



- (51) **International Patent Classification:**
H04L 12/935 (2013.01) H04L 12/771 (2013.01)
H04L 12/947 (2013.01)
- (21) **International Application Number:**
PCT/US2016/061897
- (22) **International Filing Date:**
14 November 2016 (14.11.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/255,330 13 November 2015 (13.11.2015) US
15/351,039 14 November 2016 (14.11.2016) US
- (71) **Applicant:** YAANA TECHNOLOGIES LLC [US/US];
542 Gibraltar Drive, Milpitas, CA 95035 (US).
- (72) **Inventors:** PORTER, Leigh, Alexander; Twynj olm
Church, 324 Lillie Road, London SW6 7PP (GB).
SCHWADERER, Curt; 4738 92nd Court, Urbandale, IA
50322 (US). HAMMER, Michael, P.; 1855 Saint Frances
Street, Apt. #708, Reston, VA 20190 (US). LIU, Kai,
Bang; 13449 Nw Alvada Street, Portland, OR 97229 (US).
PURI, Rajesh; 835 Yakima Drive, Fremont, CA 94536
(US). GROOTWASSINK, David; 103 Wateredge Court,
Safety Harbor, FL 34695 (US).


(74) **Agents:** DUTTA, Sanjeet et al.; Steptoe & Johnson Llp,
633 W. 5th Street, Suite 700, Los Angeles, CA 90071
(US).

(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, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,
KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
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, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

Published:

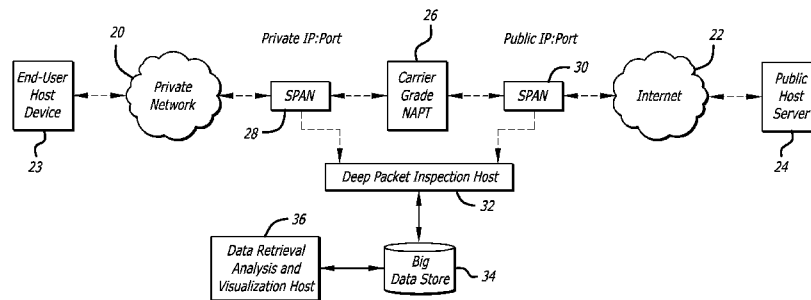
— with international search report (Art. 21(3))



WO 2017/083855 A1

(54) **Title:** SYSTEM AND METHOD FOR DISCOVERING INTERNET PROTOCOL (IP) NETWORK ADDRESS AND PORT TRANSLATION BINDINGS

FIG. 1



(57) **Abstract:** A system and method for discovering Internet Protocol (IP) network address and port translation bindings is disclosed. According to one embodiment, a system for creating IP data records, includes a carrier grade network address and port translation (CGNAPT) node in communication with a private network on an internal side and a public network on an external side. The system further includes a private Ethernet span port disposed on the internal side of the CGNAPT node that receives packet streams traversing the CGNAPT node. A public Ethernet span port is disposed on the external side of the CGNAPT node that receives packet streams traversing the CGNAPT node. The private Ethernet span port and the public Ethernet span port copy packet streams traversing the CGNAPT node and deliver the copied packet streams to a deep packet inspection host.

**SYSTEM AND METHOD FOR DISCOVERING INTERNET PROTOCOL (IP)
NETWORK ADDRESS AND PORT TRANSLATION BINDINGS**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 62/255,330, filed November 13, 2015, which is herein incorporated by reference in its entirety.

FIELD

[0002] The present disclosure relates in general to computer systems. In particular, the present disclosure relates to a system and method for discovering Internet Protocol (IP) network address and port translation bindings.

BACKGROUND

[0003] The exhaustion of IPv4 address space has driven many Communication Service Providers (CSP) to deploy Carrier Grade Network Address Translation (CGNAT) gateways within and on the edges of their networks. These gateways enable a large number of internally routable IP addresses assigned by the CSP to its customers from a private address pool to temporarily be bound to a much smaller number of publicly routable IP addresses allocated to the CSP, enabling hosts on the Internet to reach the CSP's customers and vice versa.

[0004] The maximum number of concurrently assigned address bindings possible is limited by the total number of public IP addresses, which may be owned by the CSP. As a result, many gateways also use the 65,535 possible ports for each IP address as a means to extend the possible number of concurrent bindings. These are called Network Address and Port Translation (NAPT) gateways. Furthermore, to maximize the usage of the public IP and Port range, these bindings are kept only so long as there is traffic traversing the "pinholes" through the NAPT gateway, *i.e.* the binding between the private and public addresses.

[0005] In addition, in order to solve the limits of IPv4, the IETF created IPv6 with a greatly expanded address space. However, that also leads to the issue of how to transition both clients on customer hosts and servers supporting applications on company hosts as they independently transition from IPv4 to IPv6. That transition has led to a plethora of NAPT-based services to map between the two address types, such as NAT64 and NAT464, to enable IPv4 clients to access IPv6 servers, IPv6 clients to access IPv4 servers, IPv4 hosts to communicate over IPv6

intermediate networks, and IPv6 hosts to communicate over IPv4 networks. All such address translations need to be supported and logged.

