medium storing instructions that, when executed by a processor of a media plane probe, cause the media plane probe:

to receive packets of media streams of multiple media services,

to monitor the media streams based on the packets of the media streams received by the signaling plane probe to thereby obtain information on a media data path a respective one of the media streams,

to detect the point in time at which a respective media stream of the respective service terminates, and

to generate a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

to store a media stream report for each media stream in a database, wherein each media stream report contains at least the correlation key generated for the respective media.

26. A computer-readable medium storing instructions that, when executed by a correlation unit, cause the correlation unit to correlate media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, by:

accessing a database containing signaling session reports for signaling sessions of services and a database containing media stream reports of media streams of the services, wherein the signaling session reports each contain at least one correlation key for each direction of a respective service, and the media stream reports each contain one correlation key indication the direction of the media stream,

correlating the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services,

identifying one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports, and

generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

27. The computer-readable medium according to claim 26, further storing instruction that, when executed by the processor of the correlation unit, cause the correlation unit to perform the steps of the method according to one of claims 10 to 17.




Fig. 1




Fig. 2




Fig. 3




Fig. 4




Fig. 5


**Fig. 6**

Fig. 7







Fig. 8



signaling plane probe 401**Fig. 9****media plane probe 403****Fig. 10**




Fig. 11






Fig. 12

**Fig. 13****Fig. 14**

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2011/002145

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/06 H04L12/24
 ADD. H04L12/26

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 H04L

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

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

EPO-Internal, COMPENDEX, INSPEC, IBM-TDB, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ATSUSHI KOBAYASHI ET AL: "VoIP Measurement Architecture Using Data Mediation", 29 October 2009 (2009-10-29), IP OPERATIONS AND MANAGEMENT, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 15 - 26, XP019133039, ISBN: 978-3-642-04967-5 page 18, line 5 - page 20, line 35 -----	1-27
A	US 2008/151918 A1 (FOTI GEORGE [CA]) 26 June 2008 (2008-06-26) the whole document -----	1-27
A	WO 2010/106075 A1 (KONINKL KPN NV [NL]; TNO [NL]; STOKKING HANS MAARTEN [NL]; WALRAVEN FA) 23 September 2010 (2010-09-23) page 4, line 16 - page 8, line 26 -----	1-27



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

10 January 2012

Date of mailing of the international search report

17/01/2012

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer

Aura Marcos, F

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2011/002145

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008151918 A1	26-06-2008	US 2008151918 A1	26-06-2008
		WO 2008078209 A2	03-07-2008
-----		-----	
WO 2010106075 A1	23-09-2010	EP 2409432 A1	25-01-2012
		KR 20110125674 A	21-11-2011
		WO 2010106075 A1	23-09-2010
-----		-----	

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference 007742.00043	FOR FURTHER ACTION see Form PCT/ISA/220 as well as, where applicable, item 5 below.	
International application No. PCT/US2015/024691	International filing date (<i>day/month/year</i>) 7 April 2015 (07-04-2015)	(Earliest) Priority Date (<i>day/month/year</i>) 16 April 2014 (16-04-2014)
Applicant CENTRIPETAL NETWORKS, INC.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 6 sheets.

☐

It is also accompanied by a copy of each prior art document cited in this report.

1. Basis of the report

a. With regard to the **language**, the international search was carried out on the basis of:

☒

the international application in the language in which it was filed

☐

a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))

b. ☐

This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6**bis**(a)).

c. ☐

With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. ☐

Certain claims were found unsearchable (See Box No. II)

3. ☒

Unity of invention is lacking (see Box No III)

4. With regard to the **title**,

☒

the text is approved as submitted by the applicant

☐

the text has been established by this Authority to read as follows:

5. With regard to the **abstract**,

☒

the text is approved as submitted by the applicant

☐

the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings**,

a. the figure of the **drawings** to be published with the abstract is Figure No. 11

☐

as suggested by the applicant

☐

as selected by this Authority, because the applicant failed to suggest a figure

☒

as selected by this Authority, because this figure better characterizes the invention

b. ☐

none of the figures is to be published with the abstract

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/024691

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-9

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2015/024691

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/06

ADD. H04L29/08

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2006/136987 A1 (OKUDA MASATO [JP]) 22 June 2006 (2006-06-22) paragraph [0028] - paragraph [0030]; figures 1,21 paragraphs [0033], [0036], [0040], [0093], [0095], [0135] -----	1-9
X	US 2007/240208 A1 (YU MING-CHE [TW] ET AL) 11 October 2007 (2007-10-11) abstract paragraphs [0021], [0024] - paragraph [0026] paragraph [0033] - paragraph [0036] paragraph [0039] - paragraph [0040] ----- -/--	1-9

