interleaved encrypted block is broken into a series of payloads, one for each TARP packet to be generated. Special TARP headers $IP_T$ are then added to each payload using the IP headers from the data stream packets. The TARP headers can be identical to normal IP headers or customized in some way. They should contain a formula or data for deinterleaving the data at the destination TARP terminal, a time-to-live (TTL) parameter to indicate the number of hops still to be executed, a data type identifier which indicates whether the payload contains, for example, TCP or UDP data, the sender's TARP address, the destination TARP address, and an indicator as to whether the packet contains real or decoy data or a formula for filtering out decoy data if decoy data is spread in some way through the TARP payload data.

Note that although chain-block encryption is discussed here with reference to the session key, any encryption method may be used. Preferably, as in chain block encryption, a method should be used that makes unauthorized decryption difficult without an entire result of the encryption process. Thus, by separating the encrypted block among multiple packets and making it difficult for an interloper to obtain access to all of such packets, the contents of the communications are provided an extra layer of security.

Decoy or dummy data can be added to a stream to help foil traffic analysis by reducing the peak-to-average network load. It may be desirable to provide the TARP process with an ability to respond to the time of day or other criteria to generate more decoy data during low traffic periods so that communication bursts at one point in the Internet cannot be tied to communication bursts at another point to reveal the communicating endpoints.

Dummy data also helps to break the data into a larger number of inconspicuously-sized packets permitting the interleave window size to be increased while maintaining a reasonable size for each packet. (The packet size can be a single standard size or selected from a fixed range of sizes.) One primary reason for desiring for each message to be broken into multiple packets is apparent if a chain block encryption scheme is used to form the first encryption layer prior to interleaving. A single block encryption may be applied to portion, or entirety, of a message, and that



New Bay Capital, LLC
Ex.1015-Page 2401 of 3151

portion or entirety then interleaved into a number of separate packets. Considering the agile IP routing of the packets, and the attendant difficulty of reconstructing an entire sequence of packets to form a single block-encrypted message element, decoy packets can significantly increase the difficulty of reconstructing an entire data stream.

The above scheme may be implemented entirely by processes operating between the data link layer and the network layer of each server or terminal participating in the TARP system. Because the encryption system described above is insertable between the data link and network layers, the processes involved in supporting the encrypted communication may be completely transparent to processes at the IP (network) layer and above. The TARP processes may also be completely transparent to the data link layer processes as well. Thus, no operations at or above the Network layer, or at or below the data link layer, are affected by the insertion of the TARP stack. This provides additional security to all processes at or above the network layer, since the difficulty of unauthorized penetration of the network layer (by, for example, a hacker) is increased substantially. Even newly developed servers running at the session layer leave all processes below the session layer vulnerable to attack. Note that in this architecture, security is distributed. That is, notebook computers used by executives on the road, for example, can communicate over the Internet without any compromise in security.

IP address changes made by TARP terminals and routers can be done at regular intervals, at random intervals, or upon detection of "attacks." The variation of IP addresses hinders traffic analysis that might reveal which computers are communicating, and also provides a degree of immunity from attack. The level of immunity from attack is roughly proportional to the rate at which the IP address of the host is changing.

As mentioned, IP addresses may be changed in response to attacks. An attack may be revealed, for example, by a regular series of messages indicating that a router is being probed in some way. Upon detection of an attack, the TARP layer process may respond to this event by changing its IP address. In addition, it may create a



New Bay Capital, LLC
Ex.1015-Page 2402 of 3151

subprocess that maintains the original IP address and continues interacting with the attacker in some manner.

Decoy packets may be generated by each TARP terminal on some basis determined by an algorithm. For example, the algorithm may be a random one which calls for the generation of a packet on a random basis when the terminal is idle. Alternatively, the algorithm may be responsive to time of day or detection of low traffic to generate more decoy packets during low traffic times. Note that packets are preferably generated in groups, rather than one by one, the groups being sized to simulate real messages. In addition, so that decoy packets may be inserted in normal TARP message streams, the background loop may have a latch that makes it more likely to insert decoy packets when a message stream is being received. Alternatively, if a large number of decoy packets is received along with regular TARP packets, the algorithm may increase the rate of dropping of decoy packets rather than forwarding them. The result of dropping and generating decoy packets in this way is to make the apparent incoming message size different from the apparent outgoing message size to help foil traffic analysis.

In various other embodiments of the invention, a scalable version of the system may be constructed in which a plurality of IP addresses are preassigned to each pair of communicating nodes in the network. Each pair of nodes agrees upon an algorithm for "hopping" between IP addresses (both sending and receiving), such that an eavesdropper sees apparently continuously random IP address pairs (source and destination) for packets transmitted between the pair. Overlapping or "reusable" IP addresses may be allocated to different users on the same subnet, since each node merely verifies that a particular packet includes a valid source/destination pair from the agreed-upon algorithm. Source/destination pairs are preferably not reused between any two nodes during any given end-to-end session, though limited IP block sizes or lengthy sessions might require it.

Further improvements described in this continuation-in-part application include: (1) a load balancer that distributes packets across different transmission paths according to transmission path quality; (2) a DNS proxy server that transparently



New Bay Capital, LLC
Ex.1015-Page 2403 of 3151

creates a virtual private network in response to a domain name inquiry; (3) a large-to-small link bandwidth management feature that prevents denial-of-service attacks at system chokepoints; (4) a traffic limiter that regulates incoming packets by limiting the rate at which a transmitter can be synchronized with a receiver; and (5) a signaling synchronizer that allows a large number of nodes to communicate with a central node by partitioning the communication function between two separate entities

**BRIEF DESCRIPTION OF THE DRAWINGS**

FIG. 1 is an illustration of secure communications over the Internet according to a prior art embodiment.

FIG. 2 is an illustration of secure communications over the Internet according to a an embodiment of the invention.

FIG. 3a is an illustration of a process of forming a tunneled IP packet according to an embodiment of the invention.

FIG. 3b is an illustration of a process of forming a tunneled IP packet according to another embodiment of the invention.

FIG. 4 is an illustration of an OSI layer location of processes that may be used to implement the invention.

FIG. 5 is a flow chart illustrating a process for routing a tunneled packet according to an embodiment of the invention.

FIG. 6 is a flow chart illustrating a process for forming a tunneled packet according to an embodiment of the invention.

FIG. 7 is a flow chart illustrating a process for receiving a tunneled packet according to an embodiment of the invention.

FIG. 8 shows how a secure session is established and synchronized between a client and a TARP router.

FIG. 9 shows an IP address hopping scheme between a client computer and TARP router using transmit and receive tables in each computer.

FIG. 10 shows physical link redundancy among three Internet Service Providers (ISPs) and a client computer.



New Bay Capital, LLC
Ex.1015-Page 2404 of 3151

FIG. 11 shows how multiple IP packets can be embedded into a single "frame" such as an Ethernet frame, and further shows the use of a discriminator field to camouflage true packet recipients.

FIG. 12A shows a system that employs hopped hardware addresses, hopped IP addresses, and hopped discriminator fields.

FIG. 12B shows several different approaches for hopping hardware addresses, IP addresses, and discriminator fields in combination.

FIG. 13 shows a technique for automatically re-establishing synchronization between sender and receiver through the use of a partially public sync value.

FIG. 14 shows a "checkpoint" scheme for regaining synchronization between a sender and recipient.

FIG. 15 shows further details of the checkpoint scheme of FIG. 14.

FIG. 16 shows how two addresses can be decomposed into a plurality of segments for comparison with presence vectors.

FIG. 17 shows a storage array for a receiver's active addresses.

FIG. 18 shows the receiver's storage array after receiving a sync request.

FIG. 19 shows the receiver's storage array after new addresses have been generated.

FIG. 20 shows a system employing distributed transmission paths.

FIG. 21 shows a plurality of link transmission tables that can be used to route packets in the system of FIG. 20.

FIG. 22A shows a flowchart for adjusting weight value distributions associated with a plurality of transmission links.

FIG. 22B shows a flowchart for setting a weight value to zero if a transmitter turns off.

FIG. 23 shows a system employing distributed transmission paths with adjusted weight value distributions for each path.

FIG. 24 shows an example using the system of FIG. 23.

FIG. 25 shows a conventional domain-name look-up service.



New Bay Capital, LLC
Ex.1015-Page 2405 of 3151

FIG. 26 shows a system employing a DNS proxy server with transparent VPN creation.

FIG. 27 shows steps that can be carried out to implement transparent VPN creation based on a DNS look-up function.

FIG. 28 shows a system including a link guard function that prevents packet overloading on a low-bandwidth link LOW BW.

FIG. 29 shows one embodiment of a system employing the principles of FIG. 28.

FIG. 30 shows a system that regulates packet transmission rates by throttling the rate at which synchronizations are performed.

FIG. 31 shows a signaling server 3101 and a transport server 3102 used to establish a VPN with a client computer.

FIG. 32 shows message flows relating to synchronization protocols of FIG. 31.

**DETAILED DESCRIPTION OF THE INVENTION**

Referring to FIG. 2, a secure mechanism for communicating over the internet employs a number of special routers or servers, called TARP routers 122-127 that are similar to regular IP routers 128-132 in that each has one or more IP addresses and uses normal IP protocol to send normal-looking IP packet messages, called TARP packets 140. TARP packets 140 are identical to normal IP packet messages that are routed by regular IP routers 128-132 because each TARP packet 140 contains a destination address as in a normal IP packet. However, instead of indicating a final destination in the destination field of the IP header, the TARP packet's 140 IP header always points to a next-hop in a series of TARP router hops, or the final destination, TARP terminal 110. Because the header of the TARP packet contains only the next-hop destination, there is no overt indication from an intercepted TARP packet of the true destination of the TARP packet 140 since the destination could always be the next-hop TARP router as well as the final destination, TARP terminal 110.

Each TARP packet's true destination is concealed behind an outer layer of encryption generated using a link key 146. The link key 146 is the encryption key used for encrypted communication between the end points (TARP terminals or TARP



New Bay Capital, LLC
Ex.1015-Page 2406 of 3151

routers) of a single link in the chain of hops connecting the originating TARP terminal 100 and the destination TARP terminal 110. Each TARP router 122-127, using the link key 146 it uses to communicate with the previous hop in a chain, can use the link key to reveal the true destination of a TARP packet. To identify the link key needed to decrypt the outer layer of encryption of a TARP packet, a receiving TARP or routing terminal may identify the transmitting terminal (which may indicate the link key used) by the sender field of the clear IP header. Alternatively, this identity may be hidden behind another layer of encryption in available bits in the clear IP header. Each TARP router, upon receiving a TARP message, determines if the message is a TARP message by using authentication data in the TARP packet. This could be recorded in available bytes in the TARP packet's IP header. Alternatively, TARP packets could be authenticated by attempting to decrypt using the link key 146 and determining if the results are as expected. The former may have computational advantages because it does not involve a decryption process.

Once the outer layer of decryption is completed by a TARP router 122-127, the TARP router determines the final destination. The system is preferably designed to cause each TARP packet 140 to undergo a minimum number of hops to help foil traffic analysis. The time to live counter in the IP header of the TARP message may be used to indicate a number of TARP router hops yet to be completed. Each TARP router then would decrement the counter and determine from that whether it should forward the TARP packet 140 to another TARP router 122-127 or to the destination TARP terminal 110. If the time to live counter is zero or below zero after decrementing, for an example of usage, the TARP router receiving the TARP packet 140 may forward the TARP packet 140 to the destination TARP terminal 110. If the time to live counter is above zero after decrementing, for an example of usage, the TARP router receiving the TARP packet 140 may forward the TARP packet 140 to a TARP router 122-127 that the current TARP terminal chooses at random. As a result, each TARP packet 140 is routed through some minimum number of hops of TARP routers 122-127 which are chosen at random.



New Bay Capital, LLC
Ex.1015-Page 2407 of 3151

Thus, each TARP packet, irrespective of the traditional factors determining traffic in the Internet, makes random trips among a number of geographically disparate routers before reaching its destination and each trip is highly likely to be different for each packet composing a given message because each trip is independently randomly determined as described above. This feature is called *agile routing*. For reasons that will become clear shortly, the fact that different packets take different routes provides distinct advantages by making it difficult for an interloper to obtain all the packets forming an entire multi-packet message. Agile routing is combined with another feature that furthers this purpose, a feature that ensures that any message is broken into multiple packets.

A TARP router receives a TARP packet when an IP address used by the TARP router coincides with the IP address in the TARP packet's IP header $IP_C$. The IP address of a TARP router, however, may not remain constant. To avoid and manage attacks, each TARP router, independently or under direction from another TARP terminal or router, may change its IP address. A separate, unchangeable identifier or address is also defined. This address, called the TARP address, is known only to TARP routers and terminals and may be correlated at any time by a TARP router or a TARP terminal using a Lookup Table (LUT). When a TARP router or terminal changes its IP address, it updates the other TARP routers and terminals which in turn update their respective LUTs. In reality, whenever a TARP router looks up the address of a destination in the encrypted header, it must convert a TARP address to a real IP address using its LUT.

While every TARP router receiving a TARP packet has the ability to determine the packet's final destination, the message payload is embedded behind an inner layer of encryption in the TARP packet that can only be unlocked using a session key. The session key is not available to any of the TARP routers 122-127 intervening between the originating 100 and destination 110 TARP terminals. The session key is used to decrypt the payloads of the TARP packets 140 permitting an entire message to be reconstructed.



New Bay Capital, LLC
Ex.1015-Page 2408 of 3151

In one embodiment, communication may be made private using link and session keys, which in turn may be shared and used according any desired method. For example, a public key or symmetric keys may be communicated between link or session endpoints using a public key method. Any of a variety of other mechanisms for securing data to ensure that only authorized computers can have access to the private information in the TARP packets 140 may be used as desired.

Referring to FIG. 3a, to construct a series of TARP packets, a data stream 300 of IP packets 207a, 207b, 207c, etc., such series of packets being formed by a network (IP) layer process, is broken into a series of small sized segments. In the present example, equal-sized segments 1-9 are defined and used to construct a set of interleaved data packets A, B, and C. Here it is assumed that the number of interleaved packets A, B, and C formed is three and that the number of IP packets 207a-207c used to form the three interleaved packets A, B, and C is exactly three. Of course, the number of IP packets spread over a group of interleaved packets may be any convenient number as may be the number of interleaved packets over which the incoming data stream is spread. The latter, the number of interleaved packets over which the data stream is spread, is called the *interleave window*.

To create a packet, the transmitting software interleaves the normal IP packets 207a *et. seq.* to form a new set of interleaved payload data 320. This payload data 320 is then encrypted using a session key to form a set of session-key-encrypted payload data 330, each of which, A, B, and C, will form the payload of a TARP packet. Using the IP header data, from the original packets 207a-207c, new TARP headers $IP_T$ are formed. The TARP headers $IP_T$ can be identical to normal IP headers or customized in some way. In a preferred embodiment, the TARP headers $IP_T$ are IP headers with added data providing the following information required for routing and reconstruction of messages, some of which data is ordinarily, or capable of being, contained in normal IP headers:

1. A window sequence number – an identifier that indicates where the packet belongs in the original message sequence.



New Bay Capital, LLC
Ex.1015-Page 2409 of 3151

2. An interleave sequence number – an identifier that indicates the interleaving sequence used to form the packet so that the packet can be deinterleaved along with other packets in the interleave window.

3. A time-to-live (TTL) datum – indicates the number of TARP-router-hops to be executed before the packet reaches its destination. Note that the TTL parameter may provide a datum to be used in a probabilistic formula for determining whether to route the packet to the destination or to another hop.

4. Data type identifier – indicates whether the payload contains, for example, TCP or UDP data.

5. Sender's address – indicates the sender's address in the TARP network.

6. Destination address – indicates the destination terminal's address in the TARP network.

7. Decoy/Real – an indicator of whether the packet contains real message data or dummy decoy data or a combination.

Obviously, the packets going into a single interleave window must include only packets with a common destination. Thus, it is assumed in the depicted example that the IP headers of IP packets 207a-207c all contain the same destination address or at least will be received by the same terminal so that they can be deinterleaved. Note that dummy or decoy data or packets can be added to form a larger interleave window than would otherwise be required by the size of a given message. Decoy or dummy data can be added to a stream to help foil traffic analysis by leveling the load on the network. Thus, it may be desirable to provide the TARP process with an ability to respond to the time of day or other criteria to generate more decoy data during low traffic periods so that communication bursts at one point in the Internet cannot be tied to communication bursts at another point to reveal the communicating endpoints.

Dummy data also helps to break the data into a larger number of inconspicuously-sized packets permitting the interleave window size to be increased while maintaining a reasonable size for each packet. (The packet size can be a single



New Bay Capital, LLC
Ex.1015-Page 2410 of 3151

standard size or selected from a fixed range of sizes.) One primary reason for desiring for each message to be broken into multiple packets is apparent if a chain block encryption scheme is used to form the first encryption layer prior to interleaving. A single block encryption may be applied to a portion, or the entirety, of a message, and that portion or entirety then interleaved into a number of separate packets.

Referring to FIG. 3b, in an alternative mode of TARP packet construction, a series of IP packets are accumulated to make up a predefined interleave window. The payloads of the packets are used to construct a single block 520 for chain block encryption using the session key. The payloads used to form the block are presumed to be destined for the same terminal. The block size may coincide with the interleave window as depicted in the example embodiment of FIG. 3b. After encryption, the encrypted block is broken into separate payloads and segments which are interleaved as in the embodiment of Fig 3a. The resulting interleaved packets A, B, and C, are then packaged as TARP packets with TARP headers as in the Example of FIG. 3a. The remaining process is as shown in, and discussed with reference to, FIG. 3a.

Once the TARP packets 340 are formed, each entire TARP packet 340, including the TARP header $IP_T$, is encrypted using the link key for communication with the first-hop-TARP router. The first hop TARP router is randomly chosen. A final unencrypted IP header $IP_C$ is added to each encrypted TARP packet 340 to form a normal IP packet 360 that can be transmitted to a TARP router. Note that the process of constructing the TARP packet 360 does not have to be done in stages as described. The above description is just a useful heuristic for describing the final product, namely, the TARP packet.

Note that, TARP header $IP_T$ could be a completely custom header configuration with no similarity to a normal IP header except that it contain the information identified above. This is so since this header is interpreted by only TARP routers.

The above scheme may be implemented entirely by processes operating between the data link layer and the network layer of each server or terminal participating in the TARP system. Referring to FIG. 4, a TARP transceiver 405 can be



New Bay Capital, LLC
Ex.1015-Page 2411 of 3151

an originating terminal 100, a destination terminal 110, or a TARP router 122-127. In each TARP Transceiver 405, a transmitting process is generated to receive normal packets from the Network (IP) layer and generate TARP packets for communication over the network. A receiving process is generated to receive normal IP packets containing TARP packets and generate from these normal IP packets which are "passed up" to the Network (IP) layer. Note that where the TARP Transceiver 405 is a router, the received TARP packets 140 are not processed into a stream of IP packets 415 because they need only be authenticated as proper TARP packets and then passed to another TARP router or a TARP destination terminal 110. The intervening process, a "TARP Layer" 420, could be combined with either the data link layer 430 or the Network layer 410. In either case, it would intervene between the data link layer 430 so that the process would receive regular IP packets containing embedded TARP packets and "hand up" a series of reassembled IP packets to the Network layer 410. As an example of combining the TARP layer 420 with the data link layer 430, a program may augment the normal processes running a communications card, for example, an Ethernet card. Alternatively, the TARP layer processes may form part of a dynamically loadable module that is loaded and executed to support communications between the network and data link layers.

Because the encryption system described above can be inserted between the data link and network layers, the processes involved in supporting the encrypted communication may be completely transparent to processes at the IP (network) layer and above. The TARP processes may also be completely transparent to the data link layer processes as well. Thus, no operations at or above the network layer, or at or below the data link layer, are affected by the insertion of the TARP stack. This provides additional security to all processes at or above the network layer, since the difficulty of unauthorized penetration of the network layer (by, for example, a hacker) is increased substantially. Even newly developed servers running at the session layer leave all processes below the session layer vulnerable to attack. Note that in this architecture, security is distributed. That is, notebook computers used by executives



New Bay Capital, LLC
Ex.1015-Page 2412 of 3151

on the road, for example, can communicate over the Internet without any compromise in security.