[0006] Unfortunately, many of these gateways have substantially inadequate logging facilities resulting in major logging errors or even no logs even at utilizations substantially lower than the platform maximum. Because the primary purpose of NAT is to set up and tear down bindings (pinholes), when CPU and memory resources become constrained during busy periods, secondary processes like logging stop. When logging and accounting record generation stops, that leaves network operators blind as to what traffic is entering and leaving their network and where that traffic is terminating inside their networks. As a result, any sort of forensic traffic analysis is not possible.

SUMMARY

[0007] A system and method for discovering Internet Protocol (IP) network address and port translation bindings is disclosed. According to one embodiment, a system for creating IP data records, comprises a carrier grade network address and port translation node in communication with a private network on an internal side and a public network on an external side, packet streams transmitted between the private network and the public network traverse the carrier grade network address and port translation node. The system further comprises a private Ethernet span port disposed on the internal side of the carrier grade network address and port translation node receives packet streams traversing the carrier grade network address and port translation node. A public Ethernet span port is disposed on the external side of the carrier grade network address and port translation node receiving packet streams traversing the carrier grade network address and port translation node. A deep packet inspection host is in communication with the private Ethernet span port and the public Ethernet span port. The private Ethernet span port and the public Ethernet span port copy packet streams traversing the carrier grade network address and port translation node and deliver the copied packet streams to the deep packet inspection host.

[0008] Other features and advantages will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, which illustrate by way of example, the features of the various embodiments.

BRIEF DESCRIPTION OF THE DRAWING

[0009] FIG. 1 depicts a simple example of a carrier grade Network Address and Port Translation gateways external logging system.

[0010] FIG. 2 depicts a nested example of a carrier grade Network Address and Port Translation gateways external logging system.

[0011] FIG. 3 depicts a complex example of a carrier grade Network Address and Port Translation gateways external logging system.

[0012] FIG. 4 depicts an example of packet extraction from public and private network sources and distribution by brokers such that related sets of private and public packets can be analysed and logged by load-balanced DPI hosts to compose a carrier grade Network Address and Port Translation gateway external logging system.

[0013] FIG. 5 depicts a flow chart for one embodiment of a process of correlation of private outbound packets and the public outbound packets to determine CGNAPT bindings.

DETAILED DESCRIPTION

[0014] A system and method for discovering Internet Protocol (IP) network address and port translation bindings is disclosed. According to one embodiment, a system for creating IP data records, comprises a carrier grade network address and port translation node in communication with a private network on an internal side and a public network on an external side, packet streams transmitted between the private network and the public network traverse the carrier grade network address and port translation node. The system further comprises a private Ethernet span port disposed or deployed on the internal side of the carrier grade network address and port translation node receives packet streams traversing the carrier grade network address and port translation node. A public Ethernet span port is disposed on the external side of the carrier grade network address and port translation node receiving packet streams traversing the carrier grade network address and port translation node. A deep packet inspection host is in communication with the private Ethernet span port and the public Ethernet span port. The private Ethernet span port and the public Ethernet span port copy packet streams traversing the carrier grade network address and port translation node and deliver the copied packet streams to the deep packet inspection host.

[0015] Briefly, and in general terms, various embodiments are directed to a system for creating IP data records. In one embodiment, the system includes a carrier grade network address and port translation (CGNAPT) node in communication with a private network on an internal side and a public network on an external side. Packet streams transmitted between the private network and the public network traverse the CGNAPT node. The system also includes a private Ethernet span port disposed on the internal side of the CGNAPT node and receives packet streams traversing the CGNAPT node. There is a public Ethernet span port disposed on the external side of the CGNAPT node that receives packet streams traversing the CGNAPT node. The system of one embodiment also includes a deep packet inspection host in communication with the private Ethernet span port and the public Ethernet span port. The deep packet inspection host may be a hierarchical configuration of multiple deep packet inspection hosts or a single host. The private Ethernet span port and the public Ethernet span port copy packet streams traversing the CGNAPT node and deliver the copied packet streams to the deep packet inspection host. In one embodiment of the system, the deep packet inspection host

correlates and identifies packet streams entering the CGNAPT node with packet streams exiting the carrier CGNAPT node. There may be any number of private or public networks with a plurality of CGNAPT nodes positioned between the plurality of public and private networks.

[0016] In one embodiment, the system may include a firewall combined with the CGNAPT node. The firewall may modify elements of the packet streams. The system may further include one or more host devices connected to the private network that request the CGNAPT node to open a binding with a server connected to the public network. A communications service provider may assign private IP addresses to the host devices, and the CGNAPT node may assign a public IP address and create a binding between the private IP address and the public IP address. In yet another embodiment, the CGNAPT node sets a timer to remove the binding between the private IP address and the public IP address.

[0017] The system may also include a data store in communication with the deep packet inspection host for storing the copied streaming packets or the final end-to-end path record of the processed packet streams. A data retrieval host may also be included in the system that is in communication with the data store for analyzing the copied streaming packets. The data retrieval host may be any local or remote computing device.