☒ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

3 July 2015

Date of mailing of the international search report

16/09/2015

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Barla Harter, I

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2015/024691

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2006/146879 A1 (ANTHIAS TEFCROS [US] ET AL) 6 July 2006 (2006-07-06) paragraph [0074] - paragraph [0076] paragraph [0112] - paragraph [0119] paragraphs [0207], [0208], [0240] -----	1-9
A	US 8 306 994 B2 (KENWORTHY STACY [US]) 6 November 2012 (2012-11-06) abstract; figures 4,6 column 2 - column 4 column 6 - column 9 -----	1-9

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2015/024691

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006136987 A1	22-06-2006	JP 2006174350 A US 2006136987 A1	29-06-2006 22-06-2006
US 2007240208 A1	11-10-2007	NONE	
US 2006146879 A1	06-07-2006	CN 101099345 A EP 1834453 A2 US 2006146879 A1 WO 2006073740 A2	02-01-2008 19-09-2007 06-07-2006 13-07-2006
US 8306994 B2	06-11-2012	US 6317837 B1 US 2002049899 A1 US 2010242098 A1 US 2013061294 A1	13-11-2001 25-04-2002 23-09-2010 07-03-2013

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-9

A method for packet filtering based on application-layer
packet header information

2. claims: 10-18

A method for selective packet digest logging

3. claims: 19, 20

A method for security policy update and management

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY

(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44bis)

Applicant's or agent's file reference 007742.00023	FOR FURTHER ACTION	See item 4 below
International application No. PCT/US2013/057502	International filing date (<i>day/month/year</i>) 30 August 2013 (30.08.2013)	Priority date (<i>day/month/year</i>) 22 October 2012 (22.10.2012)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1.	This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 bis.1(a).																								
2.	This REPORT consists of a total of 5 sheets, including this cover sheet.																								
	In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.																								
3.	This report contains indications relating to the following items: <table style="width: 100%;"> <tr> <td style="width: 5%; text-align: center;">☒</td> <td style="width: 25%;">Box No. I</td> <td style="width: 70%;">Basis of the report</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. II</td> <td>Priority</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. III</td> <td>Non-establishment of opinion with regard to novelty, inventive step and industrial applicability</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. IV</td> <td>Lack of unity of invention</td> </tr> <tr> <td style="text-align: center;">☒</td> <td>Box No. V</td> <td>Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. VI</td> <td>Certain documents cited</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. VII</td> <td>Certain defects in the international application</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. VIII</td> <td>Certain observations on the international application</td> </tr> </table>	☒	Box No. I	Basis of the report	☐	Box No. II	Priority	☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability	☐	Box No. IV	Lack of unity of invention	☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement	☐	Box No. VI	Certain documents cited	☐	Box No. VII	Certain defects in the international application	☐	Box No. VIII	Certain observations on the international application
☒	Box No. I	Basis of the report																							
☐	Box No. II	Priority																							
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability																							
☐	Box No. IV	Lack of unity of invention																							
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement																							
☐	Box No. VI	Certain documents cited																							
☐	Box No. VII	Certain defects in the international application																							
☐	Box No. VIII	Certain observations on the international application																							
4.	The International Bureau will communicate this report to designated Offices in accordance with Rules 44bis.3(c) and 93bis.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44bis .2).																								

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Date of issuance of this report 28 April 2015 (28.04.2015) Authorized officer Kihwan Moon e-mail: pt01.pct@wipo.int
--	---

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY (PCT Rule 43*bis*.1)

To:

see form PCT/ISA/220

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.
PCT/US2013/057502

International filing date (day/month/year)
30.08.2013

Priority date (day/month/year)
22.10.2012

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. FURTHER ACTION

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:



European Patent Office
P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040
Fax: +31 70 340 - 3016

Date of completion of
this opinion

see form
PCT/ISA/210

Authorized Officer

Veen, Gerardus

Telephone No. +31 70 340-3811



Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:
 - ☒ the international application in the language in which it was filed
 - ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).
2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))
3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:
 - a. (means)
 - ☐ on paper
 - ☐ in electronic form
 - b. (time)
 - ☐ in the international application as filed
 - ☐ together with the international application in electronic form
 - ☐ subsequently to this Authority for the purposes of search
4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.
5. Additional comments:

Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes: Claims	
	No: Claims	<u>1-20</u>
Inventive step (IS)	Yes: Claims	
	No: Claims	<u>1-20</u>
Industrial applicability (IA)	Yes: Claims	<u>1-20</u>
	No: Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 The following documents are referred to:

D1 EP 1 677 484 A2 (MICROSOFT CORP [US]) 5 July 2006 (2006-07-05)

D2 US 6 226 372 B1 (BEEBE TODD [US] ET AL) 1 May 2001 (2001-05-01)

2 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims 1-20 is not new.