Note that IP address changes made by TARP terminals and routers can be done at regular intervals, at random intervals, or upon detection of "attacks." The variation of IP addresses hinders traffic analysis that might reveal which computers are communicating, and also provides a degree of immunity from attack. The level of immunity from attack is roughly proportional to the rate at which the IP address of the host is changing.

As mentioned, IP addresses may be changed in response to attacks. An attack may be revealed, for example, by a regular series of messages indicates that a router is being probed in some way. Upon detection of an attack, the TARP layer process may respond to this event by changing its IP address. To accomplish this, the TARP process will construct a TARP-formatted message, in the style of Internet Control Message Protocol (ICMP) datagrams as an example; this message will contain the machine's TARP address, its previous IP address, and its new IP address. The TARP layer will transmit this packet to at least one known TARP router; then upon receipt and validation of the message, the TARP router will update its LUT with the new IP address for the stated TARP address. The TARP router will then format a similar message, and broadcast it to the other TARP routers so that they may update their LUTs. Since the total number of TARP routers on any given subnet is expected to be relatively small, this process of updating the LUTs should be relatively fast. It may not, however, work as well when there is a relatively large number of TARP routers and/or a relatively large number of clients; this has motivated a refinement of this architecture to provide scalability; this refinement has led to a second embodiment, which is discussed below.

Upon detection of an attack, the TARP process may also create a subprocess that maintains the original IP address and continues interacting with the attacker. The latter may provide an opportunity to trace the attacker or study the attacker's methods (called "fishbowling" drawing upon the analogy of a small fish in a fish bowl that "thinks" it is in the ocean but is actually under captive observation). A history of the



New Bay Capital, LLC
Ex.1015-Page 2413 of 3151

communication between the attacker and the abandoned (fishbowled) IP address can be recorded or transmitted for human analysis or further synthesized for purposes of responding in some way.

As mentioned above, decoy or dummy data or packets can be added to outgoing data streams by TARP terminals or routers. In addition to making it convenient to spread data over a larger number of separate packets, such decoy packets can also help to level the load on inactive portions of the Internet to help foil traffic analysis efforts.

Decoy packets may be generated by each TARP terminal 100, 110 or each router 122-127 on some basis determined by an algorithm. For example, the algorithm may be a random one which calls for the generation of a packet on a random basis when the terminal is idle. Alternatively, the algorithm may be responsive to time of day or detection of low traffic to generate more decoy packets during low traffic times. Note that packets are preferably generated in groups, rather than one by one, the groups being sized to simulate real messages. In addition, so that decoy packets may be inserted in normal TARP message streams, the background loop may have a latch that makes it more likely to insert decoy packets when a message stream is being received. That is, when a series of messages are received, the decoy packet generation rate may be increased. Alternatively, if a large number of decoy packets is received along with regular TARP packets, the algorithm may increase the rate of dropping of decoy packets rather than forwarding them. The result of dropping and generating decoy packets in this way is to make the apparent incoming message size different from the apparent outgoing message size to help foil traffic analysis. The rate of reception of packets, decoy or otherwise, may be indicated to the decoy packet dropping and generating processes through perishable decoy and regular packet counters. (A perishable counter is one that resets or decrements its value in response to time so that it contains a high value when it is incremented in rapid succession and a small value when incremented either slowly or a small number of times in rapid succession.) Note that destination TARP terminal 110 may generate decoy packets



New Bay Capital, LLC
Ex.1015-Page 2414 of 3151

equal in number and size to those TARP packets received to make it appear it is merely routing packets and is therefore not the destination terminal.

Referring to FIG. 5, the following particular steps may be employed in the above-described method for routing TARP packets.

- S0. A background loop operation is performed which applies an algorithm which determines the generation of decoy IP packets. The loop is interrupted when an encrypted TARP packet is received.
- S2. The TARP packet may be probed in some way to authenticate the packet before attempting to decrypt it using the link key. That is, the router may determine that the packet is an authentic TARP packet by performing a selected operation on some data included with the clear IP header attached to the encrypted TARP packet contained in the payload. This makes it possible to avoid performing decryption on packets that are not authentic TARP packets.
- S3. The TARP packet is decrypted to expose the destination TARP address and an indication of whether the packet is a decoy packet or part of a real message.
- S4. If the packet is a decoy packet, the perishable decoy counter is incremented.
- S5. Based on the decoy generation/dropping algorithm and the perishable decoy counter value, if the packet is a decoy packet, the router may choose to throw it away. If the received packet is a decoy packet and it is determined that it should be thrown away (S6), control returns to step S0.
- S7. The TTL parameter of the TARP header is decremented and it is determined if the TTL parameter is greater than zero.
- S8. If the TTL parameter is greater than zero, a TARP address is randomly chosen from a list of TARP addresses maintained by the router and the link key and IP address corresponding to that TARP address memorized for use in creating a new IP packet containing the TARP packet.
- S9. If the TTL parameter is zero or less, the link key and IP address corresponding to the TARP address of the destination are memorized for use in creating the new IP packet containing the TARP packet.



New Bay Capital, LLC
Ex.1015-Page 2415 of 3151

- S10. The TARP packet is encrypted using the memorized link key.
- S11. An IP header is added to the packet that contains the stored IP address, the encrypted TARP packet wrapped with an IP header. and the completed packet transmitted to the next hop or destination.

Referring to FIG. 6, the following particular steps may be employed in the above-described method for generating TARP packets.

- S20. A background loop operation applies an algorithm that determines the generation of decoy IP packets. The loop is interrupted when a data stream containing IP packets is received for transmission.
- S21. The received IP packets are grouped into a set consisting of messages with a constant IP destination address. The set is further broken down to coincide with a maximum size of an interleave window The set is encrypted, and interleaved into a set of payloads destined to become TARP packets.
- S22. The TARP address corresponding to the IP address is determined from a lookup table and stored to generate the TARP header. An initial TTL count is generated and stored in the header. The TTL count may be random with minimum and maximum values or it may be fixed or determined by some other parameter.
- S23. The window sequence numbers and interleave sequence numbers are recorded in the TARP headers of each packet.
- S24. One TARP router address is randomly chosen for each TARP packet and the IP address corresponding to it stored for use in the clear IP header. The link key corresponding to this router is identified and used to encrypt TARP packets containing interleaved and encrypted data and TARP headers.
- S25. A clear IP header with the first hop router's real IP address is generated and added to each of the encrypted TARP packets and the resulting packets.

Referring to FIG. 7, the following particular steps may be employed in the above-described method for receiving TARP packets.



New Bay Capital, LLC
Ex.1015-Page 2416 of 3151

- S40. A background loop operation is performed which applies an algorithm which determines the generation of decoy IP packets. The loop is interrupted when an encrypted TARP packet is received.
- S42. The TARP packet may be probed to authenticate the packet before attempting to decrypt it using the link key.
- S43. The TARP packet is decrypted with the appropriate link key to expose the destination TARP address and an indication of whether the packet is a decoy packet or part of a real message.
- S44. If the packet is a decoy packet, the perishable decoy counter is incremented.
- S45. Based on the decoy generation/dropping algorithm and the perishable decoy counter value, if the packet is a decoy packet, the receiver may choose to throw it away.
- S46. The TARP packets are cached until all packets forming an interleave window are received.
- S47. Once all packets of an interleave window are received, the packets are deinterleaved.
- S48. The packets block of combined packets defining the interleave window is then decrypted using the session key.
- S49. The decrypted block is then divided using the window sequence data and the $IP_T$ headers are converted into normal $IP_C$ headers. The window sequence numbers are integrated in the $IP_C$ headers.
- S50. The packets are then handed up to the IP layer processes.

## I. SCALABILITY ENHANCEMENTS

The IP agility feature described above relies on the ability to transmit IP address changes to all TARP routers. The embodiments including this feature will be referred to as "boutique" embodiments due to potential limitations in scaling these features up for a large network, such as the Internet. (The "boutique" embodiments would, however, be robust for use in smaller networks, such as small virtual private networks, for example). One problem with the boutique embodiments is that if IP



New Bay Capital, LLC
Ex.1015-Page 2417 of 3151

address changes are to occur frequently, the message traffic required to update all routers sufficiently quickly creates a serious burden on the Internet when the TARP router and/or client population gets large. The bandwidth burden added to the networks, for example in ICMP packets, that would be used to update all the TARP routers could overwhelm the Internet for a large scale implementation that approached the scale of the Internet. In other words, the boutique system's scalability is limited.

A system can be constructed which trades some of the features of the above embodiments to provide the benefits of IP agility without the additional messaging burden. This is accomplished by IP address-hopping according to shared algorithms that govern IP addresses used between links participating in communications sessions between nodes such as TARP nodes. (Note that the IP hopping technique is also applicable to the boutique embodiment.) The IP agility feature discussed with respect to the boutique system can be modified so that it becomes decentralized under this scalable regime and governed by the above-described shared algorithm. Other features of the boutique system may be combined with this new type of IP-agility.

The new embodiment has the advantage of providing IP agility governed by a local algorithm and set of IP addresses exchanged by each communicating pair of nodes. This local governance is session-independent in that it may govern communications between a pair of nodes, irrespective of the session or end points being transferred between the directly communicating pair of nodes.

In the scalable embodiments, blocks of IP addresses are allocated to each node in the network. (This scalability will increase in the future, when Internet Protocol addresses are increased to 128-bit fields, vastly increasing the number of distinctly addressable nodes). Each node can thus use any of the IP addresses assigned to that node to communicate with other nodes in the network. Indeed, each pair of communicating nodes can use a plurality of source IP addresses and destination IP addresses for communicating with each other.

Each communicating pair of nodes in a chain participating in any session stores two blocks of IP addresses, called netblocks, and an algorithm and randomization seed for selecting, from each netblock, the next pair of



New Bay Capital, LLC
Ex.1015-Page 2418 of 3151

source/destination IP addresses that will be used to transmit the next message. In other words, the algorithm governs the sequential selection of IP-address pairs, one sender and one receiver IP address, from each netblock. The combination of algorithm, seed, and netblock (IP address block) will be called a "hopblock." A router issues separate transmit and receive hopblocks to its clients. The send address and the receive address of the IP header of each outgoing packet sent by the client are filled with the send and receive IP addresses generated by the algorithm. The algorithm is "clocked" (indexed) by a counter so that each time a pair is used, the algorithm turns out a new transmit pair for the next packet to be sent.

The router's receive hopblock is identical to the client's transmit hopblock. The router uses the receive hopblock to predict what the send and receive IP address pair for the next expected packet from that client will be. Since packets can be received out of order, it is not possible for the router to predict with certainty what IP address pair will be on the next sequential packet. To account for this problem, the router generates a range of predictions encompassing the number of possible transmitted packet send/receive addresses, of which the next packet received could leap ahead. Thus, if there is a vanishingly small probability that a given packet will arrive at the router ahead of 5 packets transmitted by the client before the given packet, then the router can generate a series of 6 send/receive IP address pairs (or "hop window") to compare with the next received packet. When a packet is received, it is marked in the hop window as such, so that a second packet with the same IP address pair will be discarded. If an out-of-sequence packet does not arrive within a predetermined timeout period, it can be requested for retransmission or simply discarded from the receive table, depending upon the protocol in use for that communications session, or possibly by convention.

When the router receives the client's packet, it compares the send and receive IP addresses of the packet with the next N predicted send and receive IP address pairs and rejects the packet if it is not a member of this set. Received packets that do not have the predicted source/destination IP addresses falling with the window are rejected, thus thwarting possible hackers. (With the number of possible combinations,



New Bay Capital, LLC
Ex.1015-Page 2419 of 3151

even a fairly large window would be hard to fall into at random.) If it is a member of this set, the router accepts the packet and processes it further. This link-based IP-hopping strategy, referred to as "IHOP," is a network element that stands on its own and is not necessarily accompanied by elements of the boutique system described above. If the routing agility feature described in connection with the boutique embodiment is combined with this link-based IP-hopping strategy, the router's next step would be to decrypt the TARP header to determine the destination TARP router for the packet and determine what should be the next hop for the packet. The TARP router would then forward the packet to a random TARP router or the destination TARP router with which the source TARP router has a link-based IP hopping communication established.

Figure 8 shows how a client computer 801 and a TARP router 811 can establish a secure session. When client 801 seeks to establish an IHOP session with TARP router 811, the client 801 sends "secure synchronization" request ("SSYN") packet 821 to the TARP router 811. This SYN packet 821 contains the client's 801 authentication token, and may be sent to the router 811 in an encrypted format. The source and destination IP numbers on the packet 821 are the client's 801 current fixed IP address, and a "known" fixed IP address for the router 811. (For security purposes, it may be desirable to reject any packets from outside of the local network that are destined for the router's known fixed IP address.) Upon receipt and validation of the client's 801 SSYN packet 821, the router 811 responds by sending an encrypted "secure synchronization acknowledgment" ("SSYN ACK") 822 to the client 801. This SSYN ACK 822 will contain the transmit and receive hopblocks that the client 801 will use when communicating with the TARP router 811. The client 801 will acknowledge the TARP router's 811 response packet 822 by generating an encrypted SSYN ACK ACK packet 823 which will be sent from the client's 801 fixed IP address and to the TARP router's 811 known fixed IP address. The client 801 will simultaneously generate a SSYN ACK ACK packet; this SSYN ACK packet, referred to as the Secure Session Initiation (SSI) packet 824, will be sent with the first {sender, receiver} IP pair in the client's transmit table 921 (FIG. 9), as specified in the



New Bay Capital, LLC
Ex.1015-Page 2420 of 3151

transmit hopblock provided by the TARP router 811 in the SSYN ACK packet 822. The TARP router 811 will respond to the SSI packet 824 with an SSI ACK packet 825, which will be sent with the first {sender, receiver} IP pair in the TARP router's transmit table 923. Once these packets have been successfully exchanged, the secure communications session is established, and all further secure communications between the client 801 and the TARP router 811 will be conducted via this secure session, as long as synchronization is maintained. If synchronization is lost, then the client 801 and TARP router 802 may re-establish the secure session by the procedure outlined in Figure 8 and described above.

While the secure session is active, both the client 901 and TARP router 911 (FIG. 9) will maintain their respective transmit tables 921, 923 and receive tables 922, 924, as provided by the TARP router during session synchronization 822. It is important that the sequence of IP pairs in the client's transmit table 921 be identical to those in the TARP router's receive table 924; similarly, the sequence of IP pairs in the client's receive table 922 must be identical to those in the router's transmit table 923. This is required for the session synchronization to be maintained. The client 901 need maintain only one transmit table 921 and one receive table 922 during the course of the secure session. Each sequential packet sent by the client 901 will employ the next {send, receive} IP address pair in the transmit table, regardless of TCP or UDP session. The TARP router 911 will expect each packet arriving from the client 901 to bear the next IP address pair shown in its receive table.

Since packets can arrive out of order, however, the router 911 can maintain a "look ahead" buffer in its receive table, and will mark previously-received IP pairs as invalid for future packets; any future packet containing an IP pair that is in the look-ahead buffer but is marked as previously received will be discarded. Communications from the TARP router 911 to the client 901 are maintained in an identical manner; in particular, the router 911 will select the next IP address pair from its transmit table 923 when constructing a packet to send to the client 901, and the client 901 will maintain a look-ahead buffer of expected IP pairs on packets that it is receiving. Each



New Bay Capital, LLC
Ex.1015-Page 2421 of 3151

TARP router will maintain separate pairs of transmit and receive tables for each client that is currently engaged in a secure session with or through that TARP router.

While clients receive their hopblocks from the first server linking them to the Internet, routers exchange hopblocks. When a router establishes a link-based IP-hopping communication regime with another router, each router of the pair exchanges its transmit hopblock. The transmit hopblock of each router becomes the receive hopblock of the other router. The communication between routers is governed as described by the example of a client sending a packet to the first router.

While the above strategy works fine in the IP milieu, many local networks that are connected to the Internet are Ethernet systems. In Ethernet, the IP addresses of the destination devices must be translated into hardware addresses, and vice versa, using known processes ("address resolution protocol," and "reverse address resolution protocol"). However, if the link-based IP-hopping strategy is employed, the correlation process would become explosive and burdensome. An alternative to the link-based IP hopping strategy may be employed within an Ethernet network. The solution is to provide that the node linking the Internet to the Ethernet (call it the border node) use the link-based IP-hopping communication regime to communicate with nodes outside the Ethernet LAN. Within the Ethernet LAN, each TARP node would have a single IP address which would be addressed in the conventional way. Instead of comparing the {sender, receiver} IP address pairs to authenticate a packet, the intra-LAN TARP node would use one of the IP header extension fields to do so. Thus, the border node uses an algorithm shared by the intra-LAN TARP node to generate a symbol that is stored in the free field in the IP header, and the intra-LAN TARP node generates a range of symbols based on its prediction of the next expected packet to be received from that particular source IP address. The packet is rejected if it does not fall into the set of predicted symbols (for example, numerical values) or is accepted if it does. Communications from the intra-LAN TARP node to the border node are accomplished in the same manner, though the algorithm will necessarily be different for security reasons. Thus, each of the communicating nodes will generate transmit and receive tables in a similar manner to that of Figure 9; the intra-LAN



New Bay Capital, LLC
Ex.1015-Page 2422 of 3151

TARP nodes transmit table will be identical to the border node's receive table, and the intra-LAN TARP node's receive table will be identical to the border node's transmit table.

The algorithm used for IP address-hopping can be any desired algorithm. For example, the algorithm can be a given pseudo-random number generator that generates numbers of the range covering the allowed IP addresses with a given seed. Alternatively, the session participants can assume a certain type of algorithm and specify simply a parameter for applying the algorithm. For example the assumed algorithm could be a particular pseudo-random number generator and the session participants could simply exchange seed values.

Note that there is no permanent physical distinction between the originating and destination terminal nodes. Either device at either end point can initiate a synchronization of the pair. Note also that the authentication/synchronization-request (and acknowledgment) and hopblock-exchange may all be served by a single message so that separate message exchanges may not be required.

As another extension to the stated architecture, multiple physical paths can be used by a client, in order to provide link redundancy and further thwart attempts at denial of service and traffic monitoring. As shown in Figure 10, for example, client 1001 can establish three simultaneous sessions with each of three TARP routers provided by different ISPs 1011, 1012, 1013. As an example, the client 1001 can use three different telephone lines 1021, 1022, 1023 to connect to the ISPs, or two telephone lines and a cable modem, etc. In this scheme, transmitted packets will be sent in a random fashion among the different physical paths. This architecture provides a high degree of communications redundancy, with improved immunity from denial-of-service attacks and traffic monitoring.

## 2. FURTHER EXTENSIONS

The following describes various extensions to the techniques, systems, and methods described above. As described above, the security of communications occurring between computers in a computer network (such as the Internet, an Ethernet, or others) can be enhanced by using seemingly random source and



New Bay Capital, LLC
Ex.1015-Page 2423 of 3151

destination Internet Protocol (IP) addresses for data packets transmitted over the network. This feature prevents eavesdroppers from determining which computers in the network are communicating with each other while permitting the two communicating computers to easily recognize whether a given received data packet is legitimate or not. In one embodiment of the above-described systems, an IP header extension field is used to authenticate incoming packets on an Ethernet.

Various extensions to the previously described techniques described herein include: (1) use of hopped hardware or "MAC" addresses in broadcast type network; (2) a self-synchronization technique that permits a computer to automatically regain synchronization with a sender; (3) synchronization algorithms that allow transmitting and receiving computers to quickly re-establish synchronization in the event of lost packets or other events; and (4) a fast-packet rejection mechanism for rejecting invalid packets. Any or all of these extensions can be combined with the features described above in any of various ways.

**A. Hardware Address Hopping**

Internet protocol-based communications techniques on a LAN—or across any dedicated physical medium—typically embed the IP packets within lower-level packets, often referred to as "frames." As shown in FIG. 11, for example, a first Ethernet frame 1150 comprises a frame header 1101 and two embedded IP packets IP1 and IP2, while a second Ethernet frame 1160 comprises a different frame header 1104 and a single IP packet IP3. Each frame header generally includes a source hardware address 1101A and a destination hardware address 1101B; other well-known fields in frame headers are omitted from FIG. 11 for clarity. Two hardware nodes communicating over a physical communication channel insert appropriate source and destination hardware addresses to indicate which nodes on the channel or network should receive the frame.