[0018] Other embodiments are directed to a method for correlating private outbound packets and public outbound packets to determine the carrier grade network address and port translation bindings. The method may include receiving inbound stream packets and outbound stream packets at a deep packet inspection host from a private Ethernet span port and a public Ethernet span port. The private Ethernet span port may be disposed on an internal side of a CGNAPT node that receives packet streams traversing the CGNAPT node between a private network and a public network. The public Ethernet span port may be disposed on the external side of the CGNAPT node receiving packet streams traversing the CGNAPT node. The method includes discarding inbound stream packets collected at the deep packet inspection host and distinguishing between private outbound stream packets and public outbound stream packets. The method may also include ordering by time private outbound stream packets in a lookback buffer. Furthermore, the method includes processing the public outbound stream packets to create bindings, reading the private outbound stream packets in the lookback buffer and processing the private outbound stream packets to create bindings. The log bindings of the

public outbound stream packets and the private outbound stream packets may then be stored in memory.

[0019] In certain embodiments, the method includes determining the content size, sequence number and window size of the public outbound stream packets. In addition, destination IP addresses, destination port, and ACK number are also used for correlation. The method may include comparing a portion of one private outbound stream packet with a portion of one public outbound stream packet to determine a match.

[0020] One embodiment of the present system and method is an off-gateway Network Address and Port Translation (NAPT) logging solution that monitors private-side and public-side NAPT interfaces. The present system may output CDR-style Internet “Connection” Records (ICR) that matches internally, and possibly externally, visible Internet traffic sessions so that their original IP address and ports and all translations to external IP address and ports are recorded. The system may output other format types of records that are known in the art. Additional correlation of Communication Service Providers (CSP) assigned IP addresses with other traffic data, *e.g.* subscriber identifiers such as their MSISDN or Username are also may be performed. Combined with a high volume data retention and request management system, the present system and method is able to deliver an end to end system for NAT record extraction, retention, analysis, and disclosure. The present system and method uses various processes to reduce log volume by identifying parallel related flows belong to the same session.

[0021] In one exemplary embodiment as shown in FIG. 1, the system includes a single private internal network 20 and one public external network 22. A CSP assigns private IP addresses to end-user host devices 23 that may use one or more port addresses for each active application depending on how many communications sessions it needs to establish. Traffic leaving the private network 20 toward public host server 24 on the external network traverses the Carrier Grade Network Address and Port Translation gateway (CGNAPT) 26. The CGNAPT 26 temporarily assigns a public IP and port and creates a binding between the two addresses. The CGNAPT 26 sets a timer that upon expiry removes the binding. The timer may be set any time as desired and may be dynamic. Traffic through the binding may reset the timer. In one embodiment, TCP traffic may have a different behavior than UDP, since explicit TCP connection release traffic could also be relied on to remove the binding. In addition, some traffic

carried over the CGNAPT 26 may not use or expose UDP or TCP ports, such as RSVP, ICMP, IPSec ESP, or GRE tunnels established by PPTP. The present system logs such traffic.

[0022] Still referring to FIG. 1, the system includes a private Ethernet Span port 28 on the internal side of the CGNAPT 26 and a public Ethernet Span port 30 on the external side of the CGNAPT. It has been contemplated that electrical or optical taps could be used instead. The private and public ports 28 and 30 may transparently copy all traffic to the CGNAPT 26 and deliver the copied data to a deep packet inspection (DPI) Host 32. The ports 28 and 30 may also copy all traffic from the CGNAPT. In other embodiments, selected traffic may be copied to (or from) the CGNAPT. In one embodiment, the DPI Host 32 correlates and identifies the packet streams entering the CGNAPT 26 with the packet streams exiting the CGNAPT.

[0023] As shown in FIG. 1, the DPI host 32 is in communication with a data store 34 that stores the copied data. A data retrieval analysis and visualization host 36 is connected to the data store 34 and may be used to retrieve and analyse the copied traffic. The data store 34 may use big data tools, such as Hadoop to store data, process requests in parallel and return results. Requests could for example return all private IP addresses associated with a public IP address and Port in a given timeframe. Searches based on private IP address and other packet parameters in some combination could also be performed. Analysis could be done iteratively to arrive at a better understanding of traffic patterns.

[0024] The nature of the packet streams is not necessarily one-to-one. NATs have numerous behaviors that may or may not allow one-to-many relationships between the internal and external addresses. It is also possible that the nature of the packet streams is such that the communications could be half-duplex or full-duplex. That is, an out-going stream may or may not correlate with any incoming stream, and vice versa.

[0025] In one embodiment, the CGNAPT 26 may be combined with a firewall that performs additional operations on the messages besides the modification of IP addresses and port numbers. For example, the firewall process may add, remove or modify other elements of the packet data headers. Since such changes are driven by the specific policies of a given network operator, they are non-deterministic and unpredictable.