2.1 In the words of independent claim 1, D1 discloses:

'A method (column 4, line 13), comprising:

at each of one or more packet security gateways (c.4, l.16, 'enforcement engine') associated with a security policy management server (c.4, l.15, 'policy server'):

receiving a dynamic security policy from the security policy management server (c.4, l.16-19);

receiving packets associated with a network protected by each respective packet security gateway; and

performing, on a packet by packet basis, at least one of multiple packet transformation functions specified by the dynamic security policy on the packets (c.4, l.19-20, 'enforcing the security policies'), wherein performing the at least one of multiple packet transformation functions specified by the dynamic security policy on the packets comprises performing at least one packet transformation function other than forwarding or dropping the packets (c.5, l.30-36)'.

Claim 1 is thus not new (Art. 33(2) PCT).

2.2 The same reasoning applies, *mutatis mutandis*, to independent claims 19 and 20 which are, therefore, also not new.

2.3 Dependent claims 2-18 do not contain any features which, in combination with the features of the claims to which they refer, meet the requirements of the PCT in respect of novelty and inventive step, see D1, D2 and the corresponding citations in the search report.

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY

(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44bis)

Applicant's or agent's file reference 007742.00030	FOR FURTHER ACTION	See item 4 below
International application No. PCT/US2014/023286	International filing date (<i>day/month/year</i>) 11 March 2014 (11.03.2014)	Priority date (<i>day/month/year</i>) 12 March 2013 (12.03.2013)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1. This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 bis.1(a).
2. This REPORT consists of a total of 6 sheets, including this cover sheet.

In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.

3. This report contains indications relating to the following items:

☒	Box No. I	Basis of the report
☐	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☐	Box No. VII	Certain defects in the international application
☒	Box No. VIII	Certain observations on the international application

4. The International Bureau will communicate this report to designated Offices in accordance with Rules 44bis.3(c) and 93bis.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44bis .2).

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Date of issuance of this report 15 September 2015 (15.09.2015) Authorized officer Philippe Bécamel e-mail: pt03.pct@wipo.int
--	--

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

PCT

To:

see form PCT/ISA/220

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY (PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.
PCT/US2014/023286

International filing date (day/month/year)
11.03.2014

Priority date (day/month/year)
12.03.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☒ Box No. VIII Certain observations on the international application

2. FURTHER ACTION

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:



European Patent Office
P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk - Pays Bas
Tel. +31 70 340 - 2040
Fax: +31 70 340 - 3016

Date of completion of
this opinion

see form
PCT/ISA/210

Authorized Officer

Veen, Gerardus

Telephone No. +31 70 340-3811



Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:
- ☒ the international application in the language in which it was filed
 - ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).
2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))
3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:
- a. (means)
- ☐ on paper
 - ☐ in electronic form
- b. (time)
- ☐ in the international application as filed
 - ☐ together with the international application in electronic form
 - ☐ subsequently to this Authority for the purposes of search
4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.
5. Additional comments:

Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes: Claims	
	No: Claims	<u>1-20</u>
Inventive step (IS)	Yes: Claims	
	No: Claims	<u>1-20</u>
Industrial applicability (IA)	Yes: Claims	<u>1-20</u>
	No: Claims	

2. Citations and explanations

see separate sheet

Box No. VIII Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

The following documents are referred to:

- D1 US 2003/123456 A1 (DENZ PETER R [US] ET AL) 3 July 2003
 (2003-07-03)
- D2 US 2007/083924 A1 (LU HONGQIAN K [US]) 12 April 2007 (2007-04-12)

1 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims 1-20 is not new.

1.1 Independent claim 1 (as far as it is clear, see **Item VIII**) defines a method for filtering a packet, the method having two stages:

1) in the first stage, a first operation ('operator') is performed upon the packet, based on a first header field value;

2) in the second stage, a second operation ('transformation') is performed upon the packet, based on a second header field value.

Moreover, these two steps are interrelated in the sense that the criteria, with which the second header field value is compared during the second stage, are determined during the first stage.

Such a method is also disclosed in D1 (see the passages cited in the search report).

Present claim 1 is thus not new according to Art. 33(2) PCT.

1.2 The same reasoning applies, mutatis mutandis, to independent claims 11 and 20 which are, therefore, also not new.

1.3 Dependent claims 2-10 and 12-19 do not appear to contain any features which, in combination with the features of the claims to which they refer, meet the requirements of the PCT in respect of novelty and inventive step, see D1, D2 and the corresponding references cited in the search report.