It may be possible for a nefarious listener to acquire information about the contents of a frame and/or its communicants by examining frames on a local network rather than (or in addition to) the IP packets themselves. This is especially true in broadcast media, such as Ethernet, where it is necessary to insert into the frame



New Bay Capital, LLC
Ex.1015-Page 2424 of 3151

header the hardware address of the machine that generated the frame and the hardware address of the machine to which frame is being sent. All nodes on the network can potentially "see" all packets transmitted across the network. This can be a problem for secure communications, especially in cases where the communicants do not want for any third party to be able to identify who is engaging in the information exchange. One way to address this problem is to push the address-hopping scheme down to the hardware layer. In accordance with various embodiments of the invention, hardware addresses are "hopped" in a manner similar to that used to change IP addresses, such that a listener cannot determine which hardware node generated a particular message nor which node is the intended recipient.

FIG. 12A shows a system in which Media Access Control ("MAC") hardware addresses are "hopped" in order to increase security over a network such as an Ethernet. While the description refers to the exemplary case of an Ethernet environment, the inventive principles are equally applicable to other types of communications media. In the Ethernet case, the MAC address of the sender and receiver are inserted into the Ethernet frame and can be observed by anyone on the LAN who is within the broadcast range for that frame. For secure communications, it becomes desirable to generate frames with MAC addresses that are not attributable to any specific sender or receiver.

As shown in FIG. 12A, two computer nodes 1201 and 1202 communicate over a communication channel such as an Ethernet. Each node executes one or more application programs 1203 and 1218 that communicate by transmitting packets through communication software 1204 and 1217, respectively. Examples of application programs include video conferencing, e-mail, word processing programs, telephony, and the like. Communication software 1204 and 1217 can comprise, for example, an OSI layered architecture or "stack" that standardizes various services provided at different levels of functionality.

The lowest levels of communication software 1204 and 1217 communicate with hardware components 1206 and 1214 respectively, each of which can include one or more registers 1207 and 1215 that allow the hardware to be reconfigured or



New Bay Capital, LLC
Ex.1015-Page 2425 of 3151

controlled in accordance with various communication protocols. The hardware components (an Ethernet network interface card, for example) communicate with each other over the communication medium. Each hardware component is typically pre-assigned a fixed hardware address or MAC number that identifies the hardware component to other nodes on the network. One or more interface drivers control the operation of each card and can, for example, be configured to accept or reject packets from certain hardware addresses. As will be described in more detail below, various embodiments of the inventive principles provide for "hopping" different addresses using one or more algorithms and one or more moving windows that track a range of valid addresses to validate received packets. Packets transmitted according to one or more of the inventive principles will be generally referred to as "secure" packets or "secure communications" to differentiate them from ordinary data packets that are transmitted in the clear using ordinary, machine-correlated addresses.

One straightforward method of generating non-attributable MAC addresses is an extension of the IP hopping scheme. In this scenario, two machines on the same LAN that desire to communicate in a secure fashion exchange random-number generators and seeds, and create sequences of quasi-random MAC addresses for synchronized hopping. The implementation and synchronization issues are then similar to that of IP hopping.

This approach, however, runs the risk of using MAC addresses that are currently active on the LAN—which, in turn, could interrupt communications for those machines. Since an Ethernet MAC address is at present 48 bits in length, the chance of randomly misusing an active MAC address is actually quite small. However, if that figure is multiplied by a large number of nodes (as would be found on an extensive LAN), by a large number of frames (as might be the case with packet voice or streaming video), and by a large number of concurrent Virtual Private Networks (VPNs), then the chance that a non-secure machine's MAC address could be used in an address-hopped frame can become non-trivial. In short, any scheme that runs even a small risk of interrupting communications for other machines on the LAN is bound to receive resistance from prospective system administrators. Nevertheless, it



New Bay Capital, LLC
Ex.1015-Page 2426 of 3151

is technically feasible, and can be implemented without risk on a LAN on which there is a small number of machines, or if all of the machines on the LAN are engaging in MAC-hopped communications.

Synchronized MAC address hopping may incur some overhead in the course of session establishment, especially if there are multiple sessions or multiple nodes involved in the communications. A simpler method of randomizing MAC addresses is to allow each node to receive and process *every* incident frame on the network. Typically, each network interface driver will check the destination MAC address in the header of every incident frame to see if it matches that machine's MAC address; if there is no match, then the frame is discarded. In one embodiment, however, these checks can be disabled, and every incident packet is passed to the TARP stack for processing. This will be referred to as "promiscuous" mode, since every incident frame is processed. Promiscuous mode allows the sender to use completely random, unsynchronized MAC addresses, since the destination machine is guaranteed to process the frame. The decision as to whether the packet was truly intended for that machine is handled by the TARP stack, which checks the source and destination IP addresses for a match in its IP synchronization tables. If no match is found, the packet is discarded; if there is a match, the packet is unwrapped, the inner header is evaluated, and if the inner header indicates that the packet is destined for that machine then the packet is forwarded to the IP stack—otherwise it is discarded.

One disadvantage of purely-random MAC address hopping is its impact on processing overhead; that is, since every incident frame must be processed, the machine's CPU is engaged considerably more often than if the network interface driver is discriminating and rejecting packets unilaterally. A compromise approach is to select either a single fixed MAC address or a small number of MAC addresses (e.g., one for each virtual private network on an Ethernet) to use for MAC-hopped communications, regardless of the actual recipient for which the message is intended. In this mode, the network interface driver can check each incident frame against one (or a few) pre-established MAC addresses, thereby freeing the CPU from the task of physical-layer packet discrimination. This scheme does not betray any useful



New Bay Capital, LLC
Ex.1015-Page 2427 of 3151

information to an interloper on the LAN; in particular, every secure packet can already be identified by a unique packet type in the outer header. However, since all machines engaged in secure communications would either be using the same MAC address, or be selecting from a small pool of predetermined MAC addresses, the association between a specific machine and a specific MAC address is effectively broken.

In this scheme, the CPU will be engaged more often than it would be in non-secure communications (or in synchronized MAC address hopping), since the network interface driver cannot always unilaterally discriminate between secure packets that are destined for that machine, and secure packets from other VPNs. However, the non-secure traffic is easily eliminated at the network interface, thereby reducing the amount of processing required of the CPU. There are boundary conditions where these statements would not hold, of course—e.g., if *all* of the traffic on the LAN is secure traffic, then the CPU would be engaged to the same degree as it is in the purely-random address hopping case; alternatively, if each VPN on the LAN uses a different MAC address, then the network interface can perfectly discriminate secure frames destined for the local machine from those constituting other VPNs. These are engineering tradeoffs that might be best handled by providing administrative options for the users when installing the software and/or establishing VPNs.

Even in this scenario, however, there still remains a slight risk of selecting MAC addresses that are being used by one or more nodes on the LAN. One solution to this problem is to formally assign one address or a range of addresses for use in MAC-hopped communications. This is typically done via an assigned numbers registration authority; e.g., in the case of Ethernet, MAC address ranges are assigned to vendors by the Institute of Electrical and Electronics Engineers (IEEE). A formally-assigned range of addresses would ensure that secure frames do not conflict with any properly-configured and properly-functioning machines on the LAN.

Reference will now be made to FIGS. 12A and 12B in order to describe the many combinations and features that follow the inventive principles. As explained



New Bay Capital, LLC
Ex.1015-Page 2428 of 3151

above, two computer nodes 1201 and 1202 are assumed to be communicating over a network or communication medium such as an Ethernet. A communication protocol in each node (1204 and 1217, respectively) contains a modified element 1205 and 1216 that performs certain functions that deviate from the standard communication protocols. In particular, computer node 1201 implements a first "hop" algorithm 1208X that selects seemingly random source and destination IP addresses (and, in one embodiment, seemingly random IP header discriminator fields) in order to transmit each packet to the other computer node. For example, node 1201 maintains a transmit table 1208 containing triplets of source (S), destination (D), and discriminator fields (DS) that are inserted into outgoing IP packet headers. The table is generated through the use of an appropriate algorithm (e.g., a random number generator that is seeded with an appropriate seed) that is known to the recipient node 1202. As each new IP packet is formed, the next sequential entry out of the sender's transmit table 1208 is used to populate the IP source, IP destination, and IP header extension field (e.g., discriminator field). It will be appreciated that the transmit table need not be created in advance but could instead be created on-the-fly by executing the algorithm when each packet is formed.

At the receiving node 1202, the same IP hop algorithm 1222X is maintained and used to generate a receive table 1222 that lists valid triplets of source IP address, destination IP address, and discriminator field. This is shown by virtue of the first five entries of transmit table 1208 matching the second five entries of receive table 1222. (The tables may be slightly offset at any particular time due to lost packets, misordered packets, or transmission delays). Additionally, node 1202 maintains a receive window W3 that represents a list of valid IP source, IP destination, and discriminator fields that will be accepted when received as part of an incoming IP packet. As packets are received, window W3 slides down the list of valid entries, such that the possible valid entries change over time. Two packets that arrive out of order but are nevertheless matched to entries within window W3 will be accepted; those falling outside of window W3 will be rejected as invalid. The length of window W3 can be adjusted as necessary to reflect network delays or other factors.



New Bay Capital, LLC
Ex.1015-Page 2429 of 3151

Node 1202 maintains a similar transmit table 1221 for creating IP packets and frames destined for node 1201 using a potentially different hopping algorithm 1221X, and node 1201 maintains a matching receive table 1209 using the same algorithm 1209X. As node 1202 transmits packets to node 1201 using seemingly random IP source, IP destination, and/or discriminator fields, node 1201 matches the incoming packet values to those falling within window W1 maintained in its receive table. In effect, transmit table 1208 of node 1201 is synchronized (i.e., entries are selected in the same order) to receive table 1222 of receiving node 1202. Similarly, transmit table 1221 of node 1202 is synchronized to receive table 1209 of node 1201. It will be appreciated that although a common algorithm is shown for the source, destination and discriminator fields in FIG. 12A (using, e.g., a different seed for each of the three fields), an entirely different algorithm could in fact be used to establish values for each of these fields. It will also be appreciated that one or two of the fields can be "hopped" rather than all three as illustrated.

In accordance with another aspect of the invention, hardware or "MAC" addresses are hopped instead of or in addition to IP addresses and/or the discriminator field in order to improve security in a local area or broadcast-type network. To that end, node 1201 further maintains a transmit table 1210 using a transmit algorithm 1210X to generate source and destination hardware addresses that are inserted into frame headers (e.g., fields 1101A and 1101B in FIG. 11) that are synchronized to a corresponding receive table 1224 at node 1202. Similarly, node 1202 maintains a different transmit table 1223 containing source and destination hardware addresses that is synchronized with a corresponding receive table 1211 at node 1201. In this manner, outgoing hardware frames appear to be originating from and going to completely random nodes on the network, even though each recipient can determine whether a given packet is intended for it or not. It will be appreciated that the hardware hopping feature can be implemented at a different level in the communications protocol than the IP hopping feature (e.g., in a card driver or in a hardware card itself to improve performance).



New Bay Capital, LLC
Ex.1015-Page 2430 of 3151

FIG. 12B shows three different embodiments or modes that can be employed using the aforementioned principles. In a first mode referred to as "promiscuous" mode, a common hardware address (e.g., a fixed address for source and another for destination) or else a completely random hardware address is used by all nodes on the network, such that a particular packet cannot be attributed to any one node. Each node must initially accept all packets containing the common (or random) hardware address and inspect the IP addresses or discriminator field to determine whether the packet is intended for that node. In this regard, either the IP addresses or the discriminator field or both can be varied in accordance with an algorithm as described above. As explained previously, this may increase each node's overhead since additional processing is involved to determine whether a given packet has valid source and destination hardware addresses.

In a second mode referred to as "promiscuous per VPN" mode, a small set of fixed hardware addresses are used, with a fixed source/destination hardware address used for all nodes communicating over a virtual private network. For example, if there are six nodes on an Ethernet, and the network is to be split up into two private virtual networks such that nodes on one VPN can communicate with only the other two nodes on its own VPN, then two sets of hardware addresses could be used: one set for the first VPN and a second set for the second VPN. This would reduce the amount of overhead involved in checking for valid frames since only packets arriving from the designated VPN would need to be checked. IP addresses and one or more discriminator fields could still be hopped as before for secure communication within the VPN. Of course, this solution compromises the anonymity of the VPNs (i.e., an outsider can easily tell what traffic belongs in which VPN, though he cannot correlate it to a specific machine/person). It also requires the use of a discriminator field to mitigate the vulnerability to certain types of DoS attacks. (For example, without the discriminator field, an attacker on the LAN could stream frames containing the MAC addresses being used by the VPN; rejecting those frames could lead to excessive processing overhead. The discriminator field would provide a low-overhead means of rejecting the false packets.)



New Bay Capital, LLC
Ex.1015-Page 2431 of 3151

In a third mode referred to as "hardware hopping" mode, hardware addresses are varied as illustrated in FIG. 12A, such that hardware source and destination addresses are changed constantly in order to provide non-attributable addressing. Variations on these embodiments are of course possible, and the invention is not intended to be limited in any respect by these illustrative examples.

**B. Extending the Address Space**

Address hopping provides security and privacy. However, the level of protection is limited by the number of addresses in the blocks being hopped. A hopblock denotes a field or fields modulated on a packet-wise basis for the purpose of providing a VPN. For instance, if two nodes communicate with IP address hopping using hopblocks of 4 addresses (2 bits) each, there would be 16 possible address-pair combinations. A window of size 16 would result in most address pairs being accepted as valid most of the time. This limitation can be overcome by using a discriminator field in addition to or instead of the hopped address fields. The discriminator field would be hopped in exactly the same fashion as the address fields and it would be used to determine whether a packet should be processed by a receiver.

Suppose that two clients, each using four-bit hopblocks, would like the same level of protection afforded to clients communicating via IP hopping between two A blocks (24 address bits eligible for hopping). A discriminator field of 20 bits, used in conjunction with the 4 address bits eligible for hopping in the IP address field, provides this level of protection. A 24-bit discriminator field would provide a similar level of protection if the address fields were not hopped or ignored. Using a discriminator field offers the following advantages: (1) an arbitrarily high level of protection can be provided, and (2) address hopping is unnecessary to provide protection. This may be important in environments where address hopping would cause routing problems.



New Bay Capital, LLC
Ex.1015-Page 2432 of 3151

## C. Synchronization Techniques

It is generally assumed that once a sending node and receiving node have exchanged algorithms and seeds (or similar information sufficient to generate quasi-random source and destination tables), subsequent communication between the two nodes will proceed smoothly. Realistically, however, two nodes may lose synchronization due to network delays or outages, or other problems. Consequently, it is desirable to provide means for re-establishing synchronization between nodes in a network that have lost synchronization.

One possible technique is to require that each node provide an acknowledgment upon successful receipt of each packet and, if no acknowledgment is received within a certain period of time, to re-send the unacknowledged packet. This approach, however, drives up overhead costs and may be prohibitive in high-throughput environments such as streaming video or audio, for example.

A different approach is to employ an automatic synchronizing technique that will be referred to herein as "self-synchronization." In this approach, synchronization information is embedded into each packet, thereby enabling the receiver to re-synchronize itself upon receipt of a single packet if it determines that is has lost synchronization with the sender. (If communications are already in progress, and the receiver determines that it is still in sync with the sender, then there is no need to re-synchronize.) A receiver could detect that it was out of synchronization by, for example, employing a "dead-man" timer that expires after a certain period of time, wherein the timer is reset with each valid packet. A time stamp could be hashed into the public sync field (see below) to preclude packet-retry attacks.

In one embodiment, a "sync field" is added to the header of each packet sent out by the sender. This sync field could appear in the clear or as part of an encrypted portion of the packet. Assuming that a sender and receiver have selected a random-number generator (RNG) and seed value, this combination of RNG and seed can be used to generate a random-number sequence (RNS). The RNS is then used to generate a sequence of source/destination IP pairs (and, if desired, discriminator fields and hardware source and destination addresses), as described above. It is not necessary,



New Bay Capital, LLC
Ex.1015-Page 2433 of 3151

however, to generate the entire sequence (or the first N-1 values) in order to generate the Nth random number in the sequence; if the sequence index N is known, the random value corresponding to that index can be directly generated (see below). Different RNGs (and seeds) with different fundamental periods could be used to generate the source and destination IP sequences, but the basic concepts would still apply. For the sake of simplicity, the following discussion will assume that IP source and destination address pairs (only) are hopped using a single RNG sequencing mechanism.

In accordance with a "self-synchronization" feature, a sync field in each packet header provides an index (i.e., a sequence number) into the RNS that is being used to generate IP pairs. Plugging this index into the RNG that is being used to generate the RNS yields a specific random number value, which in turn yields a specific IP pair. That is, an IP pair can be generated directly from knowledge of the RNG, seed, and index number; it is not necessary, in this scheme, to generate the entire sequence of random numbers that precede the sequence value associated with the index number provided.

Since the communicants have presumably previously exchanged RNGs and seeds, the only new information that must be provided in order to generate an IP pair is the sequence number. If this number is provided by the sender in the packet header, then the receiver need only plug this number into the RNG in order to generate an IP pair – and thus verify that the IP pair appearing in the header of the packet is valid. In this scheme, if the sender and receiver lose synchronization, the receiver can immediately re-synchronize upon receipt of a single packet by simply comparing the IP pair in the packet header to the IP pair generated from the index number. Thus, synchronized communications can be resumed upon receipt of a single packet, making this scheme ideal for multicast communications. Taken to the extreme, it could obviate the need for synchronization tables entirely; that is, the sender and receiver could simply rely on the index number in the sync field to validate the IP pair on each packet, and thereby eliminate the tables entirely.



New Bay Capital, LLC
Ex.1015-Page 2434 of 3151

The aforementioned scheme may have some inherent security issues associated with it — namely, the placement of the sync field. If the field is placed in the outer header, then an interloper could observe the values of the field and their relationship to the IP stream. This could potentially compromise the algorithm that is being used to generate the IP-address sequence, which would compromise the security of the communications. If, however, the value is placed in the inner header, then the sender must decrypt the inner header before it can extract the sync value and validate the IP pair; this opens up the receiver to certain types of denial-of-service (DoS) attacks, such as packet replay. That is, if the receiver must decrypt a packet before it can validate the IP pair, then it could potentially be forced to expend a significant amount of processing on decryption if an attacker simply retransmits previously valid packets. Other attack methodologies are possible in this scenario.

A possible compromise between algorithm security and processing speed is to split up the sync value between an inner (encrypted) and outer (unencrypted) header. That is, if the sync value is sufficiently long, it could potentially be split into a rapidly-changing part that can be viewed in the clear, and a fixed (or very slowly changing) part that must be protected. The part that can be viewed in the clear will be called the "public sync" portion and the part that must be protected will be called the "private sync" portion.

Both the public sync and private sync portions are needed to generate the complete sync value. The private portion, however, can be selected such that it is fixed or will change only occasionally. Thus, the private sync value can be stored by the recipient, thereby obviating the need to decrypt the header in order to retrieve it. If the sender and receiver have previously agreed upon the frequency with which the private part of the sync will change, then the receiver can selectively decrypt a single header in order to extract the new private sync if the communications gap that has led to lost synchronization has exceeded the lifetime of the previous private sync. This should not represent a burdensome amount of decryption, and thus should not open up the receiver to denial-of-service attack simply based on the need to occasionally decrypt a single header.