[0026] In certain embodiments, there are protocols by which the users on the private network 20 may either request the CGNAPT 26 to open a binding or to communicate with a server on the

public side of the CGNAPT to punch a hole (open a binding) through the CGNAPT and learn of the public address of that binding, such a NAT-PTP, STUN, TURN, ICE, *etc.*

[0027] Another embodiment of a system is shown in FIG. 2. In this nested example, there are multiple private networks with CGNAPT between them as well as a CGNAPT between the private and public globally routable Internet. This may occur in very large networks where regional access networks may be aggregated by an overall national network before reaching a publicly routable international network like the Internet. FIG. 2 shows a system that includes a private access network 40 and a private aggregation network 42 separated by a first CGNAPT 44. A second CGNAPT 46 separates the private aggregation network 42 from a public network 48. The system shown in FIG. 2 also includes a first private Ethernet Span port 50 on the internal side of the first CGNAPT 44 and a public Ethernet Span port 52 on the external side of the second CGNAPT 46. It has been contemplated that electrical or optical taps could be used instead. In certain embodiments, there is also a second private Ethernet Span port 54 on the aggregation network 42 side of all gateways.

[0028] The present methods are used even if multiple networks and gateways are encompassed, since the aggregation network and the gateways would just appear like a very large gateway. Also the system depicted in FIG. 2 iteratively correlates traffic and records. As shown, the private and public ports 50 and 52 may transparently copy all traffic to the first and second CGNAPT 44 and 46, and deliver the copied data to a deep packet inspection (DPI) host 56. The ports 50 and 52 may also copy all traffic from the CGNAPT. In an embodiment including second private ports 54 on the aggregation network side of the gateways, the second private port 54 also copies all traffic to and from the CGNAPT 44 and 46 and send the copied traffic to the DPI host 56. In other embodiments, selected traffic may be copied to (or from) the CGNAPT. In one embodiment, the DPI host 56 correlates and identifies the packet streams entering the first and second CGNAPT 44 and 46 with the packet streams exiting the CGNAPT 44 and 46.

[0029] As shown in FIG. 2, the DPI host 56 is in communication with a data store 58 that stores the copied data. A data retrieval analysis and visualization host 60 is connected to the data store 58 and may be used to retrieve and analyse the copied traffic.

[0030] In addition, though the DPI host 56 is shown as a single host, multiple hosts in clusters or in different hierarchical configurations may be used to process traffic and intermediate records to result in the final end-to-end path record.

[0031] In yet another embodiment shown in FIG. 3, a complex system involves the multiplicity of traffic streams that may exist amongst three or more hosts and three or more private and public networks. A given packet stream could traverse a single private network, multiple private networks, and combinations of multiple private access, multiple aggregations networks, and multiple public networks.

[0032] Similar to previous embodiments, the system of FIG. 3 shows span ports 70, 72, 74, and 76, at the inner-most and outer-most edges of the gateways. As with the other embodiments, the span ports could be on any inside or outside of any gateway. Also similar to the previous embodiments, a DPI host 78 represents the multiple possible configurations of DPI hosts needed to analyze the entire CSP network. In FIG. 3, the End-User host devices 80 and 82 are shown connected to the Private Access Networks 84 and 86, respectively, but they could also be connected to the Private Aggregation Network 88. In certain embodiment, there could be more than one level of access or aggregation networks. Also, as shown, there are public host servers 90 and 92 connected to public networks 94 and 96, respectively. Multiple CGNAPT 98, 100, 102, and 104 are located between the private and public networks 84, 86, 88, 94, and 96 as shown in FIG. 3. Similar to the other embodiments, the DPI host 78 is in communication with a data store 106 that stores the copied data. A data retrieval analysis and visualization host 108 is connected to the data store 58 and may be used to retrieve and analyse the copied traffic.

[0033] The present methods enable the derivation of logging records for any communications patterns of any combination of hosts, on any combination of networks, interconnected by any types of CGNAPT.

[0034] In general, current systems rely on the CGNAPT/firewall generating logs of the bindings as they occur and exporting those bindings as communications detail records, similar to call detail records (CDR) for telephone systems. However, those logs are not crucial to the billing of the customer, and are forfeited if computer and storage resources are needed for more basic processes associated with service delivery. Thus, while binding are still setup and released, no logging occurs.

[0035] The present system and method enables discovery and logging by a node external to the CGNAPT through the use of DPI analysis of the packet streams on both the private and public sides of the CGNAPT. The present system associates the more permanent private IP address and temporary port of the end-user with the temporary IP and port appearing on the public Internet, or the IPX exchange between operators.

[0036] Embodiments of the current system have at least two points for intercepting and copying the packet streams and delivering them to a common DPI host. At least one intercept point is on each network of an end-user and covers all egress from that network through CGNAPT to other networks. At least one intercept point will be on each egress from an operator's network to the Internet (public network) or packet exchanges beyond which the operator has no further control. Additional intercept points may be on links by CGNAPT to intermediate networks. More than one DPI host may be used and interconnected via a common data store.