1.4 It is further noted that a reasoning based on D2 would lead to very similar conclusions.

Re Item VIII

Certain observations on the international application

2 Claims 1-20 do not meet the requirements of Article 6 PCT for, at least, the following reasons:

- in claim 1, the expressions 'header field' and 'application header field', and especially how they differ from each other, are not clear;
- a similar remark applies to the expressions 'operator' and 'packet transformation function', respectively.

Similar observations apply to independent claims 11 and 20.

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY

(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44bis)

Applicant's or agent's file reference 007742.00031	FOR FURTHER ACTION	See item 4 below
International application No. PCT/US2014/027723	International filing date (<i>day/month/year</i>) 14 March 2014 (14.03.2014)	Priority date (<i>day/month/year</i>) 15 March 2013 (15.03.2013)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1.	This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 bis.1(a).																								
2.	This REPORT consists of a total of 7 sheets, including this cover sheet.																								
In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.																									
3.	This report contains indications relating to the following items: <table style="width: 100%;"> <tr> <td style="width: 10%; text-align: center;">☒</td> <td style="width: 30%;">Box No. I</td> <td style="width: 60%;">Basis of the report</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. II</td> <td>Priority</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. III</td> <td>Non-establishment of opinion with regard to novelty, inventive step and industrial applicability</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. IV</td> <td>Lack of unity of invention</td> </tr> <tr> <td style="text-align: center;">☒</td> <td>Box No. V</td> <td>Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. VI</td> <td>Certain documents cited</td> </tr> <tr> <td style="text-align: center;">☒</td> <td>Box No. VII</td> <td>Certain defects in the international application</td> </tr> <tr> <td style="text-align: center;">☐</td> <td>Box No. VIII</td> <td>Certain observations on the international application</td> </tr> </table>	☒	Box No. I	Basis of the report	☐	Box No. II	Priority	☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability	☐	Box No. IV	Lack of unity of invention	☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement	☐	Box No. VI	Certain documents cited	☒	Box No. VII	Certain defects in the international application	☐	Box No. VIII	Certain observations on the international application
☒	Box No. I	Basis of the report																							
☐	Box No. II	Priority																							
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability																							
☐	Box No. IV	Lack of unity of invention																							
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement																							
☐	Box No. VI	Certain documents cited																							
☒	Box No. VII	Certain defects in the international application																							
☐	Box No. VIII	Certain observations on the international application																							
4.	The International Bureau will communicate this report to designated Offices in accordance with Rules 44bis.3(c) and 93bis.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44bis .2).																								

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Date of issuance of this report 15 September 2015 (15.09.2015) Authorized officer Nora Lindner e-mail: pt05.pct@wipo.int
--	--

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

PCT

To:

see form PCT/ISA/220

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY (PCT Rule 43bis.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.
PCT/US2014/027723

International filing date (day/month/year)
14.03.2014

Priority date (day/month/year)
15.03.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☒ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. FURTHER ACTION

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1bis(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:



European Patent Office
D-80298 Munich
Tel. +49 89 2399 - 0
Fax: +49 89 2399 - 4465

Date of completion of
this opinion

see form
PCT/ISA/210

Authorized Officer

Olaechea, Javier

Telephone No. +49 89 2399-8233



Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:
 - ☒ the international application in the language in which it was filed
 - ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).
2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))
3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:
 - a. (means)
 - ☐ on paper
 - ☐ in electronic form
 - b. (time)
 - ☐ in the international application as filed
 - ☐ together with the international application in electronic form
 - ☐ subsequently to this Authority for the purposes of search
4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.
5. Additional comments:

Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes: Claims	
	No: Claims	<u>1-20</u>
Inventive step (IS)	Yes: Claims	
	No: Claims	<u>1-20</u>
Industrial applicability (IA)	Yes: Claims	<u>1-20</u>
	No: Claims	

2. Citations and explanations

see separate sheet

Box No. VII Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 Reference is made to the following document:

D1 US 2006/048142 A1 (ROESE JOHN J [US] ET AL) 2 March 2006 (2006-03-02)

D2 Sean Moore: "SBIR Case Study: Centripetal Networks Subtitle: How CNI Leveraged DHS S&T SBIR Funding to Launch a Successful Cyber Security Company Cyber Security Division 2012 Principal Investigators' Meeting 10-October-2012", , 10 October 2012 (2012-10-10), XP055122227, Retrieved from the Internet:
URL:<http://www.dhs.gov/sites/default/files/publications/csd-day-2-04-sbir-centripetal-price.pdf>
[retrieved on 2014-06-06]

D3 Cisco: "Control Plane Policing Implementation Best Practices", , 13 March 2013 (2013-03-13), XP002725728, Retrieved from the Internet:
URL:https://web.archive.org/web/20130313135143/http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html
[retrieved on 2014-06-12]

2 The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of claims 1, 11 and 16 is not new in the sense of Article 33(2) PCT.