New Bay Capital, LLC
Ex.1015-Page 2435 of 3151

One implementation of this is to use a hashing function with a one-to-one mapping to generate the private and public sync portions from the sync value. This implementation is shown in FIG. 13, where (for example) a first ISP 1302 is the sender and a second ISP 1303 is the receiver. (Other alternatives are possible from FIG. 13.) A transmitted packet comprises a public or "outer" header 1305 that is not encrypted, and a private or "inner" header 1306 that is encrypted using for example a link key. Outer header 1305 includes a public sync portion while inner header 1306 contains the private sync portion. A receiving node decrypts the inner header using a decryption function 1307 in order to extract the private sync portion. This step is necessary only if the lifetime of the currently buffered private sync has expired. (If the currently-buffered private sync is still valid, then it is simply extracted from memory and "added" (which could be an inverse hash) to the public sync, as shown in step 1308.) The public and decrypted private sync portions are combined in function 1308 in order to generate the combined sync 1309. The combined sync (1309) is then fed into the RNG (1310) and compared to the IP address pair (1311) to validate or reject the packet.

An important consideration in this architecture is the concept of "future" and "past" where the public sync values are concerned. Though the sync values, themselves, should be random to prevent spoofing attacks, it may be important that the receiver be able to quickly identify a sync value that has already been sent — even if the packet containing that sync value was never actually received by the receiver. One solution is to hash a time stamp or sequence number into the public sync portion, which could be quickly extracted, checked, and discarded, thereby validating the public sync portion itself.

In one embodiment, packets can be checked by comparing the source/destination IP pair generated by the sync field with the pair appearing in the packet header. If (1) they match, (2) the time stamp is valid, and (3) the dead-man timer has expired, then re-synchronization occurs; otherwise, the packet is rejected. If enough processing power is available, the dead-man timer and synchronization tables



New Bay Capital, LLC
Ex.1015-Page 2436 of 3151

can be avoided altogether, and the receiver would simply resynchronize (e.g., validate) on every packet.

The foregoing scheme may require large-integer (e.g., 160-bit) math, which may affect its implementation. Without such large-integer registers, processing throughput would be affected, thus potentially affecting security from a denial-of-service standpoint. Nevertheless, as large-integer math processing features become more prevalent, the costs of implementing such a feature will be reduced.

**D. Other Synchronization Schemes**

As explained above, if W or more consecutive packets are lost between a transmitter and receiver in a VPN (where W is the window size), the receiver's window will not have been updated and the transmitter will be transmitting packets not in the receiver's window. The sender and receiver will not recover synchronization until perhaps the random pairs in the window are repeated by chance. Therefore, there is a need to keep a transmitter and receiver in synchronization whenever possible and to re-establish synchronization whenever it is lost.

A "checkpoint" scheme can be used to regain synchronization between a sender and a receiver that have fallen out of synchronization. In this scheme, a checkpoint message comprising a random IP address pair is used for communicating synchronization information. In one embodiment, two messages are used to communicate synchronization information between a sender and a recipient:

1. SYNC_REQ is a message used by the sender to indicate that it wants to synchronize; and
2. SYNC_ACK is a message used by the receiver to inform the transmitter that it has been synchronized.

According to one variation of this approach, both the transmitter and receiver maintain three checkpoints (see FIG. 14):

1. In the transmitter, ckpt_o ("checkpoint old") is the IP pair that was used to re-send the last SYNC_REQ packet to the receiver. In the receiver, ckpt_o ("checkpoint old") is the IP pair that receives repeated SYNC_REQ packets from the transmitter.



New Bay Capital, LLC
Ex.1015-Page 2437 of 3151

2. In the transmitter, ckpt_n ("checkpoint new") is the IP pair that will be used to send the next SYNC_REQ packet to the receiver. In the receiver, ckpt_n ("checkpoint new") is the IP pair that receives a new SYNC_REQ packet from the transmitter and which causes the receiver's window to be re-aligned, ckpt_o set to ckpt_n, a new ckpt_n to be generated and a new ckpt_r to be generated.

3. In the transmitter, ckpt_r is the IP pair that will be used to send the next SYNC_ACK packet to the receiver. In the receiver, ckpt_r is the IP pair that receives a new SYNC_ACK packet from the transmitter and which causes a new ckpt_n to be generated. Since SYNC_ACK is transmitted from the receiver ISP to the sender ISP, the transmitter ckpt_r refers to the ckpt_r of the receiver and the receiver ckpt_r refers to the ckpt_r of the transmitter (see FIG. 14).

When a transmitter initiates synchronization, the IP pair it will use to transmit the next data packet is set to a predetermined value and when a receiver first receives a SYNC_REQ, the receiver window is updated to be centered on the transmitter's next IP pair. This is the primary mechanism for checkpoint synchronization.

Synchronization can be initiated by a packet counter (e.g., after every N packets transmitted, initiate a synchronization) or by a timer (every S seconds, initiate a synchronization) or a combination of both. See FIG. 15. From the transmitter's perspective, this technique operates as follows: (1) Each transmitter periodically transmits a "sync request" message to the receiver to make sure that it is in sync. (2) If the receiver is still in sync, it sends back a "sync ack" message. (If this works, no further action is necessary). (3) If no "sync ack" has been received within a period of time, the transmitter retransmits the sync request again. If the transmitter reaches the next checkpoint without receiving a "sync ack" response, then synchronization is broken, and the transmitter should stop transmitting. The transmitter will continue to send sync_reqs until it receives a sync_ack , at which point transmission is reestablished.



New Bay Capital, LLC
Ex.1015-Page 2438 of 3151

From the receiver's perspective, the scheme operates as follows: (1) when it receives a "sync request" request from the transmitter, it advances its window to the next checkpoint position (even skipping pairs if necessary), and sends a "sync ack" message to the transmitter. If sync was never lost, then the "jump ahead" really just advances to the next available pair of addresses in the table (i.e., normal advancement).

If an interloper intercepts the "sync request" messages and tries to interfere with communication by sending new ones, it will be ignored if the synchronization has been established or it it will actually help to re-establish synchronization.

A window is realigned whenever a re-synchronization occurs. This realignment entails updating the receiver's window to straddle the address pairs used by the packet transmitted immediately after the transmission of the SYNC_REQ packet. Normally, the transmitter and receiver are in synchronization with one another. However, when network events occur, the receiver's window may have to be advanced by many steps during resynchronization. In this case, it is desirable to move the window ahead without having to step through the intervening random numbers sequentially. (This feature is also desirable for the auto-sync approach discussed above).

## E. Random Number Generator with a Jump-Ahead capability

An attractive method for generating randomly hopped addresses is to use identical random number generators in the transmitter and receiver and advance them as packets are transmitted and received. There are many random number generation algorithms that could be used. Each one has strengths and weaknesses for address hopping applications.

Linear congruential random number generators (LCRs) are fast, simple and well characterized random number generators that can be made to jump ahead $n$ steps efficiently. An LCR generates random numbers $X_1$, $X_2$, $X_3$ ... $X_k$ starting with seed $X_0$ using a recurrence

$$X_i=(a\,X_{i-1} + b) \bmod c, \qquad (1)$$

where a, b and c define a particular LCR. Another expression for $X_i$,



New Bay Capital, LLC
Ex.1015-Page 2439 of 3151

$X_i=((a^i(X_0+b)-b)/(a-1)) \bmod c$ (2)

enables the jump-ahead capability. The factor $a^i$ can grow very large even for modest i if left unfettered. Therefore some special properties of the modulo operation can be used to control the size and processing time required to compute (2). (2) can be rewritten as:

$X_i=(a^i (X_0(a-1)+b)-b)/(a-1) \bmod c.$ (3)

It can be shown that:

$(a^i(X_0(a-1)+b)-b)/(a-1) \bmod c =$

$((a^i \bmod((a-1)c)(X_0(a-1)+b) -b) /(a-1)) \bmod c$ (4).

$(X_0(a-1)+b)$ can be stored as $(X_0(a-1)+b) \bmod c$, b as b mod c and compute $a^i \bmod((a-1)c)$ (this requires O(log(i)) steps).

A practical implementation of this algorithm would jump a fixed distance, n, between synchronizations; this is tantamount to synchronizing every *n* packets. The window would commence *n* IP pairs from the start of the previous window. Using $X_j^w$, the random number at the $j^{th}$ checkpoint, as $X_0$ and *n* as *i*, a node can store $a^n \bmod((a-1)c)$ once per LCR and set

$X_{j+1}^w=X_{n(j+1)}=((a^n \bmod((a-1)c) (X_j^w (a-1)+b)-b)/(a-1))\bmod c,$ (5)

to generate the random number for the $j+1^{th}$ synchronization. Using this construction, a node could jump ahead an arbitrary (but fixed) distance between synchronizations in a constant amount of time (independent of n).

Pseudo-random number generators, in general, and LCRs, in particular, will eventually repeat their cycles. This repetition may present vulnerability in the IP hopping scheme. An adversary would simply have to wait for a repeat to predict future sequences. One way of coping with this vulnerability is to create a random number generator with a known long cycle. A random sequence can be replaced by a new random number generator before it repeats. LCRs can be constructed with known long cycles. This is not currently true of many random number generators.

Random number generators can be cryptographically insecure. An adversary can derive the RNG parameters by examining the output or part of the output. This is true of LCGs. This vulnerability can be mitigated by incorporating an encryptor,



New Bay Capital, LLC
Ex.1015-Page 2440 of 3151

designed to scramble the output as part of the random number generator. The random number generator prevents an adversary from mounting an attack—e.g., a known plaintext attack—against the encryptor.

## F. Random Number Generator Example

Consider a RNG where a=31,b=4 and c=15. For this case equation (1) becomes:

$X_i=(31\ X_{i-1} + 4) \bmod 15.$ (6)

If one sets $X_0=1$, equation (6) will produce the sequence 1, 5, 9, 13, 2, 6, 10, 14, 3, 7, 11, 0, 4, 8, 12. This sequence will repeat indefinitely. For a jump ahead of 3 numbers in this sequence $a^n= 31^3=29791$, $c^*(a-1)=15*30=450$ and $a^n \bmod((a-1)c) = 31^3 \bmod(15*30)=29791 \bmod(450)=91$. Equation (5) becomes:

$((91\ (X_i30+4)-4)/30) \bmod 15$ (7).

Table 1 shows the jump ahead calculations from (7) . The calculations start at 5 and jump ahead 3.

TABLE 1

| I | $X_i$ | $(X_i30+4)$ | $91\ (X_i30+4)-4$ | $((91\ (X_i30+4)-4)/30$ | $X_{i+3}$ |
|---|---|---|---|---|---|
| 1 | 5 | 154 | 14010 | 467 | 2 |
| 4 | 2 | 64 | 5820 | 194 | 14 |
| 7 | 14 | 424 | 38580 | 1286 | 11 |
| 10 | 11 | 334 | 30390 | 1013 | 8 |
| 13 | 8 | 244 | 22200 | 740 | 5 |

## G. Fast Packet Filter

Address hopping VPNs must rapidly determine whether a packet has a valid header and thus requires further processing, or has an invalid header (a hostile packet) and should be immediately rejected. Such rapid determinations will be referred to as "fast packet filtering." This capability protects the VPN from attacks by an adversary who streams hostile packets at the receiver at a high rate of speed in the hope of saturating the receiver's processor (a so-called "denial of service" attack). Fast packet



New Bay Capital, LLC
Ex.1015-Page 2441 of 3151

filtering is an important feature for implementing VPNs on shared media such as Ethernet.

Assuming that all participants in a VPN share an unassigned "A" block of addresses, one possibility is to use an experimental "A" block that will never be assigned to any machine that is not address hopping on the shared medium. "A" blocks have a 24 bits of address that can be hopped as opposed to the 8 bits in "C" blocks. In this case a hopblock will be the "A" block. The use of the experimental "A" block is a likely option on an Ethernet because:

1. The addresses have no validity outside of the Ethernet and will not be routed out to a valid outside destination by a gateway.
2. There are $2^{24}$ (~16 million) addresses that can be hopped within each "A" block. This yields >280 trillion possible address pairs making it very unlikely that an adversary would guess a valid address. It also provides acceptably low probability of collision between separate VPNs (all VPNs on a shared medium independently generate random address pairs from the same "A" block).
3. The packets will not be received by someone on the Ethernet who is not on a VPN (unless the machine is in promiscuous mode) minimizing impact on non-VPN computers.

The Ethernet example will be used to describe one implementation of fast packet filtering. The ideal algorithm would quickly examine a packet header, determine whether the packet is hostile, and reject any hostile packets or determine which active IP pair the packet header matches. The problem is a classical associative memory problem. A variety of techniques have been developed to solve this problem (hashing, B–trees etc). Each of these approaches has its strengths and weaknesses. For instance, hash tables can be made to operate quite fast in a statistical sense, but can occasionally degenerate into a much slower algorithm. This slowness can persist for a period of time. Since there is a need to discard hostile packets quickly at all times, hashing would be unacceptable.

## H. Presence Vector Algorithm



New Bay Capital, LLC
Ex.1015-Page 2442 of 3151

A presence vector is a bit vector of length $2^n$ that can be indexed by $n$-bit numbers (each ranging from 0 to $2^n-1$). One can indicate the presence of $k$ $n$-bit numbers (not necessarily unique), by setting the bits in the presence vector indexed by each number to 1. Otherwise, the bits in the presence vector are 0. An $n$-bit number, $x$, is one of the $k$ numbers if and only if the $x^{th}$ bit of the presence vector is 1. A fast packet filter can be implemented by indexing the presence vector and looking for a 1, which will be referred to as the "test."

For example, suppose one wanted to represent the number 135 using a presence vector. The $135^{th}$ bit of the vector would be set. Consequently, one could very quickly determine whether an address of 135 was valid by checking only one bit: the $135^{th}$ bit. The presence vectors could be created in advance corresponding to the table entries for the IP addresses. In effect, the incoming addresses can be used as indices into a long vector, making comparisons very fast. As each RNG generates a new address, the presence vector is updated to reflect the information. As the window moves, the presence vector is updated to zero out addresses that are no longer valid.

There is a trade-off between efficiency of the test and the amount of memory required for storing the presence vector(s). For instance, if one were to use the 48 bits of hopping addresses as an index, the presence vector would have to be 35 terabytes. Clearly, this is too large for practical purposes. Instead, the 48 bits can be divided into several smaller fields. For instance, one could subdivide the 48 bits into four 12-bit fields (see FIG. 16). This reduces the storage requirement to 2048 bytes at the expense of occasionally having to process a hostile packet. In effect, instead of one long presence vector, the decomposed address portions must match all four shorter presence vectors before further processing is allowed. (If the first part of the address portion doesn't match the first presence vector, there is no need to check the remaining three presence vectors).

A presence vector will have a 1 in the $y^{th}$ bit if and only if one or more addresses with a corresponding field of y are active. An address is active only if each presence vector indexed by the appropriate sub-field of the address is 1.



New Bay Capital, LLC
Ex.1015-Page 2443 of 3151

Consider a window of 32 active addresses and 3 checkpoints. A hostile packet will be rejected by the indexing of one presence vector more than 99% of the time. A hostile packet will be rejected by the indexing of all 4 presence vectors more than 99.9999995% of the time. On average, hostile packets will be rejected in less than 1.02 presence vector index operations.

The small percentage of hostile packets that pass the fast packet filter will be rejected when matching pairs are not found in the active window or are active checkpoints. Hostile packets that serendipitously match a header will be rejected when the VPN software attempts to decrypt the header. However, these cases will be extremely rare. There are many other ways this method can be configured to arbitrate the space/speed tradeoffs.

## I. Further Synchronization Enhancements

A slightly modified form of the synchronization techniques described above can be employed. The basic principles of the previously described checkpoint synchronization scheme remain unchanged. The actions resulting from the reception of the checkpoints are, however, slightly different. In this variation, the receiver will maintain between OoO ("Out of Order") and 2×WINDOW_SIZE+OoO active addresses ($1 \leq OoO \leq WINDOW\_SIZE$ and $WINDOW\_SIZE \geq 1$). OoO and WINDOW_SIZE are engineerable parameters, where OoO is the minimum number of addresses needed to accommodate lost packets due to events in the network or out of order arrivals and WINDOW_SIZE is the number of packets transmitted before a SYNC_REQ is issued. FIG. 17 depicts a storage array for a receiver's active addresses.

The receiver starts with the first 2×WINDOW_SIZE addresses loaded and active (ready to receive data). As packets are received, the corresponding entries are marked as "used" and are no longer eligible to receive packets. The transmitter maintains a packet counter, initially set to 0, containing the number of data packets transmitted since the last *initial* transmission of a SYNC_REQ for which SYNC_ACK has been received. When the transmitter packet counter equals WINDOW_SIZE, the transmitter generates a SYNC_REQ and does its initial



New Bay Capital, LLC
Ex.1015-Page 2444 of 3151

transmission. When the receiver receives a SYNC_REQ corresponding to its current CKPT_N, it generates the next WINDOW_SIZE addresses and starts loading them in order starting at the first location after the last active address wrapping around to the beginning of the array after the end of the array has been reached. The receiver's array might look like FIG. 18 when a SYNC_REQ has been received. In this case a couple of packets have been either lost or will be received out of order when the SYNC_REQ is received.

FIG. 19 shows the receiver's array after the new addresses have been generated. If the transmitter does not receive a SYNC_ACK, it will re-issue the SYNC_REQ at regular intervals. When the transmitter receives a SYNC_ACK, the packet counter is decremented by WINDOW_SIZE. If the packet counter reaches 2×WINDOW_SIZE – OoO then the transmitter ceases sending data packets until the appropriate SYNC_ACK is finally received. The transmitter then resumes sending data packets. Future behavior is essentially a repetition of this initial cycle. The advantages of this approach are:

1. There is no need for an efficient jump ahead in the random number generator,
2. No packet is ever transmitted that does not have a corresponding entry in the receiver side
3. No timer based re-synchronization is necessary. This is a consequence of 2.
4. The receiver will always have the ability to accept data messages transmitted within OoO messages of the most recently transmitted message.

### J. Distributed Transmission Path Variant

Another embodiment incorporating various inventive principles is shown in FIG. 20. In this embodiment, a message transmission system includes a first computer 2001 in communication with a second computer 2002 through a network 2011 of intermediary computers. In one variant of this embodiment, the network includes two edge routers 2003 and 2004 each of which is linked to a plurality of Internet Service Providers (ISPs) 2005 through 2010. Each ISP is coupled to a plurality of other ISPs in an arrangement as shown in FIG. 20, which is a representative configuration only and is not intended to be limiting. Each connection



New Bay Capital, LLC
Ex.1015-Page 2445 of 3151

between ISPs is labeled in FIG. 20 to indicate a specific physical transmission path (e.g., AD is a physical path that links ISP A (element 2005) to ISP D (element 2008)). Packets arriving at each edge router are selectively transmitted to one of the ISPs to which the router is attached on the basis of a randomly or quasi-randomly selected basis.

As shown in FIG. 21, computer 2001 or edge router 2003 incorporates a plurality of link transmission tables 2100 that identify, for each potential transmission path through the network, valid sets of IP addresses that can be used to transmit the packet. For example, AD table 2101 contains a plurality of IP source/destination pairs that are randomly or quasi-randomly generated. When a packet is to be transmitted from first computer 2001 to second computer 2002, one of the link tables is randomly (or quasi-randomly) selected, and the next valid source/destination address pair from that table is used to transmit the packet through the network. If path AD is randomly selected, for example, the next source/destination IP address pair (which is pre-determined to transmit between ISP A (element 2005) and ISP B (element 2008)) is used to transmit the packet. If one of the transmission paths becomes degraded or inoperative, that link table can be set to a "down" condition as shown in table 2105, thus preventing addresses from being selected from that table. Other transmission paths would be unaffected by this broken link.

## 3. CONTINUATION-IN-PART IMPROVEMENTS

The following describes various improvements and features that can be applied to the embodiments described above. The improvements include: (1) a load balancer that distributes packets across different transmission paths according to transmission path quality; (2) a DNS proxy server that transparently creates a virtual private network in response to a domain name inquiry; (3) a large-to-small link bandwidth management feature that prevents denial-of-service attacks at system chokepoints; (4) a traffic limiter that regulates incoming packets by limiting the rate at which a transmitter can be synchronized with a receiver; and (5) a signaling synchronizer that allows a large number of nodes to communicate with a central node



New Bay Capital, LLC
Ex.1015-Page 2446 of 3151

by partitioning the communication function between two separate entities. Each is discussed separately below.