[0037] In the embodiments shown in FIGS. 1-3, the DPI host 32, 56, and 78 is a single node. However, due to scaling and operator routing issues, the DPI host may really be a complex of nodes as shown in FIG. 4. In one embodiment, the private side traffic 120 could originate or terminate on many private VLANs distributed across an operator's network. Similarly the public side traffic 122 could enter or leave the operator's network at various egress points. Once intercepted, the replicated traffic is unidirectionally routed to the correlation complex.

[0038] Still referring to FIG. 4, in one embodiment, private Ethernet Span ports 124 are in communication with a CGNAPT and a first packet router broker 126. Public Ethernet Span ports 128 are in communication with a separate CGNAPT and a second packet router broker 130. The role of the packet router brokers 126 and 130 is to distribute the traffic across multiple DPI hosts 132 while ensuring that both the private side traffic 120 and public side traffic 122 of a CGNAPT corresponding to each unique end-to-end packet session go to the same DPI host. That is accomplished by using the Least Significant Bits (LSB) of the Destination IP Address of Out-Bound packets, which is identical on both Private and Public side of the CGNAPT. Similarly, for In-Bound packets, we use the LSB of the Source IP Address. From the DPI hosts 132, the traffic is stored in a data store 134, which can be accessed by a local or remote computer for retrieval and analysis.

[0039] In one embodiment, VLAN tags determine whether the packets are In-Bound or Out-Bound. The VLAN tags are different for In-Bound or Out-Bound, on both the private and public sides of the CGNAPT. That leads to a second function of the packet router brokers 126 and 130, which is to map the In-Bound/Out-Bound VLANs coming from the intercepted network to In-Bound and Out-Bound VLANs to the DPI hosts 132. In that way, the DPI hosts 132 can focus on determining the CGNAPT bindings using only the Out-Bound packets, e.g. TCP SYN.

[0040] In certain embodiments, private side (of the CGNAPT) traffic packets are determined through provisioning of the Private Address spaces used by the operator. Likewise, the public side (of the CGNAPT) traffic packets may be identified through the provisioning of the Public IP Address space assigned to the CGNAPT. When the packet stream arrives at the packet broker, the packet broker needs to understand if the source is on the private or public side of the CGNAPT.

[0041] Any in-bound packet that can get through without prior out-bound packets does so via static provisioned holes, and as such are known and therefore do not need to be discovered. This type of permanent binding may be monitored, and if necessary measured, counted, and logged at the packet router broker. Such traffic may be discarded by the packet router broker to reduce load on the DPI host.

[0042] In certain embodiments, the packet router brokers also may be controlled by the DPI host to measure, count, log, and discard or truncate mid-call flows, or packets of a certain type, such as video frames, to reduce load on the DPI host.

[0043] The distribution by packet routers broker 126 and 130 across multiple DPI hosts 132 may be done by using any number of LSBs of the IP addresses. By way of example only, the following shows how 3 bits determine what traffic is directed to which DPI host.

DPI Host 1 receives if:

LSB=000

LSB=001

LSB=010

DPI Host 2 receives if:

LSB=011

LSB=100

LSB=101

DPI Host 3 receives if:

LSB=110

LSB=111

[0044] In another embodiment, variable numbers of bits could be used to determine what traffic is directed to which DPI host so long as the full number space is covered as shown in the following:

DPI Host 1 receives if: LSB=0

LSB=000

LSB=100

LSB=010

LSB=110

DPI Host 2 receives if: LSB=01

LSB=001

LSB=101

DPI Host 3 receives if: LSB=11

LSB=011

LSB=111

[0045] Referring now to FIG. 5, a flow chart of one embodiment of a process of correlation of the private outbound packets and the public outbound packets to determine the CGNAPT bindings is shown. As shown in the flow chart, the first step of the process is to filter out the inbound packets either using the VLAN tags or identifying that the destination address belongs to

the CGNAT or the private IP address space used by the operator. As shown, at step 200 a packet is received and it is determined at step 202 if it is an out-bound packet. If the packet is not an out-bound packet, but rather an in-bound packet, the in-bound packet is discarded at step 204. Next, the process distinguishes the private out-bound packets from the public out-bound packets at step 206. As shown in the flow chart of FIG. 5, the private-side packets are ordered by time at step 208 and maintained in a time-based multi-map lookback buffer at step 210. Private-side packets exceeding the look-back times may eventually be discarded, however, in other embodiments the private-side packets may have no time limit for being discarded.

[0046] It is then determined at step 212 if the public-side packets include an existing binding, and if not, the packets are ordered by 5-Tuple at step 214 and maintained in a 5-Tuple hash map buffer at step 216 as shown in FIG. 5. At step 218, the content size, sequence number, and window size of the public-side packets are determined, and then at step 220, it is determined if there are any candidates for the public-side packet. In certain embodiments, destination IP addresses, destination port, and ACK number are also used for correlation. If there are no candidates, the packets are logged as errors at step 222. Public-side packets with candidates are then processed to create bindings. At step 224, the content of the public-side packet is compared byte by byte, and if there is no match at step 226, the unmatched public-side packet is logged as an error at step 222. The byte-by-byte comparison of the remaining candidates of the packet content at step 224 is performed in order to eliminate false positives and determine a match. Some minimum number of bytes beyond which a single candidate is found or until the end of the content to confirm the match. Often 64 bytes is sufficient, however, this may vary.