2.1 The document D1 discloses (the references in parentheses applying to this document) in terms of claim 1:

A method, comprising:

at a computing platform comprising at least one processor, memory, and communication interface (*network device in Figures 1 and 2*):

receiving, via the communication interface, a plurality of packets (paragraph 55; *traffic is received by the network device*);

responsive to a determination that an overload condition has occurred in one or more networks associated with the computing platform, applying, by the processor and to at least some of the plurality of packets, a first group of packet-filtering rules that comprises at least one five-tuple indicating a first set of packets that should be allowed to continue toward their respective destinations, wherein applying the first group of packet-filtering rules includes allowing at least a first portion of the plurality of packets that fall within the first set of packets to continue toward their respective destinations (paragraphs 24, 56 and table in pages 9-10; *in response to an attack trigger, a rule set is activated by the network device, said rule set including disabling all traffic excluding the management traffic, RR1*); and

responsive to a determination that the overload condition has been mitigated to a first degree, applying, to at least some of the plurality of packets, a second group of packet-filtering rules that comprises at least one five-tuple indicating a second set of packets that should be allowed to continue toward their respective destinations, wherein applying the second group of packet-filtering rules includes allowing at least a second portion of the plurality of packets that fall within the second set of packets to continue toward their respective destinations, and wherein the second group of packet-filtering rules is less restrictive than the first group of packet-filtering rules (paragraph 24; *upon detection of a new triggering condition, a less restrictive set of filtering rules is implemented by the networking device*).

Since all the features of claim 1 are known in combination from document D1, the subject-matter of claim 1 is not new (Article 33(2) PCT). Therefore, claim 1 is not allowable (Article 33(1) PCT).

- 2.2 The same reasoning applies, *mutatis mutandis*, to the subject-matter of the corresponding independent claims 11 and 16, which therefore are also considered not new (Article 33(2) PCT).

- 2.3 Furthermore, it is noted that the disclosure of document D2 (see page 4) would also render the subject-matter of claims 1, 11 and 16 not new or at least not inventive.
- 3 None of the dependent claims 2-10, 12-15 and 17-20 contain additional features which lead the subject-matter to be both new and inventive (see documents D1-D3 and the corresponding passages cited in the search report).

Re Item VII

Certain defects in the international application

The independent claims are not in the two-part form in accordance with Rule 6.3(b) PCT, which in the present case would be appropriate, with those features known in combination from the prior art being placed in the preamble (Rule 6.3(b)(i) PCT) and with the remaining features being included in the characterising part (Rule 6.3(b)(ii) PCT).

The features of the claims are not provided with reference signs placed in parentheses (Rule 6.2(b) PCT).

The description is not in conformity with the claims as required by Rule 5.1(a)(iii) PCT.

Contrary to the requirements of Rule 5.1(a)(ii) PCT, the relevant background art disclosed in the documents D1-D3 is not mentioned in the description, nor are these documents identified therein.

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2016/026339

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/06 H04L12/26
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L

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

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

EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/093513 A1 (CANTRELL CRAIG [US] ET AL) 13 May 2004 (2004-05-13)	1,31,32
Y	abstract paragraph [0012] - paragraph [0044]; figure 1 paragraph [0065] - paragraph [0069] -----	2-30, 33-43
X	US 2012/023576 A1 (SORENSEN AMANDA [US] ET AL) 26 January 2012 (2012-01-26)	1,31,32
Y	abstract paragraph [0006] - paragraph [0012]; figure 3 -----	2-30, 33-43
X	US 2008/163333 A1 (KASRALIKAR RAHUL [US]) 3 July 2008 (2008-07-03)	1,31,32
A	abstract paragraph [0009] - paragraph [0010]; figure 2 -----	2-30, 33-43

☐

Further documents are listed in the continuation of Box C.