### A. Load Balancer

Various embodiments described above include a system in which a transmitting node and a receiving node are coupled through a plurality of transmission paths, and wherein successive packets are distributed quasi-randomly over the plurality of paths. See, for example, FIGS. 20 and 21 and accompanying description. The improvement extends this basic concept to encompass distributing packets across different paths in such a manner that the loads on the paths are generally balanced according to transmission link quality.

In one embodiment, a system includes a transmitting node and a receiving node that are linked via a plurality of transmission paths having potentially varying transmission quality. Successive packets are transmitted over the paths based on a weight value distribution function for each path. The rate that packets will be transmitted over a given path can be different for each path. The relative "health" of each transmission path is monitored in order to identify paths that have become degraded. In one embodiment, the health of each path is monitored in the transmitter by comparing the number of packets transmitted to the number of packet acknowledgements received. Each transmission path may comprise a physically separate path (e.g., via dial-up phone line, computer network, router, bridge, or the like), or may comprise logically separate paths contained within a broadband communication medium (e.g., separate channels in an FDM, TDM, CDMA, or other type of modulated or unmodulated transmission link).

When the transmission quality of a path falls below a predetermined threshold and there are other paths that can transmit packets, the transmitter changes the weight value used for that path, making it less likely that a given packet will be transmitted over that path. The weight will preferably be set no lower than a minimum value that keeps nominal traffic on the path. The weights of the other available paths are altered to compensate for the change in the affected path. When the quality of a path degrades to where the transmitter is turned off by the synchronization function (i.e., no packets



New Bay Capital, LLC
Ex.1015-Page 2447 of 3151

are arriving at the destination), the weight is set to zero. If all transmitters are turned off, no packets are sent.

Conventional TCP/IP protocols include a "throttling" feature that reduces the transmission rate of packets when it is determined that delays or errors are occurring in transmission. In this respect, timers are sometimes used to determine whether packets have been received. These conventional techniques for limiting transmission of packets, however, do not involve multiple transmission paths between two nodes wherein transmission across a particular path relative to the others is changed based on link quality.

According to certain embodiments, in order to damp oscillations that might otherwise occur if weight distributions are changed drastically (e.g., according to a step function), a linear or an exponential decay formula can be applied to gradually decrease the weight value over time that a degrading path will be used. Similarly, if the health of a degraded path improves, the weight value for that path is gradually increased.

Transmission link health can be evaluated by comparing the number of packets that are acknowledged within the transmission window (see embodiments discussed above) to the number of packets transmitted within that window and by the state of the transmitter (i.e., on or off). In other words, rather than accumulating general transmission statistics over time for a path, one specific implementation uses the "windowing" concepts described above to evaluate transmission path health.

The same scheme can be used to shift virtual circuit paths from an "unhealthy" path to a "healthy" one, and to select a path for a new virtual circuit.

FIG. 22A shows a flowchart for adjusting weight values associated with a plurality of transmission links. It is assumed that software executing in one or more computer nodes executes the steps shown in FIG. 22A. It is also assumed that the software can be stored on a computer-readable medium such as a magnetic or optical disk for execution by a computer.

Beginning in step 2201, the transmission quality of a given transmission path is measured. As described above, this measurement can be based on a comparison



New Bay Capital, LLC
Ex.1015-Page 2448 of 3151

between the number of packets transmitted over a particular link to the number of packet acknowledgements received over the link (e.g., per unit time, or in absolute terms). Alternatively, the quality can be evaluated by comparing the number of packets that are acknowledged within the transmission window to the number of packets that were transmitted within that window. In yet another variation, the number of missed synchronization messages can be used to indicate link quality. Many other variations are of course possible.

In step 2202, a check is made to determine whether more than one transmitter (e.g., transmission path) is turned on. If not, the process is terminated and resumes at step 2201.

In step 2203, the link quality is compared to a given threshold (e.g., 50%, or any arbitrary number). If the quality falls below the threshold, then in step 2207 a check is made to determine whether the weight is above a minimum level (e.g., 1%). If not, then in step 2209 the weight is set to the minimum level and processing resumes at step 2201. If the weight is above the minimum level, then in step 2208 the weight is gradually decreased for the path, then in step 2206 the weights for the remaining paths are adjusted accordingly to compensate (e.g., they are increased).

If in step 2203 the quality of the path was greater than or equal to the threshold, then in step 2204 a check is made to determine whether the weight is less than a steady-state value for that path. If so, then in step 2205 the weight is increased toward the steady-state value, and in step 2206 the weights for the remaining paths are adjusted accordingly to compensate (e.g., they are decreased). If in step 2204 the weight is not less than the steady-state value, then processing resumes at step 2201 without adjusting the weights.

The weights can be adjusted incrementally according to various functions, preferably by changing the value gradually. In one embodiment, a linearly decreasing function is used to adjust the weights; according to another embodiment, an exponential decay function is used. Gradually changing the weights helps to damp oscillators that might otherwise occur if the probabilities were abruptly.



New Bay Capital, LLC
Ex.1015-Page 2449 of 3151

Although not explicitly shown in FIG. 22A the process can be performed only periodically (e.g., according to a time schedule), or it can be continuously run, such as in a background mode of operation. In one embodiment, the combined weights of all potential paths should add up to unity (e.g., when the weighting for one path is decreased, the corresponding weights that the other paths will be selected will increase).

Adjustments to weight values for other paths can be prorated. For example, a decrease of 10% in weight value for one path could result in an evenly distributed increase in the weights for the remaining paths. Alternatively, weightings could be adjusted according to a weighted formula as desired (e.g., favoring healthy paths over less healthy paths). In yet another variation, the difference in weight value can be amortized over the remaining links in a manner that is proportional to their traffic weighting.

FIG. 22B shows steps that can be executed to shut down transmission links where a transmitter turns off. In step 2210, a transmitter shut-down event occurs. In step 2211, a test is made to determine whether at least one transmitter is still turned on. If not, then in step 2215 all packets are dropped until a transmitter turns on. If in step 2211 at least one transmitter is turned on, then in step 2212 the weight for the path is set to zero, and the weights for the remaining paths are adjusted accordingly.

FIG. 23 shows a computer node 2301 employing various principles of the above-described embodiments. It is assumed that two computer nodes of the type shown in FIG. 23 communicate over a plurality of separate physical transmission paths. As shown in FIG. 23, four transmission paths X1 through X4 are defined for communicating between the two nodes. Each node includes a packet transmitter 2302 that operates in accordance with a transmit table 2308 as described above. (The packet transmitter could also operate without using the IP-hopping features described above, but the following description assumes that some form of hopping is employed in conjunction with the path selection mechanism.). The computer node also includes a packet receiver 2303 that operates in accordance with a receive table 2309, including a moving window W that moves as valid packets are received. Invalid



New Bay Capital, LLC
Ex.1015-Page 2450 of 3151

packets having source and destination addresses that do not fall within window W are rejected.

As each packet is readied for transmission, source and destination IP addresses (or other discriminator values) are selected from transmit table 2308 according to any of the various algorithms described above, and packets containing these source/destination address pairs, which correspond to the node to which the four transmission paths are linked, are generated to a transmission path switch 2307. Switch 2307, which can comprise a software function, selects from one of the available transmission paths according to a weight distribution table 2306. For example, if the weight for path X1 is 0.2, then every fifth packet will be transmitted on path X1. A similar regime holds true for the other paths as shown. Initially, each link's weight value can be set such that it is proportional to its bandwidth, which will be referred to as its "steady-state" value.

Packet receiver 2303 generates an output to a link quality measurement function 2304 that operates as described above to determine the quality of each transmission path. (The input to packet receiver 2303 for receiving incoming packets is omitted for clarity). Link quality measurement function 2304 compares the link quality to a threshold for each transmission link and, if necessary, generates an output to weight adjustment function 2305. If a weight adjustment is required, then the weights in table 2306 are adjusted accordingly, preferably according to a gradual (e.g., linearly or exponentially declining) function. In one embodiment, the weight values for all available paths are initially set to the same value, and only when paths degrade in quality are the weights changed to reflect differences.

Link quality measurement function 2304 can be made to operate as part of a synchronizer function as described above. That is, if resynchronization occurs and the receiver detects that synchronization has been lost (e.g., resulting in the synchronization window W being advanced out of sequence), that fact can be used to drive link quality measurement function 2304. According to one embodiment, load balancing is performed using information garnered during the normal synchronization, augmented slightly to communicate link health from the receiver to



New Bay Capital, LLC
Ex.1015-Page 2451 of 3151

the transmitter. The receiver maintains a count, MESS_R(W), of the messages received in synchronization window W. When it receives a synchronization request (SYNC_REQ) corresponding to the end of window W, the receiver includes counter MESS_R in the resulting synchronization acknowledgement (SYNC_ACK) sent back to the transmitter. This allows the transmitter to compare messages sent to messages received in order to asses the health of the link.

If synchronization is completely lost, weight adjustment function 2305 decreases the weight value on the affected path to zero. When synchronization is regained, the weight value for the affected path is gradually increased to its original value. Alternatively, link quality can be measured by evaluating the length of time required for the receiver to acknowledge a synchronization request. In one embodiment, separate transmit and receive tables are used for each transmission path.

When the transmitter receives a SYNC_ACK, the MESS_R is compared with the number of messages transmitted in a window (MESS_T). When the transmitter receives a SYNC_ACK, the traffic probabilities will be examined and adjusted if necessary. MESS_R is compared with the number of messages transmitted in a window (MESS_T). There are two possibilities:

1. If MESS_R is less than a threshold value, THRESH, then the link will be deemed to be unhealthy. If the transmitter was turned off, the transmitter is turned on and the weight P for that link will be set to a minimum value MIN. This will keep a trickle of traffic on the link for monitoring purposes until it recovers. If the transmitter was turned on, the weight P for that link will be set to:

$$P' = \alpha \times MIN + (1 - \alpha) \times P \quad (1)$$

Equation 1 will exponentially damp the traffic weight value to MIN during sustained periods of degraded service.

2. If MESS_R for a link is greater than or equal to THRESH, the link will be deemed healthy. If the weight P for that link is greater than or equal to the steady state value S for that link, then P is left unaltered. If the weight P for that link is less than THRESH then P will be set to:

$$P' = \beta \times S + (1 - \beta) \times P \quad (2)$$



New Bay Capital, LLC
Ex.1015-Page 2452 of 3151

where $\beta$ is a parameter such that $0<=\beta<=1$ that determines the damping rate of P.

Equation 2 will increase the traffic weight to S during sustained periods of acceptable service in a damped exponential fashion.

A detailed example will now be provided with reference to FIG. 24. As shown in FIG. 24, a first computer 2401 communicates with a second computer 2402 through two routers 2403 and 2404. Each router is coupled to the other router through three transmission links. As described above, these may be physically diverse links or logical links (including virtual private networks).

Suppose that a first link L1 can sustain a transmission bandwidth of 100 Mb/s and has a window size of 32; link L2 can sustain 75 Mb/s and has a window size of 24; and link L3 can sustain 25 Mb/s and has a window size of 8. The combined links can thus sustain 200Mb/s. The steady state traffic weights are 0.5 for link L1; 0.375 for link L2. and 0.125 for link L3. MIN=1Mb/s, THRESH =0.8 MESS_T for each link, $\alpha$=.75 and $\beta$=.5. These traffic weights will remain stable until a link stops for synchronization or reports a number of packets received less than its THRESH. Consider the following sequence of events:

1. Link L1 receives a SYNC_ACK containing a MESS_R of 24, indicating that only 75% of the MESS_T (32) messages transmitted in the last window were successfully received. Link 1 would be below THRESH (0.8). Consequently, link L1's traffic weight value would be reduced to 0.12825, while link L2's traffic weight value would be increased to 0.65812 and link L3's traffic weight value would be increased to 0.217938.

2. Link L2 and L3 remained healthy and link L1 stopped to synchronize. Then link L1's traffic weight value would be set to 0, link L2's traffic weight value would be set to 0.75, and link L33's traffic weight value would be set to 0.25.

3. Link L1 finally received a SYNC_ACK containing a MESS_R of 0 indicating that none of the MESS_T (32) messages transmitted in the last window were successfully received. Link L1 would be below THRESH. Link L1's traffic weight value would be increased to .005, link L2's traffic weight value would be



New Bay Capital, LLC
Ex.1015-Page 2453 of 3151

decreased to 0.74625, and link L3's traffic weight value would be decreased to 0.24875.

4. Link L1 received a SYNC_ACK containing a MESS_R of 32 indicating that 100% of the MESS_T (32) messages transmitted in the last window were successfully received. Link L1 would be above THRESH. Link L1's traffic weight value would be increased to 0.2525, while link L2's traffic weight value would be decreased to 0.560625 and link L3's traffic weight value would be decreased to .186875.

5. Link L1 received a SYNC_ACK containing a MESS_R of 32 indicating that 100% of the MESS_T (32) messages transmitted in the last window were successfully received. Link L1 would be above THRESH. Link L1's traffic weight value would be increased to 0.37625; link L2's traffic weight value would be decreased to 0.4678125, and link L3's traffic weight value would be decreased to 0.1559375.

6. Link L1 remains healthy and the traffic probabilities approach their steady state traffic probabilities.

**B. Use of a DNS Proxy to Transparently Create Virtual Private Networks**

A second improvement concerns the automatic creation of a virtual private network (VPN) in response to a domain-name server look-up function.

Conventional Domain Name Servers (DNSs) provide a look-up function that returns the IP address of a requested computer or host. For example, when a computer user types in the web name "Yahoo.com," the user's web browser transmits a request to a DNS, which converts the name into a four-part IP address that is returned to the user's browser and then used by the browser to contact the destination web site.

This conventional scheme is shown in FIG. 25. A user's computer 2501 includes a client application 2504 (for example, a web browser) and an IP protocol stack 2505. When the user enters the name of a destination host, a request DNS REQ is made (through IP protocol stack 2505) to a DNS 2502 to look up the IP address associated with the name. The DNS returns the IP address DNS RESP to client



New Bay Capital, LLC
Ex.1015-Page 2454 of 3151

application 2504, which is then able to use the IP address to communicate with the host 2503 through separate transactions such as PAGE REQ and PAGE RESP.

In the conventional architecture shown in FIG. 25, nefarious listeners on the Internet could intercept the DNS REQ and DNS RESP packets and thus learn what IP addresses the user was contacting. For example, if a user wanted to set up a secure communication path with a web site having the name "Target.com," when the user's browser contacted a DNS to find the IP address for that web site, the true IP address of that web site would be revealed over the Internet as part of the DNS inquiry. This would hamper anonymous communications on the Internet.

One conventional scheme that provides secure virtual private networks over the Internet provides the DNS server with the public keys of the machines that the DNS server has the addresses for. This allows hosts to retrieve automatically the public keys of a host that the host is to communicate with so that the host can set up a VPN without having the user enter the public key of the destination host. One implementation of this standard is presently being developed as part of the FreeS/WAN project(RFC 2535).

The conventional scheme suffers from certain drawbacks. For example, any user can perform a DNS request. Moreover, DNS requests resolve to the same value for all users.

According to certain aspects of the invention, a specialized DNS server traps DNS requests and, if the request is from a special type of user (e.g., one for which secure communication services are defined), the server does not return the true IP address of the target node, but instead automatically sets up a virtual private network between the target node and the user. The VPN is preferably implemented using the IP address "hopping" features of the basic invention described above, such that the true identity of the two nodes cannot be determined even if packets during the communication are intercepted. For DNS requests that are determined to not require secure services (e.g., an unregistered user), the DNS server transparently "passes through" the request to provide a normal look-up function and return the IP address of the target web server, provided that the requesting host has permissions to resolve



New Bay Capital, LLC
Ex.1015-Page 2455 of 3151

unsecured sites. Different users who make an identical DNS request could be provided with different results.

FIG. 26 shows a system employing various principles summarized above. A user's computer 2601 includes a conventional client (e.g., a web browser) 2605 and an IP protocol stack 2606 that preferably operates in accordance with an IP hopping function 2607 as outlined above. A modified DNS server 2602 includes a conventional DNS server function 2609 and a DNS proxy 2610. A gatekeeper server 2603 is interposed between the modified DNS server and a secure target site 2704. An "unsecure" target site 2611 is also accessible via conventional IP protocols.

According to one embodiment, DNS proxy 2610 intercepts all DNS lookup functions from client 2605 and determines whether access to a secure site has been requested. If access to a secure site has been requested (as determined, for example, by a domain name extension, or by reference to an internal table of such sites), DNS proxy 2610 determines whether the user has sufficient security privileges to access the site. If so, DNS proxy 2610 transmits a message to gatekeeper 2603 requesting that a virtual private network be created between user computer 2601 and secure target site 2604. In one embodiment, gatekeeper 2603 creates "hopblocks" to be used by computer 2601 and secure target site 2604 for secure communication. Then, gatekeeper 2603 communicates these to user computer 2601. Thereafter, DNS proxy 2610 returns to user computer 2601 the resolved address passed to it by the gatekeeper (this address could be different from the actual target computer) 2604, preferably using a secure administrative VPN. The address that is returned need not be the actual address of the destination computer.

Had the user requested lookup of a non-secure web site such as site 2611, DNS proxy would merely pass through to conventional DNS server 2609 the look-up request, which would be handled in a conventional manner, returning the IP address of non-secure web site 2611. If the user had requested lookup of a secure web site but lacked credentials to create such a connection, DNS proxy 2610 would return a "host unknown" error to the user. In this manner, different users requesting access to the same DNS name could be provided with different look-up results.



New Bay Capital, LLC
Ex.1015-Page 2456 of 3151

Gatekeeper 2603 can be implemented on a separate computer (as shown in FIG. 26) or as a function within modified DNS server 2602. In general, it is anticipated that gatekeeper 2703 facilitates the allocation and exchange of information needed to communicate securely, such as using "hopped" IP addresses. Secure hosts such as site 2604 are assumed to be equipped with a secure communication function such as an IP hopping function 2608.

It will be appreciated that the functions of DNS proxy 2610 and DNS server 2609 can be combined into a single server for convenience. Moreover, although element 2602 is shown as combining the functions of two servers, the two servers can be made to operate independently.

FIG. 27 shows steps that can be executed by DNS proxy server 2610 to handle requests for DNS look-up for secure hosts. In step 2701, a DNS look-up request is received for a target host. In step 2702, a check is made to determine whether access to a secure host was requested. If not, then in step 2703 the DNS request is passed to conventional DNS server 2609, which looks up the IP address of the target site and returns it to the user's application for further processing.

In step 2702, if access to a secure host was requested, then in step 2704 a further check is made to determine whether the user is authorized to connect to the secure host. Such a check can be made with reference to an internally stored list of authorized IP addresses, or can be made by communicating with gatekeeper 2603 (e.g., over an "administrative" VPN that is secure). It will be appreciated that different levels of security can also be provided for different categories of hosts. For example, some sites may be designated as having a certain security level, and the security level of the user requesting access must match that security level. The user's security level can also be determined by transmitting a request message back to the user's computer requiring that it prove that it has sufficient privileges.

If the user is not authorized to access the secure site, then a "host unknown" message is returned (step 2705). If the user has sufficient security privileges, then in step 2706 a secure VPN is established between the user's computer and the secure target site. As described above, this is preferably done by allocating a hopping regime



New Bay Capital, LLC
Ex.1015-Page 2457 of 3151

that will be carried out between the user's computer and the secure target site, and is preferably performed transparently to the user (i.e., the user need not be involved in creating the secure link). As described in various embodiments of this application, any of various fields can be "hopped" (e.g., IP source/destination addresses; a field in the header; etc.) in order to communicate securely.

Some or all of the security functions can be embedded in gatekeeper 2603, such that it handles all requests to connect to secure sites. In this embodiment, DNS proxy 2610 communicates with gatekeeper 2603 to determine (preferably over a secure administrative VPN) whether the user has access to a particular web site.

Various scenarios for implementing these features are described by way of example below:

Scenario #1: Client has permission to access target computer, and gatekeeper has a rule to make a VPN for the client. In this scenario, the client's DNS request would be received by the DNS proxy server 2610, which would forward the request to gatekeeper 2603. The gatekeeper would establish a VPN between the client and the requested target. The gatekeeper would provide the address of the destination to the DNS proxy, which would then return the resolved name as a result. The resolved address can be transmitted back to the client in a secure administrative VPN.