[0047] If there is a match at step 226, then a log binding is created at step 228, which is logged in the CGNAPT at step 230. After that the DPI host may use the information in the log to inform the DPDI host data plane (cavium code) to discard any additional mid-session packets associated with those bindings. It is possible that packets may be queued before the previous steps to discover the binding and before the data plane is informed. But since the public hash table entry is marked as correlation done with private side IP/port recorded (this hash table is also the CGNAPT binding table), any queued public packets belong to the binding would not trigger another round of correlation search. These public packets would be dropped in one embodiment.

[0048] For both public and private packets, packets of a flow may be dropped except for packets 1, 2, and 3. Packet 4 also may not be dropped if the flow ends with TCP FIN packet. The system may determine to analyze another 3 packets for both private and public side flows if packet 1, 2, and 3 fail to yield a successful correlation. Correlation can be initiated from either private or public side flows on reception of each packet. On the public side, packet 1, 2, and 3 also are stored in the lookback buffers. Comparison is not ordered. This is very useful to correlate out of order packets or out of order private and public flow arriving time. In one embodiment, incoming packets are kept for a short period of time to determine whether the packets are needed. During the desired amount of time, the packets are used for correlation or the packets are discarded. It is sufficient to determine a correlation between private and public flow using the disclosed metrics.

[0049] In this embodiment of the process, a timer based on an over-sizing factor of the provisioned or learned CGNAPT traversal times is used to fetch previously received private out-bound packets, to compare similar values sequentially to eliminate mismatches and narrow down the pool of possible matches. The private-side packets logged in the lookback buffer are read within lookback time frame at step 232 and are processed to receive a binding using steps 220-230 as described above. Normally, a match is found on average around the average latency across the CGNAPT. In one embodiment, conservative look-back accounts are used for slower than normal latency times. Error conditions or missed private side packets may occur. However, if an error does occur, a new public side packet arriving will trigger another correlation operation as described. As long as one packet in a public 5-tuple based flow gets correlated, the job is done for the flow. In this way, the DPI host system itself has very large error recovery margin.

[0050] If a binding for the public out-bound packet does exist at step 212, it is determined if the binding is TCP FIN or RST at step 234, then the binding is removed at step 236 and logged in the CGNAPT at step 230. Otherwise, if the binding is determined to be UDP at step 238, then the UDP timeout is reset at step 240. If the UDP timeout at step 242, then the binding is removed at step 236 and logged in the external solution that discovers and logs the start and stop of bindings in the CGNAPT at step 230. If there is also no UDP, then the public out-bound packet is discarded at step 204. That is, this example embodiment of an algorithm is demonstrating the tracking of TCP and UDP flows that create bindings, while discarding other types of packets not of interest. Note that other embodiments may also look for Stream Control

Transmission Protocol (SCTP) or other types of packet flows. TCP and UDP are shown here as they are the most widely used.

[0051] The various embodiments described above are provided by way of illustration only and should not be construed to limit the claimed invention. Those skilled in the art will readily recognize various modifications and changes that may be made to the claimed invention without following the example embodiments and applications illustrated and described herein, and without departing from the true spirit and scope of the claimed invention, which is set forth in the following claims.

WHAT IS CLAIMED:

1. A system for creating IP data records, comprising:
 - a carrier grade network address and port translation node in communication with a private network on an internal side and a public network on an external side, packet streams transmitted between the private network and the public network traverse the carrier grade network address and port translation node;
 - a private Ethernet span port disposed on the internal side of the carrier grade network address and port translation node receives packet streams traversing the carrier grade network address and port translation node;
 - a public Ethernet span port disposed on the external side of the carrier grade network address and port translation node receiving packet streams traversing the carrier grade network address and port translation node; and
 - a deep packet inspection host in communication with the private Ethernet span port and the public Ethernet span port;
 - wherein the private Ethernet span port and the public Ethernet span port copy packet streams traversing the carrier grade network address and port translation node and deliver the copied packet streams to the deep packet inspection host.
2. The system of claim 1, wherein the deep packet inspection host correlates and identifies packet streams entering the carrier grade network address and port translation node with packet streams exiting the carrier grade network address and port translation node.
3. The system of claim 1, further comprising a firewall associated with the carrier grade network address and port translation node that modifies elements of the packet stream.
4. The system of claim 1, further comprising host devices connected to the private network sending packets to the host server in the public network through the CGNAPT node triggering it to assign public address and port to that packet stream creating the binding.
5. The system of claim 4, further comprising a communications service provider that assigns private IP addresses to the host devices.