☒

See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

1 June 2016

Date of mailing of the international search report

09/06/2016

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

San Millán Maeso, J

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/026339

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004093513	A1	13-05-2004	AR 042020 A1 08-06-2005
		AU 2003290674 A1 03-06-2004	
		CN 1720459 A 11-01-2006	
		EP 1558937 A2 03-08-2005	
		JP 2006506853 A 23-02-2006	
		JP 2010268483 A 25-11-2010	
		KR 20050086441 A 30-08-2005	
		KR 20100132079 A 16-12-2010	
		US 2004093513 A1 13-05-2004	
		US 2005028013 A1 03-02-2005	
		US 2005044422 A1 24-02-2005	
		WO 2004045126 A2 27-05-2004	
US 2012023576	A1	26-01-2012	AU 2011279907 A1 28-02-2013
		CA 2805823 A1 26-01-2012	
		SG 187068 A1 28-02-2013	
		US 2012023576 A1 26-01-2012	
		WO 2012012280 A2 26-01-2012	
US 2008163333	A1	03-07-2008	NONE

TITLE: RULE-BASED NETWORK-THREAT DETECTION

APPLICANT: CENTRIPETAL NETWORKS, INC.

IPC CLASSIFICATION: H04L29/06, H04L12/26

EXAMINER: San Millán Maeso, J

CONSULTED DATABASES:

CLASSIFICATION SYMBOLS DEFINING EXTENT OF THE SEARCH:

IPC:

CPC: TXTE, EPODOC, INSPEC, WPI, XPESP, XPI3E, XPIEE, H04L63/0227, H04L43/06, H04L12/2602, H04L43/00, H04L63/0263, H04L63/1416, H04L63/1425, H04L63/1441

FI/F-TERMS:

KEYWORDS OR OTHER ELEMENTS FEATURING THE INVENTION:

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

PCT

To:

Wright, Bradley
BANNER & WIT
1100 13th Street
Washington, DC
ETATS-UNIS D

DUE DATE

ART 19 AMEND
Due: 5-6-2017

NOTIFICATION OF TRANSMITTAL OF
THE INTERNATIONAL SEARCH REPORT AND
THE WRITTEN OPINION OF THE INTERNATIONAL
SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

Applicant's or agent's file reference
007742.00083

Date of mailing
(day/month/year)

6 March 2017 (06-03-2017)

FOR FURTHER ACTION

See paragraphs 1 and 4 below

International application No.
PCT/US2016/068008

International filing date
(day/month/year)

21 December 2016 (21-12-2016)

Applicant

CENTRIPETAL NETWORKS, INC.

1. ☒ The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

Filing of amendments and statement under Article 19:

The applicant is entitled, if he so wishes, to amend the claims of the International Application (see Rule 46):

When? The time limit for filing such amendments is normally two months from the date of transmittal of the International Search Report.

How? Directly to the International Bureau of WIPO, 34 chemin des Colombettes
1211 Geneva 20, Switzerland, Facsimile No.: (41-22) 338.82.70

For more detailed instructions, see PCT Applicant's Guide, International Phase, paragraphs 9.004 - 9.011.

2. ☐ The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.
3. ☐ **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

☐ the protest together with the decision thereon has been transmitted to the International Bureau together with any applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

☐ no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. Reminders

The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. These comments will be made available to the public after international publication. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established.

Shortly after the expiration of 18 months from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before the completion of the technical preparations for international publication (Rules 90bis.1 and 90bis.3).

Within 19 months from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase until 30 months from the priority date (in some Offices even later); otherwise, the applicant must, within 20 months from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices. In respect of other designated Offices, the time limit of 30 months (or later) will apply even if no demand is filed within 19 months. For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the PCT Applicant's Guide, National Chapters.

Within 19 months from the priority date, the applicant may request that a supplementary international search be carried out by a different International Searching Authority that offers this service (Rule 45bis.1). The procedure for requesting supplementary international search is described in the PCT Applicant's Guide, International Phase, paragraphs 8.006-8.032.

Name and mailing address of the International Searching Authority



European Patent Office, P.B. 5818 Patentlaan 2
NL-2280 HV Rijswijk
Tel. (+31-70) 340-2040
Fax: (+31-70) 340-3016

Authorized officer

EOZENOU, Bernard
Tel: +49 (0)89 2399-6022

RECEIVED

MAR 17 2017

MAR 17 2017

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference 007742.00083	FOR FURTHER ACTION see Form PCT/ISA/220 as well as, where applicable, item 5 below.	
International application No. PCT/US2016/068008	International filing date (day/month/year) 21 December 2016 (21-12-2016)	(Earliest) Priority Date (day/month/year) 4 January 2016 (04-01-2016)
Applicant CENTRIPETAL NETWORKS, INC.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 4 sheets.

☒ It is also accompanied by a copy of each prior art document cited in this report.