Scenario #2: Client does not have permission to access target computer. In this scenario, the client's DNS request would be received by the DNS proxy server 2610, which would forward the request to gatekeeper 2603. The gatekeeper would reject the request, informing DNS proxy server 2610 that it was unable to find the target computer. The DNS proxy 2610 would then return a "host unknown" error message to the client.

Scenario #3: Client has permission to connect using a normal non-VPN link, and the gatekeeper does not have a rule to set up a VPN for the client to the target site. In this scenario, the client's DNS request is received by DNS proxy server 2610, which would check its rules and determine that no VPN is needed. Gatekeeper 2603 would then inform the DNS proxy server to forward the request to conventional DNS



New Bay Capital, LLC
Ex.1015-Page 2458 of 3151

server 2609, which would resolve the request and return the result to the DNS proxy server and then back to the client.

Scenario #4: Client does not have permission to establish a normal/non-VPN link, and the gatekeeper does not have a rule to make a VPN for the client to the target site. In this scenario, the DNS proxy server would receive the client's DNS request and forward it to gatekeeper 2603. Gatekeeper 2603 would determine that no special VPN was needed, but that the client is not authorized to communicate with non-VPN members. The gatekeeper would reject the request, causing DNS proxy server 2610 to return an error message to the client.

### C. Large Link to Small Link Bandwidth Management

One feature of the basic architecture is the ability to prevent so-called "denial of service" attacks that can occur if a computer hacker floods a known Internet node with packets, thus preventing the node from communicating with other nodes. Because IP addresses or other fields are "hopped" and packets arriving with invalid addresses are quickly discarded, Internet nodes are protected against flooding targeted at a single IP address.

In a system in which a computer is coupled through a link having a limited bandwidth (e.g., an edge router) to a node that can support a much higher-bandwidth link (e.g., an Internet Service Provider), a potential weakness could be exploited by a determined hacker. Referring to FIG. 28, suppose that a first host computer 2801 is communicating with a second host computer 2804 using the IP address hopping principles described above. The first host computer is coupled through an edge router 2802 to an Internet Service Provider (ISP) 2803 through a low bandwidth link (LOW BW), and is in turn coupled to second host computer 2804 through parts of the Internet through a high bandwidth link (HIGH BW). In this architecture, the ISP is able to support a high bandwidth to the internet, but a much lower bandwidth to the edge router 2802.

Suppose that a computer hacker is able to transmit a large quantity of dummy packets addressed to first host computer 2801 across high bandwidth link HIGH BW. Normally, host computer 2801 would be able to quickly reject the packets since they



New Bay Capital, LLC
Ex.1015-Page 2459 of 3151

would not fall within the acceptance window permitted by the IP address hopping scheme. However, because the packets must travel across low bandwidth link LOW BW, the packets overwhelm the lower bandwidth link before they are received by host computer 2801. Consequently, the link to host computer 2801 is effectively flooded before the packets can be discarded.

According to one inventive improvement, a "link guard" function 2805 is inserted into the high-bandwidth node (e.g., ISP 2803) that quickly discards packets destined for a low-bandwidth target node if they are not valid packets. Each packet destined for a low-bandwidth node is cryptographically authenticated to determine whether it belongs to a VPN. If it is not a valid VPN packet, the packet is discarded at the high-bandwidth node. If the packet is authenticated as belonging to a VPN, the packet is passed with high preference. If the packet is a valid non-VPN packet, it is passed with a lower quality of service (e.g., lower priority).

In one embodiment, the ISP distinguishes between VPN and non-VPN packets using the protocol of the packet. In the case of IPSEC [rfc 2401], the packets have IP protocols 420 and 421. In the case of the TARP VPN, the packets will have an IP protocol that is not yet defined. The ISP's link guard, 2805, maintains a table of valid VPNs which it uses to validate whether VPN packets are cryptographically valid.

According to one embodiment, packets that do not fall within any hop windows used by nodes on the low-bandwidth link are rejected, or are sent with a lower quality of service. One approach for doing this is to provide a copy of the IP hopping tables used by the low-bandwidth nodes to the high-bandwidth node, such that both the high-bandwidth and low-bandwidth nodes track hopped packets (e.g., the high-bandwidth node moves its hopping window as valid packets are received). In such a scenario, the high-bandwidth node discards packets that do not fall within the hopping window before they are transmitted over the low-bandwidth link. Thus, for example, ISP 2903 maintains a copy 2910 of the receive table used by host computer 2901. Incoming packets that do not fall within this receive table are discarded. According to a different embodiment, link guard 2805 validates each VPN packet using a keyed hashed message authentication code (HMAC) [rfc 2104]. According



New Bay Capital, LLC
Ex.1015-Page 2460 of 3151

to another embodiment, separate VPNs (using, for example, hopblocks) can be established for communicating between the low-bandwidth node and the high-bandwidth node (i.e., packets arriving at the high-bandwidth node are converted into different packets before being transmitted to the low-bandwidth node).

As shown in FIG. 29, for example, suppose that a first host computer 2900 is communicating with a second host computer 2902 over the Internet, and the path includes a high bandwidth link HIGH BW to an ISP 2901 and a low bandwidth link LOW BW through an edge router 2904. In accordance with the basic architecture described above, first host computer 2900 and second host computer 2902 would exchange hopblocks (or a hopblock algorithm) and would be able to create matching transmit and receive tables 2905, 2906, 2912 and 2913. Then in accordance with the basic architecture, the two computers would transmit packets having seemingly random IP source and destination addresses, and each would move a corresponding hopping window in its receive table as valid packets were received.

Suppose that a nefarious computer hacker 2903 was able to deduce that packets having a certain range of IP addresses (e.g., addresses 100 to 200 for the sake of simplicity) are being transmitted to ISP 2901, and that these packets are being forwarded over a low-bandwidth link. Hacker computer 2903 could thus "flood" packets having addresses falling into the range 100 to 200, expecting that they would be forwarded along low bandwidth link LOW BW, thus causing the low bandwidth link to become overwhelmed. The fast packet reject mechanism in first host computer 3000 would be of little use in rejecting these packets, since the low bandwidth link was effectively jammed before the packets could be rejected. In accordance with one aspect of the improvement, however, VPN link guard 2911 would prevent the attack from impacting the performance of VPN traffic because the packets would either be rejected as invalid VPN packets or given a lower quality of service than VPN traffic over the lower bandwidth link. A denial-of-service flood attack could, however, still disrupt non-VPN traffic.

According to one embodiment of the improvement, ISP 2901 maintains a separate VPN with first host computer 2900, and thus translates packets arriving at the



New Bay Capital, LLC
Ex.1015-Page 2461 of 3151

ISP into packets having a different IP header before they are transmitted to host computer 2900. The cryptographic keys used to authenticate VPN packets at the link guard 2911 and the cryptographic keys used to encrypt and decrypt the VPN packets at host 2902 and host 2901 can be different, so that link guard 2911 does not have access to the private host data; it only has the capability to authenticate those packets.

According to yet a third embodiment, the low-bandwidth node can transmit a special message to the high-bandwidth node instructing it to shut down all transmissions on a particular IP address, such that only hopped packets will pass through to the low-bandwidth node. This embodiment would prevent a hacker from flooding packets using a single IP address. According to yet a fourth embodiment, the high-bandwidth node can be configured to discard packets transmitted to the low-bandwidth node if the transmission rate exceeds a certain predetermined threshold for any given IP address; this would allow hopped packets to go through. In this respect, link guard 2911 can be used to detect that the rate of packets on a given IP address are exceeding a threshold rate; further packets addressed to that same IP address would be dropped or transmitted at a lower priority (e.g., delayed).

**D. Traffic Limiter**

In a system in which multiple nodes are communicating using "hopping" technology, a treasonous insider could internally flood the system with packets. In order to prevent this possibility, one inventive improvement involves setting up "contracts" between nodes in the system, such that a receiver can impose a bandwidth limitation on each packet sender. One technique for doing this is to delay acceptance of a checkpoint synchronization request from a sender until a certain time period (e.g., one minute) has elapsed. Each receiver can effectively control the rate at which its hopping window moves by delaying "SYNC ACK" responses to "SYNC_REQ" messages.

A simple modification to the checkpoint synchronizer will serve to protect a receiver from accidental or deliberate overload from an internally treasonous client. This modification is based on the observation that a receiver will not update its tables until a SYNC_REQ is received on hopped address CKPT_N. It is a simple matter of



New Bay Capital, LLC
Ex.1015-Page 2462 of 3151

deferring the generation of a new CKPT_N until an appropriate interval after previous checkpoints.

Suppose a receiver wished to restrict reception from a transmitter to 100 packets a second, and that checkpoint synchronization messages were triggered every 50 packets. A compliant transmitter would not issue new SYNC_REQ messages more often than every 0.5 seconds. The receiver could delay a non-compliant transmitter from synchronizing by delaying the issuance of CKPT_N for 0.5 second after the last SYNC_REQ was accepted.

In general, if M receivers need to restrict N transmitters issuing new SYNC_REQ messages after every W messages to sending R messages a second in aggregate, each receiver could defer issuing a new CKPT_N until MxNxW/R seconds have elapsed since the last SYNC_REQ has been received and accepted. If the transmitter exceeds this rate between a pair of checkpoints, it will issue the new checkpoint before the receiver is ready to receive it, and the SYNC_REQ will be discarded by the receiver. After this, the transmitter will re-issue the SYNC_REQ every T1 seconds until it receives a SYNC_ACK. The receiver will eventually update CKPT_N and the SYNC_REQ will be acknowledged. If the transmission rate greatly exceeds the allowed rate, the transmitter will stop until it is compliant. If the transmitter exceeds the allowed rate by a little, it will eventually stop after several rounds of delayed synchronization until it is in compliance. Hacking the transmitter's code to not shut off only permits the transmitter to lose the acceptance window. In this case it can recover the window and proceed only after it is compliant again.

Two practical issues should be considered when implementing the above scheme:

1. The receiver rate should be slightly higher than the permitted rate in order to allow for statistical fluctuations in traffic arrival times and non-uniform load balancing.

2. Since a transmitter will rightfully continue to transmit for a period after a SYNC_REQ is transmitted, the algorithm above can artificially reduce the transmitter's bandwidth. If events prevent a compliant transmitter from synchronizing



New Bay Capital, LLC
Ex.1015-Page 2463 of 3151

for a period (e.g. the network dropping a SYNC_REQ or a SYNC_ACK) a SYNC_REQ will be accepted later than expected. After this, the transmitter will transmit fewer than expected messages before encountering the next checkpoint. The new checkpoint will not have been activated and the transmitter will have to retransmit the SYNC_REQ. This will appear to the receiver as if the transmitter is not compliant. Therefore, the next checkpoint will be accepted late from the transmitter's perspective. This has the effect of reducing the transmitter's allowed packet rate until the transmitter transmits at a packet rate below the agreed upon rate for a period of time.

To guard against this, the receiver should keep track of the times that the last C SYNC_REQs were received and accepted and use the minimum of MxNxW/R seconds after the last SYNC_REQ has been received and accepted, 2xMxNxW/R seconds after next to the last SYNC_REQ has been received and accepted, CxMxNxW/R seconds after $(C-1)^{th}$ to the last SYNC_REQ has been received, as the time to activate CKPT_N. This prevents the receiver from inappropriately limiting the transmitter's packet rate if at least one out of the last C SYNC_REQs was processed on the first attempt.

FIG. 30 shows a system employing the above-described principles. In FIG. 30, two computers 3000 and 3001 are assumed to be communicating over a network N in accordance with the "hopping" principles described above (e.g., hopped IP addresses, discriminator values, etc.). For the sake of simplicity, computer 3000 will be referred to as the receiving computer and computer 3001 will be referred to as the transmitting computer, although full duplex operation is of course contemplated. Moreover, although only a single transmitter is shown, multiple transmitters can transmit to receiver 3000.

As described above, receiving computer 3000 maintains a receive table 3002 including a window W that defines valid IP address pairs that will be accepted when appearing in incoming data packets. Transmitting computer 3001 maintains a transmit table 3003 from which the next IP address pairs will be selected when transmitting a packet to receiving computer 3000. (For the sake of illustration,



New Bay Capital, LLC
Ex.1015-Page 2464 of 3151

window W is also illustrated with reference to transmit table 3003). As transmitting computer moves through its table, it will eventually generate a SYNC_REQ message as illustrated in function 3010. This is a request to receiver 3000 to synchronize the receive table 3002, from which transmitter 3001 expects a response in the form of a CKPT_N (included as part of a SYNC_ACK message). If transmitting computer 3001 transmits more messages than its allotment, it will prematurely generate the SYNC_REQ message. (If it has been altered to remove the SYNC_REQ message generation altogether, it will fall out of synchronization since receiver 3000 will quickly reject packets that fall outside of window W, and the extra packets generated by transmitter 3001 will be discarded).

In accordance with the improvements described above, receiving computer 3000 performs certain steps when a SYNC_REQ message is received, as illustrated in FIG. 30. In step 3004, receiving computer 3000 receives the SYNC_REQ message. In step 3005, a check is made to determine whether the request is a duplicate. If so, it is discarded in step 3006. In step 3007, a check is made to determine whether the SYNC_REQ received from transmitter 3001 was received at a rate that exceeds the allowable rate R (i.e., the period between the time of the last SYNC_REQ message). The value R can be a constant, or it can be made to fluctuate as desired. If the rate exceeds R, then in step 3008 the next activation of the next CKPT_N hopping table entry is delayed by W/R seconds after the last SYNC_REQ has been accepted.

Otherwise, if the rate has not been exceeded, then in step 3109 the next CKPT_N value is calculated and inserted into the receiver's hopping table prior to the next SYNC_REQ from thetransmitter 3101. Transmitter 3101 then processes the SYNC_REQ in the normal manner.

### E. Signaling Synchronizer

In a system in which a large number of users communicate with a central node using secure hopping technology, a large amount of memory must be set aside for hopping tables and their supporting data structures. For example, if one million subscribers to a web site occasionally communicate with the web site, the site must maintain one million hopping tables, thus using up valuable computer resources, even



New Bay Capital, LLC
Ex.1015-Page 2465 of 3151

though only a small percentage of the users may actually be using the system at any one time. A desirable solution would be a system that permits a certain maximum number of simultaneous links to be maintained, but which would "recognize" millions of registered users at any one time. In other words, out of a population of a million registered users, a few thousand at a time could simultaneously communicate with a central server, without requiring that the server maintain one million hopping tables of appreciable size.

One solution is to partition the central node into two nodes: a signaling server that performs session initiation for user log-on and log-off (and requires only minimally sized tables), and a transport server that contains larger hopping tables for the users. The signaling server listens for the millions of known users and performs a fast-packet reject of other (bogus) packets. When a packet is received from a known user, the signaling server activates a virtual private link (VPL) between the user and the transport server, where hopping tables are allocated and maintained. When the user logs onto the signaling server, the user's computer is provided with hop tables for communicating with the transport server, thus activating the VPL. The VPLs can be torn down when they become inactive for a time period, or they can be torn down upon user log-out. Communication with the signaling server to allow user log-on and log-off can be accomplished using a specialized version of the checkpoint scheme described above.

FIG. 31 shows a system employing certain of the above-described principles. In FIG. 31, a signaling server 3101 and a transport server 3102 communicate over a link. Signaling server 3101 contains a large number of small tables 3106 and 3107 that contain enough information to authenticate a communication request with one or more clients 3103 and 3104. As described in more detail below, these small tables may advantageously be constructed as a special case of the synchronizing checkpoint tables described previously. Transport server 3102, which is preferably a separate computer in communication with signaling server 3101, contains a smaller number of larger hopping tables 3108, 3109, and 3110 that can be allocated to create a VPN with one of the client computers.



New Bay Capital, LLC
Ex.1015-Page 2466 of 3151

According to one embodiment, a client that has previously registered with the system (e.g., via a system administration function, a user registration procedure, or some other method) transmits a request for information from a computer (e.g., a web site). In one variation, the request is made using a "hopped" packet, such that signaling server 3101 will quickly reject invalid packets from unauthorized computers such as hacker computer 3105. An "administrative" VPN can be established between all of the clients and the signaling server in order to ensure that a hacker cannot flood signaling server 3101 with bogus packets. Details of this scheme are provided below.

Signaling server 3101 receives the request 3111 and uses it to determine that client 3103 is a validly registered user. Next, signaling server 3101 issues a request to transport server 3102 to allocate a hopping table (or hopping algorithm or other regime) for the purpose of creating a VPN with client 3103. The allocated hopping parameters are returned to signaling server 3101 (path 3113), which then supplies the hopping parameters to client 3103 via path 3114, preferably in encrypted form.

Thereafter, client 3103 communicates with transport server 3102 using the normal hopping techniques described above. It will be appreciated that although signaling server 3101 and transport server 3102 are illustrated as being two separate computers, they could of course be combined into a single computer and their functions performed on the single computer. Alternatively, it is possible to partition the functions shown in FIG. 31 differently from as shown without departing from the inventive principles.

One advantage of the above-described architecture is that signaling server 3101 need only maintain a small amount of information on a large number of potential users, yet it retains the capability of quickly rejecting packets from unauthorized users such as hacker computer 3105. Larger data tables needed to perform the hopping and synchronization functions are instead maintained in a transport server 3102, and a smaller number of these tables are needed since they are only allocated for "active" links. After a VPN has become inactive for a certain time period (e.g., one hour), the VPN can be automatically torn down by transport server 3102 or signaling server 3101.



New Bay Capital, LLC
Ex.1015-Page 2467 of 3151

A more detailed description will now be provided regarding how a special case of the checkpoint synchronization feature can be used to implement the signaling scheme described above.

The signaling synchronizer may be required to support many (millions) of standing, low bandwidth connections. It therefore should minimize per-VPL memory usage while providing the security offered by hopping technology. In order to reduce memory usage in the signaling server, the data hopping tables can be completely eliminated and data can be carried as part of the SYNC_REQ message. The table used by the server side (receiver) and client side (transmitter) is shown schematically as element 3106 in FIG. 31.

The meaning and behaviors of CKPT_N, CKPT_O and CKPT_R remain the same from the previous description, except that CKPT_N can receive a combined data and SYNC_REQ message or a SYNC_REQ message without the data.

The protocol is a straightforward extension of the earlier synchronizer. Assume that a client transmitter is on and the tables are synchronized. The initial tables can be generated "out of band." For example, a client can log into a web server to establish an account over the Internet. The client will receive keys etc encrypted over the Internet. Meanwhile, the server will set up the signaling VPN on the signaling server.

Assuming that a client application wishes to send a packet to the server on the client's standing signaling VPL:

1. The client sends the message marked as a data message on the inner header using the transmitter's CKPT_N address. It turns the transmitter off and starts a timer T1 noting CKPT_O. Messages can be one of three types: DATA, SYNC_REQ and SYNC_ACK. In the normal algorithm, some potential problems can be prevented by identifying each message type as part of the encrypted inner header field. In this algorithm, it is important to distinguish a data packet and a SYNC_REQ in the signaling synchronizer since the data and the SYNC_REQ come in on the same address.



New Bay Capital, LLC
Ex.1015-Page 2468 of 3151

2. When the server receives a data message on its CKPT_N, it verifies the message and passes it up the stack. The message can be verified by checking message type and and other information (i.e user credentials) contained in the inner header It replaces its CKPT_O with CKPT_N and generates the next CKPT_N. It updates its transmitter side CKPT_R to correspond to the client's receiver side CKPT_R and transmits a SYNC_ACK containing CKPT_O in its payload.

3. When the client side receiver receives a SYNC_ACK on its CKPT_R with a payload matching its transmitter side CKPT_O and the transmitter is off, the transmitter is turned on and the receiver side CKPT_R is updated. If the SYNC_ACK's payload does not match the transmitter side CKPT_O or the transmitter is on, the SYNC_ACK is simply discarded.

4. T1 expires: If the transmitter is off and the client's transmitter side CKPT_O matches the CKPT_O associated with the timer, it starts timer T1 noting CKPT_O again, and a SYNC_REQ is sent using the transmitter's CKPT_O address. Otherwise, no action is taken.