6. The system of claim 5, wherein the carrier grade network address and port translation node assigns a public IP address and creates a binding between the private IP address and the public IP address.
7. The system of claim 6, wherein the carrier grade network address and port translation node sets a timer to remove the binding between the private IP address and the public IP address.
8. The system of claim 1, further comprising a data store in communication with the deep packet inspection host for storing binding records of the copied streaming packets.
9. The system of claim 8, further comprising a data retrieval host in communication with the data store for analyzing the binding records of the copied streaming packets.
10. The system of claim 1, wherein the deep packet inspection host is a hierarchical configuration of multiple deep packet inspection hosts.
11. A method for correlating private outbound packets and public outbound packets to determine the carrier grade network address and port translation bindings, the method comprising:
 - receiving inbound stream packets and outbound stream packets at a deep packet inspection host from a private Ethernet span port and a public Ethernet span port, the private Ethernet span port disposed on an internal side of a carrier grade network address and port translation node that receives packet streams traversing the carrier grade network address and port translation node between a private network and a public network, and the public Ethernet span port disposed on the external side of the carrier grade network address and port translation node receiving packet streams traversing the carrier grade network address and port translation node;
 - discarding inbound stream packets collected at the deep packet inspection host;
 - distinguishing between private outbound stream packets and public outbound stream packets;
 - ordering by time private outbound stream packets in a lookback buffer;
 - processing the public outbound stream packets to create bindings;

reading the private outbound stream packets in the lookback buffer and processing the private outbound stream packets to create bindings;

storing the log bindings of the public outbound stream packets and the private outbound stream packets.

12. The method of claim 11, further comprising determining the content size, sequence number and window size of the public outbound stream packets.

13. The method of claim 11, further comprising comparing a portion of one private outbound stream packet with a portion of one public outbound stream packet to determine a match.