1. Basis of the report

a. With regard to the **language**, the international search was carried out on the basis of:

- ☒ the international application in the language in which it was filed
☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))

b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6bis(a)).

c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. ☐ **Certain claims were found unsearchable** (See Box No. II)

3. ☐ **Unity of invention is lacking** (see Box No. III)

4. With regard to the **title**,

- ☒ the text is approved as submitted by the applicant
☐ the text has been established by this Authority to read as follows:

5. With regard to the **abstract**,

- ☒ the text is approved as submitted by the applicant
☐ the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings**,

- a. the figure of the **drawings** to be published with the abstract is Figure No. 3
☐ as suggested by the applicant
☐ as selected by this Authority, because the applicant failed to suggest a figure
☒ as selected by this Authority, because this figure better characterizes the invention
b. ☐ none of the figures is to be published with the abstract

RECEIVED

MAR 17 2017

BANNER REGISTRATION

Page 1681

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2016/068008

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L29/06

ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L G06F

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

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

EPO-Internal, WPI Data, COMPENDEX, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/350229 A1 (MITCHELL DAVID JAMES [US]) 3 December 2015 (2015-12-03) abstract; figures 1, 3 paragraphs [0012], [0014], [0026], [0030] - [0034], [0038], [0039], [0066] paragraphs [0074], [0080], [0083], [0086] -----	1-7
A	US 2005/108557 A1 (KAYO DAVID G [US] ET AL) 19 May 2005 (2005-05-19) abstract; figures 3, 6, 7 paragraphs [0004], [0010] - [0016], [0019] - [0021], [0027] - [0029], [0058] - [0059] paragraphs [0062], [0063] ----- -/-	1-7



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

24 February 2017

Date of mailing of the international search report

06/03/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Hristova, Ana

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/068008

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2015/373043 A1 (WANG JISHENG [US] ET AL) 24 December 2015 (2015-12-24) abstract; figures 1, 2A, 2B, 3 paragraphs [0002], [0034] - [0042], [0044] - [0057], [0063] - [0072] -----	1-7

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/068008

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015350229 A1	03-12-2015	NONE	
US 2005108557 A1	19-05-2005	NONE	
US 2015373043 A1	24-12-2015	US 2015373043 A1	24-12-2015
		WO 2016025081 A1	18-02-2016

Information on Search Strategy - Pilot phase (see OJ 2015, A86)

The type of information contained in this sheet may change during the pilot for improving the usefulness of this new service.

Application Number

PCT/US2016/068008

TITLE: EFFICIENT PACKET CAPTURE FOR CYBER THREAT ANALYSIS

APPLICANT: CENTRIPETAL NETWORKS, INC.

IPC CLASSIFICATION: H04L29/06

EXAMINER: Hristova, Ana

CONSULTED DATABASES: WPI, COMPDX, INSPEC, TDB, XP3GPP, XPCRYPT, XPESP, XPETSI, XPI3E, XPIEE, XPIETF, XPLNCS, XPNPL, EPODOC

CLASSIFICATION SYMBOLS DEFINING EXTENT OF THE SEARCH:

IPC:

CPC: H04L63/0236, H04L63/0263, H04L63/1408, H04L63/1416, H04L63/20, H04L63/101, G06F21/566

FI/F-TERMS:


KEYWORDS OR OTHER ELEMENTS FEATURING THE INVENTION:

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

PCT

To:



WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43bis.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.
PCT/US2016/068008

International filing date (day/month/year)
21.12.2016

Priority date (day/month/year)
04.01.2016

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:


- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☒ Box No. VII Certain defects in the international application
- ☒ Box No. VIII Certain observations on the international application

2. FURTHER ACTION

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1bis(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.



Name and mailing address of the ISA:



European Patent Office
D-80298 Munich
Tel. +49 89 2399 - 0
Fax: +49 89 2399 - 4465

Date of completion of
this opinion

see form
PCT/ISA/210

Authorized Officer

Hristova, Ana

Telephone No. +49 89 2399-0



Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:
 - ☒ the international application in the language in which it was filed.
 - ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).
2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))
3. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing:
 - a. ☐ forming part of the international application as filed:
 - ☐ in the form of an Annex C/ST.25 text file.
 - ☐ on paper or in the form of an image file.
 - b. ☐ furnished together with the international application under PCT Rule 13*ter*.1(a) for the purposes of international search only in the form of an Annex C/ST.25 text file.
 - c. ☐ furnished subsequent to the international filing date for the purposes of international search only:
 - ☐ in the form of an Annex C/ST.25 text file (Rule 13*ter*.1(a)).
 - ☐ on paper or in the form of an image file (Rule 13*ter*.1(b) and Administrative Instructions, Section 713).
4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that forming part of the application as filed or does not go beyond the application as filed, as appropriate, were furnished.
5. Additional comments:

**WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY**

International application No.
PCT/US2016/068008

Box No. V Reasoned statement under Rule 43bis.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes: Claims	<u>2-7</u>
	No: Claims	<u>1</u>
Inventive step (IS)	Yes: Claims	
	No: Claims	<u>1-7</u>
Industrial applicability (IA)	Yes: Claims	<u>1-7</u>
	No: Claims	

2. Citations and explanations

see separate sheet

Box No. VII Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Box No. VIII Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 Reference is made to the following documents:

- D1 US 2015/350229 A1 (MITCHELL DAVID JAMES [US]) 3 December 2015 (2015-12-03)
- D2 US 2005/108557 A1 (KAYO DAVID G [US] ET AL) 19 May 2005 (2005-05-19)
- D3 US 2015/373043 A1 (WANG JISHENG [US] ET AL) 24 December 2015 (2015-12-24)

2 Independent claims

- 2.1 The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of claim 1 is not new in the sense of Article 33(2) PCT.
- 2.2 Document D1 discloses all the features of claim 1 (the references in parentheses applying to this document):

A method comprising:

receiving, by a computing device, one or more threat indicators from one or more threat intelligence providers, wherein the computing device is located at a boundary between a protected network associated with an enterprise and an unprotected network ([0074], lines: 1-7; [0039], lines: 1-11);

in response to receiving the threat indicators, combining the threat indicators based on a source address and port, a destination address and port, and a protocol type associated with the threat indicators ([0066], lines: 1-7; [0032], lines: 13-18);

generating one or more packet capture and packet filtering rules based on the combined threat indicators ([0038], lines: 7-10; [0074], lines: 12-14); and

filtering, by the computing device, on a packet-by-packet basis, at least one packet based on at least one of the one or more packet capture and packet filtering rules ([0083], lines: 3-12; [0039], lines: 13-16; [0080]).

Hence, the subject-matter of claim 1 does not meet the requirements of the PCT concerning novelty.

- 2.3 Document D1 also discloses the following features of claim 7 (the references in parentheses applying to this document):

A method comprising:

generating, by a computing device, a packet capture file comprising raw packet content and threat context information, wherein the threat context information comprises a filtering rule and associated threat indicator that caused a packet to be captured ([0033], lines: 1-5).

The difference between document D1 and the subject-matter of claim 7 resides in generating a threat context information, wherein the threat context information comprises a filtering rule and associated threat indicator that caused a packet to be captured.

The objective problem to be solved by claim 7 may therefore be regarded as how to associate a packet with its corresponding detected threats.

However, generating a threat context information, wherein the threat context information comprises a filtering rule and associated threat indicator that caused a packet to be captured is a common design measure for associating a packet with its corresponding detected threats. This measure is also known from D2 (see e.g. [0058], [0059]). Hence, the subject-matter of claim 7 does not meet the requirements of the PCT concerning inventive step.

3 Dependent claims

Additionally, the dependent claims do not appear to contain any additional technical features which, either alone or in combination with the features of any claim to which they refer, render the present subject-matter novel (Article 33(2) PCT) or inventive (Article 33(3) PCT) since their features (e.g. combining of the threat indicators based on a hostname or an URI etc.) are either known from the prior art (documents D1-D3; see cited passages in the International Search Report) or merely represent minor implementation details to the person skilled in the art.

Re Item VII

Certain defects in the international application

- 1 The independent claims are not in the two-part form vis-à-vis D1 in accordance with Rule 6.3(b) PCT, which in the present case would be appropriate.
- 2 The features of the claims are not provided with reference signs placed in parentheses to increase the intelligibility of the claims (Rule 6.2(b) PCT).

- 3 Contrary to the requirements of Rule 5.1(a)(ii) PCT, the relevant background art disclosed in D1 is not mentioned in the description, nor is this document identified therein.
- 4 The expression "incorporated by reference" in paragraphs [0001], [0027], [0029], [0033], [0045] and [0048] of the description does not render the application self-contained according to Article 5 PCT.

Re Item VIII

Certain observations on the international application

- 1 The statement "within the scope and spirit of the appended claims" in the description in paragraph [0066] on lines: 2-3 implies that the subject-matter for which protection is sought may be different to that defined by the claims, thereby resulting in lack of clarity (Article 6 PCT) when used to interpret them.
- 2 It is unclear in claims 1 and 7 whether the protection sought is limited to the method per se or whether protection is sought for use in a specific device. For the sake of clarity, this claim should be directed to a method for use in a computing device (Article 6 PCT).
- 3 A subset of claims of the present application is not concise since the application comprises two independent claims of the same category, namely claims 1 and 7, neither defining complementary functions nor reflecting separate alternative technical solutions (Article 6 PCT). The number of independent claims should be limited to one independent claim in the method category.