5. When the server receives a SYNC_REQ on its CKPT_N, it replaces its CKPT_O with CKPT_N and generates the next CKPT_N. It updates its transmitter side CKPT_R to correspond to the client's receiver side CKPT_R and transmits a SYNC_ACK containing CKPT_O in its payload.

6. When the server receives a SYNC_REQ on its CKPT_O, it updates its transmitter side CKPT_R to correspond to the client's receiver side CKPT_R and transmits a SYNC_ACK containing CKPT_O in its payload.

FIG. 32 shows message flows to highlight the protocol. Reading from top to bottom, the client sends data to the server using its transmitter side CKPT_N. The client side transmitter is turned off and a retry timer is turned off. The transmitter will not transmit messages as long as the transmitter is turned off. The client side transmitter then loads CKPT_N into CKPT_O and updates CKPT_N. This message is successfully received and a passed up the stack. It also synchronizes the receiver i.e, the server loads CKPT_N into CKPT_O and generates a new CKPT_N, it generates a new CKPT_R in the server side transmitter and transmits a SYNC_ACK containing



New Bay Capital, LLC
Ex.1015-Page 2469 of 3151

the server side receiver's CKPT_O the server. The SYNC_ACK is successfully received at the client. The client side receiver's CKPT_R is updated, the transmitter is turned on and the retry timer is killed. The client side transmitter is ready to transmit a new data message.

Next, the client sends data to the server using its transmitter side CKPT_N. The client side transmitter is turned off and a retry timer is turned off. The transmitter will not transmit messages as long as the transmitter is turned off. The client side transmitter then loads CKPT_N into CKPT_O and updates CKPT_N. This message is lost. The client side timer expires and as a result a SYNC_REQ is transmitted on the client side transmitter's CKPT_O (this will keep happening until the SYNC_ACK has been received at the client). The SYNC_REQ is successfully received at the server. It synchronizes the receiver i.e, the server loads CKPT_N into CKPT_O and generates a new CKPT_N, it generates an new CKPT_R in the server side transmitter and transmits a SYNC_ACK containing the server side receiver's CKPT_O the server. The SYNC_ACK is successfully received at the client. The client side receiver's CKPT_R is updated, the transmitter is turned off and the retry timer is killed. The client side transmitter is ready to transmit a new data message.

There are numerous other scenarios that follow this flow. For example, the SYNC_ACK could be lost. The transmitter would continue to re-send the SYNC_REQ until the receiver synchronizes and responds.

The above-described procedures allow a client to be authenticated at signaling server 3201 while maintaining the ability of signaling server 3201 to quickly reject invalid packets, such as might be generated by hacker computer 3205. In various embodiments, the signaling synchronizer is really a derivative of the synchronizer. It provides the same protection as the hopping protocol, and it does so for a large number of low bandwidth connections.



New Bay Capital, LLC
Ex.1015-Page 2470 of 3151

## CLAIMS

1. A method of transmitting data packets between a first computer and a second computer, wherein the first computer and the second computer are linked via a plurality of separate transmission paths, the method comprising the steps of:

(1) assigning a weight value to each of the plurality of transmission paths, wherein each respective weight value represents the relative number of packets that a respective transmission path will transmit;

(2) for each data packet that is to be transmitted from the first computer to the second computer, selecting one of the plurality of transmission paths on the basis of each respective transmission path's assigned weight value;

(3) measuring the transmission quality for each of the plurality of transmission paths; and

(4) adjusting downwardly to a non-zero value the assigned weight value for a transmission path for which the transmission quality has declined.

2. The method of claim 1, wherein step (4) comprises the step of gradually decreasing over time the assigned weight value in relation to weight values assigned to the remaining transmission paths.

3. The method of claim 2, wherein step (4) comprises the step of gradually decreasing the assigned weight value according to an incrementally decreasing function.

4. The method of claim 2, wherein step (4) comprises the step of gradually decreasing the assigned weight value according to an exponentially decaying function.

5. The method of claim 1, wherein step (3) comprises the step of determining that one or more packets transmitted to the second computer was not acknowledged by the second computer.

6. The method of claim 1, wherein step (3) comprises the step of evaluating the contents of a synchronization packet that maintains synchronization with a moving window of valid values.



New Bay Capital, LLC
Ex.1015-Page 2471 of 3151

7. The method of claim 1, further comprising the step of inserting into each data packet a source and destination IP address pair that is selected according to a pseudo-random sequence.

8. The method of claim 1, wherein step (4) comprises the step of adjusting downwardly the assigned weight value for a transmission path only if the transmission quality has declined below a predetermined threshold.

9. The method of claim 1, further comprising the step of adjusting upwardly the assigned weight value that was adjusted in step (4) if it is later determined that the transmission quality has improved.

10. The method of claim 1, further comprising the step of adjusting upwardly the weight values of the remaining transmission links in an amount that compensates for the downwardly adjusted weight value.

11. The method of claim 10, wherein the step of adjusting upwardly comprises the step of equally distributing the amount that was downwardly adjusted across the remaining transmission links.

12. The method of claim 1, further comprising the step of adjusting downwardly to zero the assigned weight value for any transmission link whose quality has degraded below a preset threshold.

13. The method of claim 1, wherein steps (2) through (4) are repeated periodically.

14. A first computer that transmits data packets to a second computer over a plurality of separate transmission paths, wherein the first computer performs the steps of:

(1) assigning a weight value to each of the plurality of transmission paths, wherein each respective weight value represents the relative number of packets that a respective transmission path will transmit;

(2) for each data packet that is to be transmitted to the second computer, selecting one of the plurality of transmission paths on the basis of each respective transmission path's assigned weight value;



New Bay Capital, LLC
Ex.1015-Page 2472 of 3151

(3) measuring the transmission quality for each of the plurality of transmission paths; and

(4) adjusting downwardly to a non-zero value the assigned weight value for a transmission path for which the transmission quality has declined.

15. The first computer of claim 14, wherein the first computer gradually decreases over time the assigned weight value in relation to weight values assigned to the remaining transmission paths.

16. The first computer of claim 15, wherein the first computer gradually decreases the assigned weight value according to an incrementally decreasing function.

17. The first computer of claim 15, wherein the first computer gradually decreases the assigned weight value according to an exponentially decaying function.

18. The first computer of claim 14, wherein the first computer measures the transmission quality by determining that one or more packets transmitted to the second computer was not acknowledged by the second computer.

19. The first computer of claim 14, wherein the first computer measures the transmission quality by evaluating the contents of a synchronization packet that maintains synchronization with a moving window of valid values.

20. The first computer of claim 14, wherein the first computer inserts into each data packet a source and destination IP address pair that is selected according to a pseudo-random sequence.

21. The first computer of claim 14, wherein the first computer adjusts downwardly the assigned weight value for any transmission path only if the transmission quality has declined below a predetermined threshold.

22. The first computer of claim 14, wherein the first computer adjusts upwardly the assigned weight value that was adjusted in step (4) if it is later determined that the transmission quality has improved.

23. The first computer of claim 14, wherein the first computer adjusts upwardly the weight values of the remaining transmission links in an amount that compensates for the downwardly adjusted weight value.



New Bay Capital, LLC
Ex.1015-Page 2473 of 3151

24. The first computer of claim 23, wherein the first computer upwardly adjusts probabilities across the remaining transmission links in an amount equal to the downwardly adjusted weight value.

25. The first computer of claim 14, wherein the first computer adjusts downwardly to zero the assigned weight value for any transmission link whose quality has degraded below a preset threshold.

26. The first computer of claim 14, wherein the first computer repeats steps (2) through (4) periodically.

27. A system comprising the first computer of claim 14 and a second computer constructed in accordance with the first computer of claim 14.

28. A method of transparently creating a virtual private network (VPN) between a client computer and a target computer, comprising the steps of:

(1) generating from the client computer a Domain Name Service (DNS) request that requests an IP address corresponding to a domain name associated with the target computer;

(2) determining whether the DNS request transmitted in step (1) is requesting access to a secure web site; and

(3) in response to determining that the DNS request in step (2) is requesting access to a secure target web site, automatically initiating the VPN between the client computer and the target computer.

29. The method of claim 28, wherein steps (2) and (3) are performed at a DNS server separate from the client computer.

30. The method of claim 28, further comprising the step of:

(4) in response to determining that the DNS request in step (2) is not requesting access to a secure target web site, resolving the IP address for the domain name and returning the IP address to the client computer.

31. The method of claim 28, wherein step (3) comprises the step of, prior to automatically initiating the VPN between the client computer and the target computer, determining whether the client computer is authorized to establish a VPN with the target computer and, if not so authorized, returning an error from the DNS request.



New Bay Capital, LLC
Ex.1015-Page 2474 of 3151

32. The method of claim 28, wherein step (3) comprises the step of, prior to automatically initiating the VPN between the client computer and the target computer, determining whether the client computer is authorized to resolve addresses of non secure target computers and, if not so authorized, returning an error from the DNS request.

33. The method of claim 28, wherein step (3) comprises the step of establishing the VPN by creating an IP address hopping scheme between the client computer and the target computer.

34. The method of claim 28, wherein step (3) comprises the step of using a gatekeeper computer that allocates VPN resources for communicating between the client computer and the target computer.

35. The method of claim 28, wherein step (2) is performed in a DNS proxy server that passes through the request to a DNS server if it is determined in step (3) that access is not being requested to a secure target web site.

36. The method of claim 32, wherein step (3) comprises the step of transmitting a message to the client computer to determine whether the client computer is authorized to establish the VPN target computer.

37. A system that transparently creates a virtual private network (VPN) between a client computer and a secure target computer, comprising:

a DNS proxy server that receives a request from the client computer to look up an IP address for a domain name, wherein the DNS proxy server returns the IP address for the requested domain name if it is determined that access to a non-secure web site has been requested, and wherein the DNS proxy server generates a request to create the VPN between the client computer and the secure target computer if it is determined that access to a secure web site has been requested; and

a gatekeeper computer that allocates resources for the VPN between the client computer and the secure web computer in response to the request by the DNS proxy server.

38. The system of claim 37, wherein the gatekeeper computer creates the VPN by establishing an IP address hopping regime that is used to pseudorandomly



New Bay Capital, LLC
Ex.1015-Page 2475 of 3151

change IP addresses in packets transmitted between the client computer and the secure target computer.

39. The system of claim 37, wherein the gatekeeper computer determines whether the client computer has sufficient security privileges to create the VPN and, if the client computer lacks sufficient security privileges, rejecting the request to create the VPN.

40. A method of preventing data packets received from a high bandwidth link from flooding a low bandwidth link, comprising the steps of:

(1) receiving data packets from the high bandwidth link that are ostensibly addressed to a computer residing on the low-bandwidth link;

(2) for each data packet, determining whether the data packet is validly addressed to the computer on the low-bandwidth link;

(3) in response to determining that the data packet is not validly addressed to the computer on the low-bandwidth link, rejecting the data packet; and

(4) in response to determining that the data packet is validly addressed to the computer on the low-bandwidth link, forwarding the data packet to the computer over the low-bandwidth link.

41. The method of claim 40, wherein step (3) comprises the step of comparing a value in a header of each data packet to a set of valid values maintained for the computer on the low-bandwidth link.

42. The method of claim 41, wherein step (3) comprises the step of comparing a value in a header of each data packet to a moving window of valid values.

43. The method of claim 42, wherein step (3) comprises the step of comparing the IP address in the header of each data packet to a moving window of valid IP addresses, wherein the moving window is also maintained by the computer on the low-bandwidth link.

44. The method of claim 40, wherein step (3) comprises the step of reducing a priority level of the packet in relation to other data packets, wherein the priority level determines whether a particular data packet will be transmitted before another data packet having a different priority level.



New Bay Capital, LLC
Ex.1015-Page 2476 of 3151

45. The method of claim 40, wherein step (3) comprises the step of performing a cryptographic check on each data packet to determine whether each data packet is validly addressed.

46. The method of claim 40, wherein step (3) comprises the step of receiving a message from the computer on the low-bandwidth link to stop accepting messages having a particular characteristic.

47. The method of claim 46, wherein step (3) comprises the step of receiving a message from the computer on the low-bandwidth link to stop accepting messages addressed to a particular IP address.

48. The method of claim 40, wherein step (3) comprises the step of determining that a packet transmission rate has been exceeded for a given packet parameter.

49. The method of claim 48, wherein step (3) comprises the step of determining that a packet transmission rate has been exceeded for a given IP destination address.

50. In a system having a low bandwidth data link, a first computer coupled to the low bandwidth data link, and a high bandwidth data link, an improvement comprising:

a second computer coupled between the low bandwidth data link and the high bandwidth data link, wherein the second computer receives data packets from the high bandwidth data link and, if they are addressed to the first computer, routes them to the first computer over the low bandwidth data link,

wherein the second computer prevents invalid data packets ostensibly addressed to the first computer from being transmitted over the low bandwidth data link.

51. The system of claim 50, wherein the second computer prevents invalid data packets from being transmitted over the low bandwidth data link by comparing a discriminator field in a header of each data packet to a table of valid discriminator fields maintained for the first computer.



New Bay Capital, LLC
Ex.1015-Page 2477 of 3151

52. The system of claim 50, wherein the second computer compares an Internet Protocol (IP) address in a header of each data packet to a table of valid IP addresses.

53. The system of claim 52, wherein the second computer compares the IP address in the header of each data packet to a moving window of valid IP addresses, wherein the moving window is also maintained by the first computer.

54. The system of claim 50, wherein the second computer reduces a priority level of a data packet in relation to other data packets, wherein the priority level determines whether a particular data packet will be transmitted before another data packet having a different priority level.

55. The system of claim 50, wherein the second computer performs a cryptographic check on each data packet to determine whether each data packet is validly addressed.

56. The system of claim 50, wherein the second computer receives a message from the first computer that causes the second computer to stop accepting messages having a particular characteristic.

57. The system of claim 56, wherein the second computer receiving a message from the first computer to stop accepting messages addressed to a particular IP address.

58. The system of claim 50, wherein the second computer rejects invalid packets by determining that a packet transmission rate has been exceeded for a given packet parameter.

59. The system of claim 58, wherein the second computer determines that a packet transmission rate has been exceeded for a given IP destination address.

60. In a system comprising a first computer that transmits data packets to a second computer over a network according to a scheme by which at least one field in a series of data packets is periodically changed according to a sequence known by the first and second computers, and wherein the second computer periodically receives a synchronization request from the first computer to maintain synchronization of the sequence between the first and second computers, a method comprising the steps of:



New Bay Capital, LLC
Ex.1015-Page 2478 of 3151

(1) receiving at the first computer the synchronization request from the second computer;

(2) determining whether the synchronization request was received in less than a predetermined interval;

(3) in response to determining that the synchronization request was received in less than the predetermined interval, ignoring the synchronization request; and

(4) in response to determining that the synchronization request was not received in less than the predetermined interval, providing the synchronization response to the first computer.

61. The method of claim 60, wherein step (3) comprises the step of delaying the acceptance of a SYNC_REQ for W/R seconds, where W is the number of data packets between synchronization requests according to an agreed schedule, and R is the agreed rate at which synchronization requests should be received according to the agreed schedule.

62. The method of claim 60, further comprising the step of determining whether the synchronization request is a duplicate of a previously received synchronization request and, if it is a duplicate, discarding it.

63. The method of claim 60, wherein step (4) comprises the step of providing a response that includes a new checkpoint for synchronizing a window in a hopping table.

64. A computer that receives data packets from a second computer over a network according to a scheme by which at least one field in a series of data packets is periodically changed according to a known sequence, wherein the second computer periodically transmits a synchronization request to maintain synchronization of the sequence, wherein the computer performs the steps of:

(1) receiving the synchronization request from the second computer;

(2) determining whether the synchronization request was received in less than a predetermined interval;

(3) in response to determining that the synchronization request was received in less than a predetermined interval ignoring the synchronization request; and



New Bay Capital, LLC
Ex.1015-Page 2479 of 3151

(4) in response to determining that the synchronization request was not received in less than a predetermined interval, providing the response to the first computer.

65. The computer of claim 64, wherein the computer delays the acceptance of a SYNC_REQ in step (3) for W/R seconds, where W is the number of data packets between synchronization requests according to an agreed schedule, and R is the agreed rate at which synchronization requests should be received according to the agreed schedule.

66. The computer of claim 64, wherein the computer further performs the step of determining whether the synchronization request is a duplicate of a previously received synchronization request and, if it is a duplicate, discarding it.

67. A method of establishing communication between one of a plurality of client computers and a central computer that maintains a plurality of authentication tables each corresponding to one of the client computers, the method comprising the steps of:

(1) in the central computer, receiving from one of the plurality of client computers a request to establish a connection;

(2) authenticating, with reference to one of the plurality of authentication tables, that the request received in step (1) is from an authorized client;

(3) responsive to a determination that the request is from an authorized client, allocating resources to establish a virtual private link between the client and a second computer; and

(4) communicating between the authorized client and the second computer using the virtual private link.

68. The method of claim 67, wherein step (4) comprises the step of communicating according to a scheme by which at least one field in a series of data packets is periodically changed according to a known sequence.

69. The method of claim 68, wherein step (4) comprises the step of comparing an Internet Protocol (IP) address in a header of each data packet to a table of valid IP addresses maintained in a table in the second computer.



New Bay Capital, LLC
Ex.1015-Page 2480 of 3151

70. The method of claim 69, wherein step (4) comprises the step of comparing the IP address in the header of each data packet to a moving window of valid IP addresses, and rejecting data packets having IP addresses that do not fall within the moving window.

71. The method of claim 67, wherein step (2) comprises the step of using a checkpoint data structure that maintains synchronization of a periodically changing parameter known by the central computer and the client computer to authenticate the client.