14. The method of claim 11, wherein filtering out the inbound packets using VLAN tags.

15. The method of claim 11, wherein filtering out the inbound packets by identifying a destination IP address.

16. The method of claim 11, further comprising buffering the public outbound steam packet according to normal 5-tuple.

FIG. 1





FIG. 2






FIG. 3


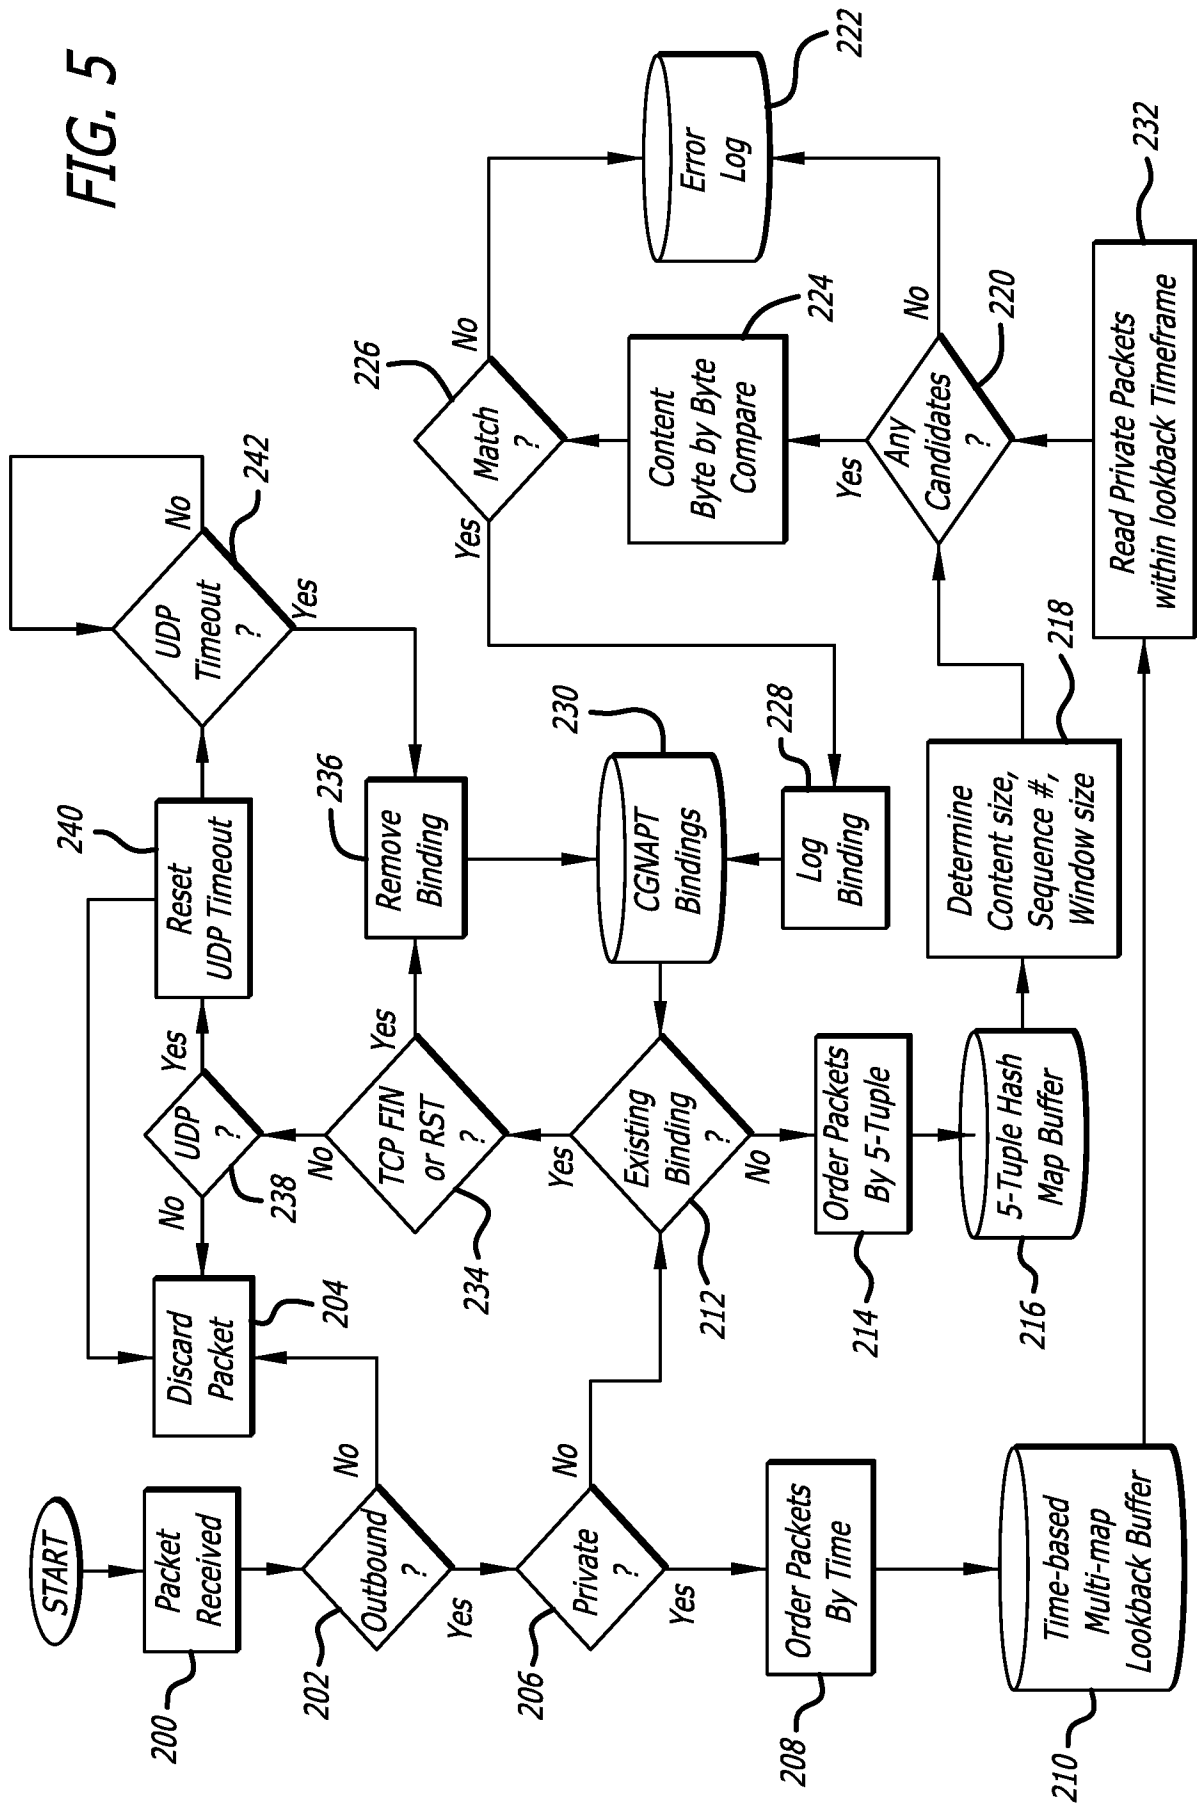


FIG. 4

FIG. 5



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US16/61897

A. CLASSIFICATION OF SUBJECT MATTER
 IPC - H04L12/935; H04L12/947; H04L12/771 (2017.01)
 CPC - H04L61/256; H04L61/2517; H04L63/02; H04L63/1408

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)
 See Search History document

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2011/0182183 A1 (PERKINS, C) 28 July 2011, fig. 1A, paragraphs [0030], [0035], [0038], [0061]	1-16
A	CN 104869181A (BEIJING GEO POLYMERIZATION TECHNOLOGY CO LTD)26 August 2015; see machine translation	1-16
A	US 2014/0226478 A1 (VMWARE, INC) 14 August 2014, fig. 3, paragraph [0013]	1-16
A	US 2013/0343388 A1 (STROUD, J et al.) 26 December 2013, paragraphs [0408], [0428]	1-16

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
 09 January 2017 (09.01.2017)

Date of mailing of the international search report

03 FEB 2017

Name and mailing address of the ISA/
 Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
 P.O. Box 1450, Alexandria, Virginia 22313-1450
 Facsimile No. 571-273-8300

Authorized officer
 Shane Thomas

PCT Helpdesk: 571-272-4300
 PCT OSP: 571-272-7774