New Bay Capital, LLC
Ex.1015-Page 2481 of 3151

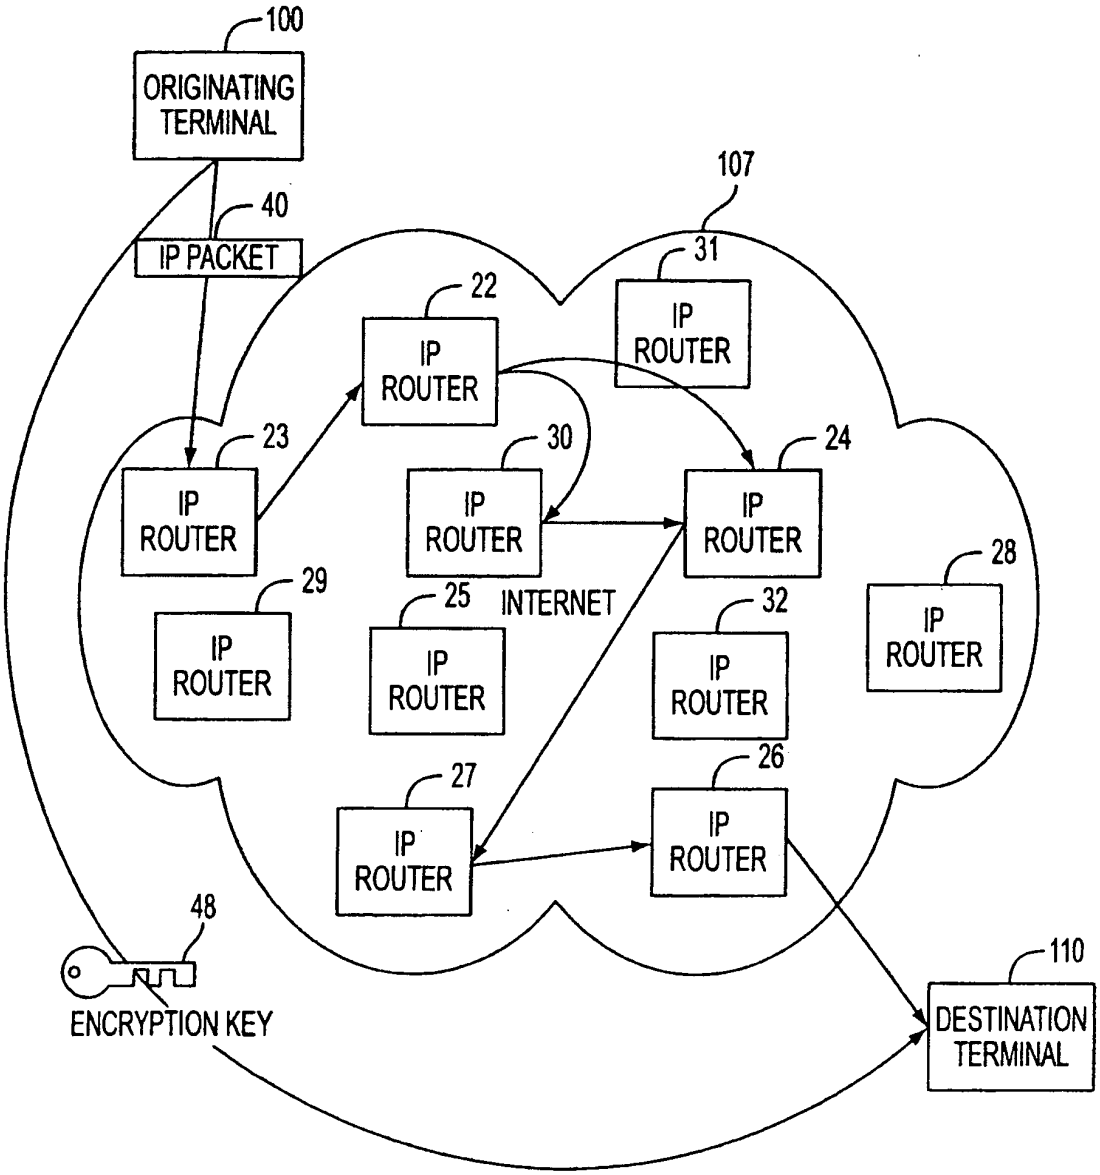
100
ORIGINATING TERMINAL
40
IP PACKET
107
31
IP ROUTER
22
IP ROUTER
23
IP ROUTER
30
IP ROUTER
24
IP ROUTER
28
IP ROUTER
29
IP ROUTER
25
IP ROUTER
INTERNET
32
IP ROUTER
27
IP ROUTER
26
IP ROUTER
48
ENCRYPTION KEY
110
DESTINATION TERMINAL

FIG. 1

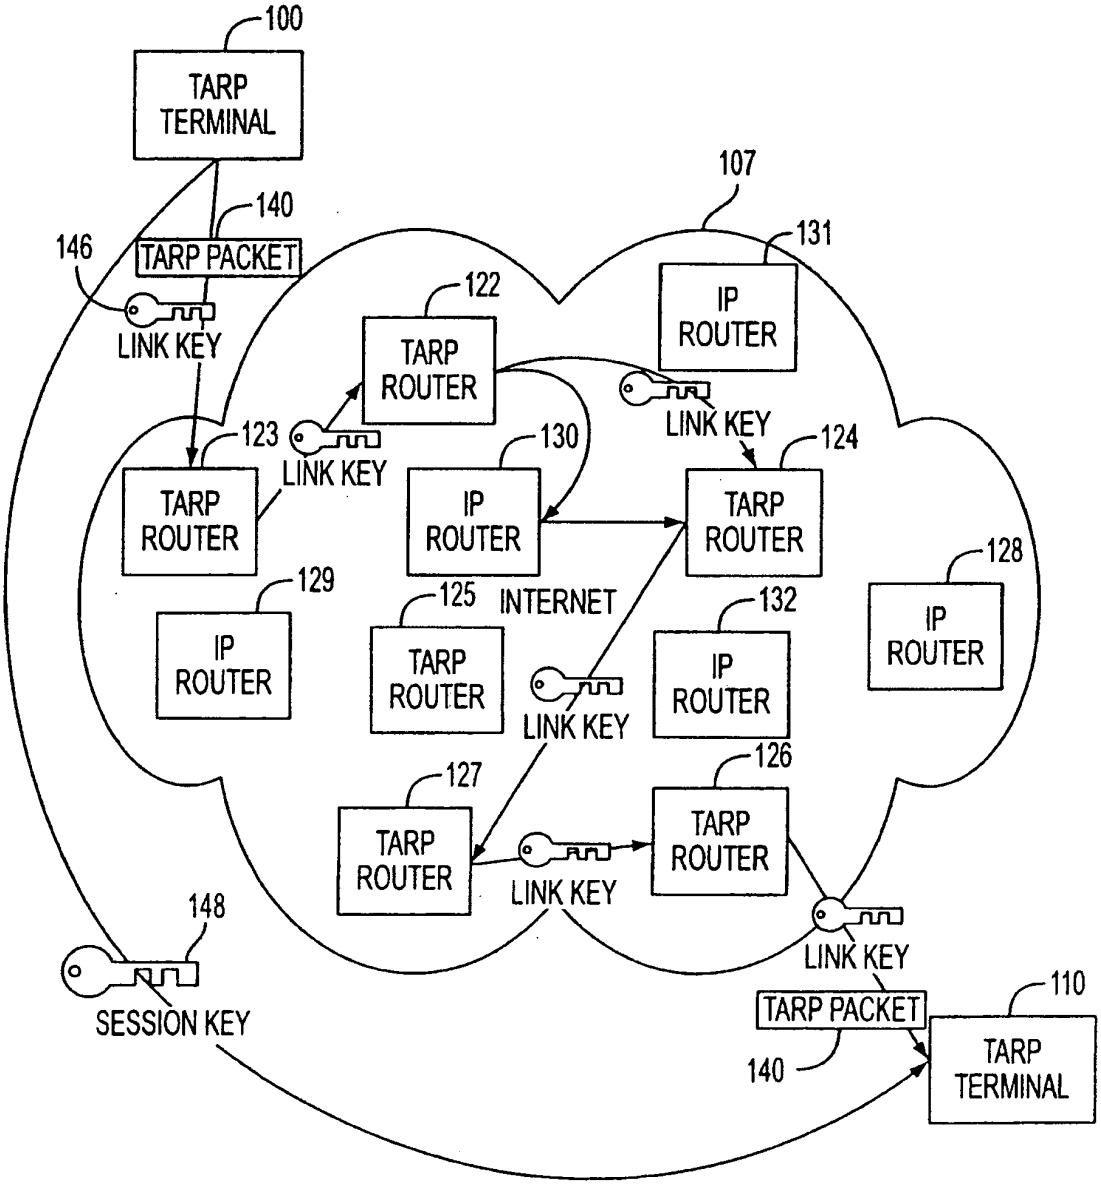


FIG. 2



New Bay Capital, LLC
Ex.1015-Page 2483 of 3151

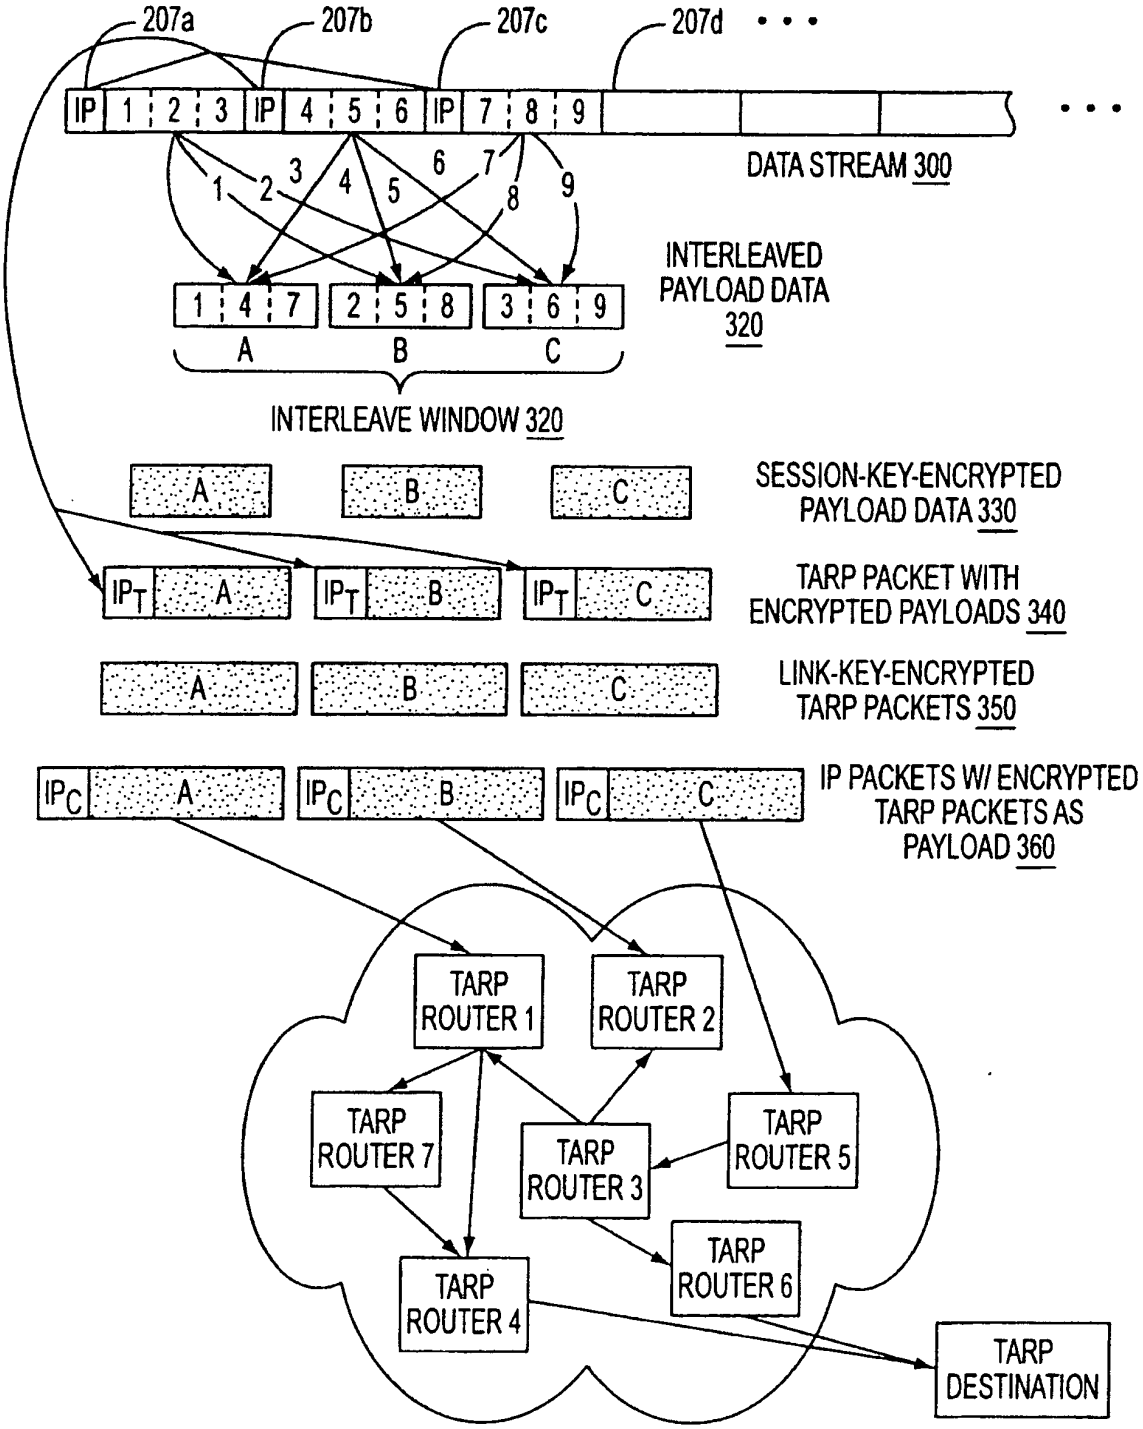


FIG. 3A



New Bay Capital, LLC
Ex.1015-Page 2484 of 3151

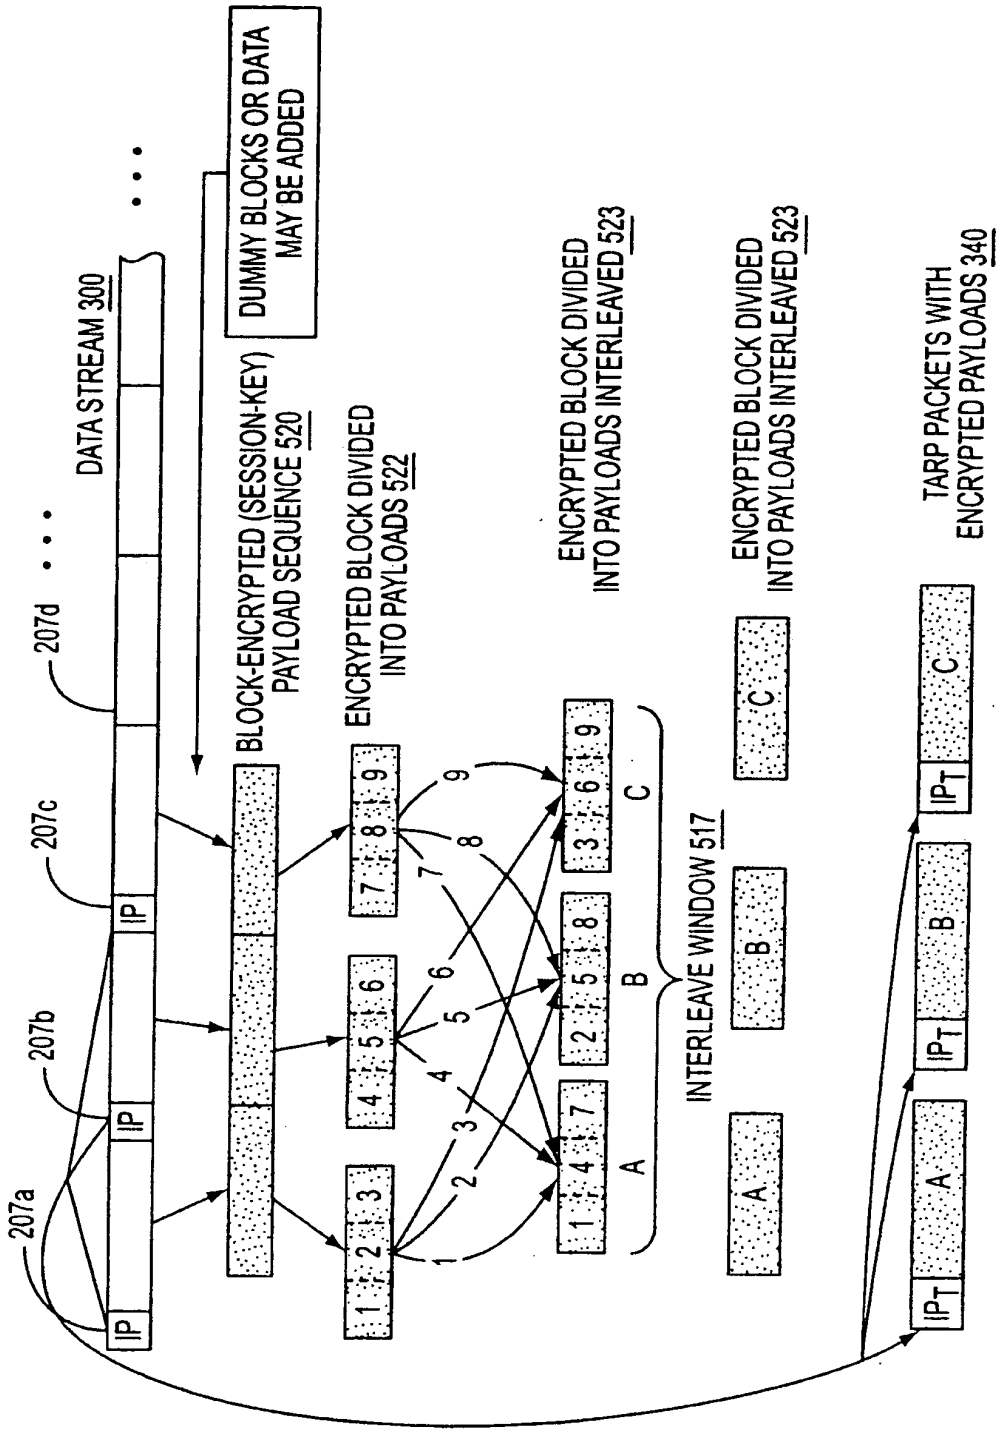


FIG. 3B



New Bay Capital, LLC
Ex.1015-Page 2485 of 3151

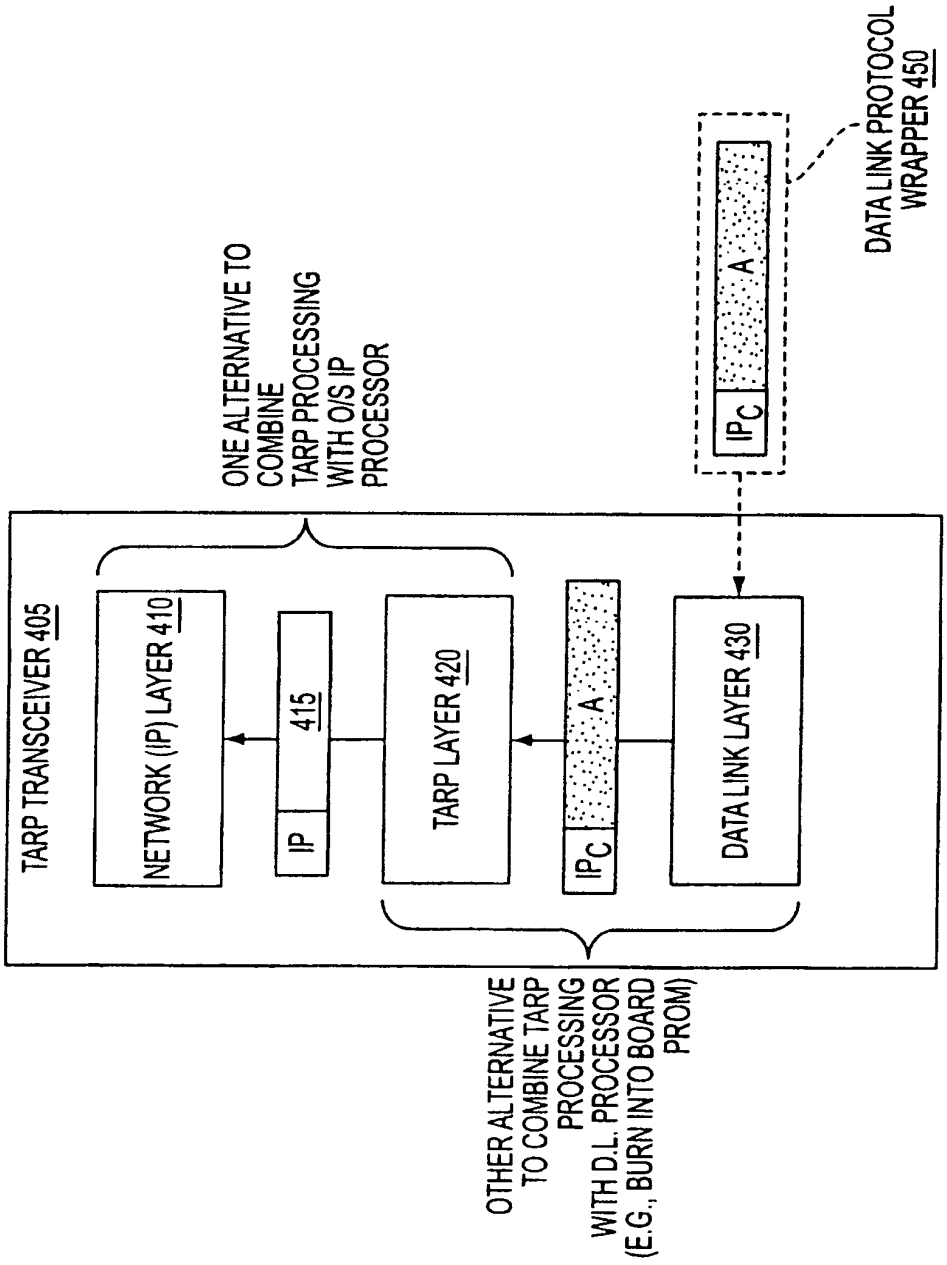


FIG. 4



New Bay Capital, LLC
Ex.1015-Page 2486 of 3151

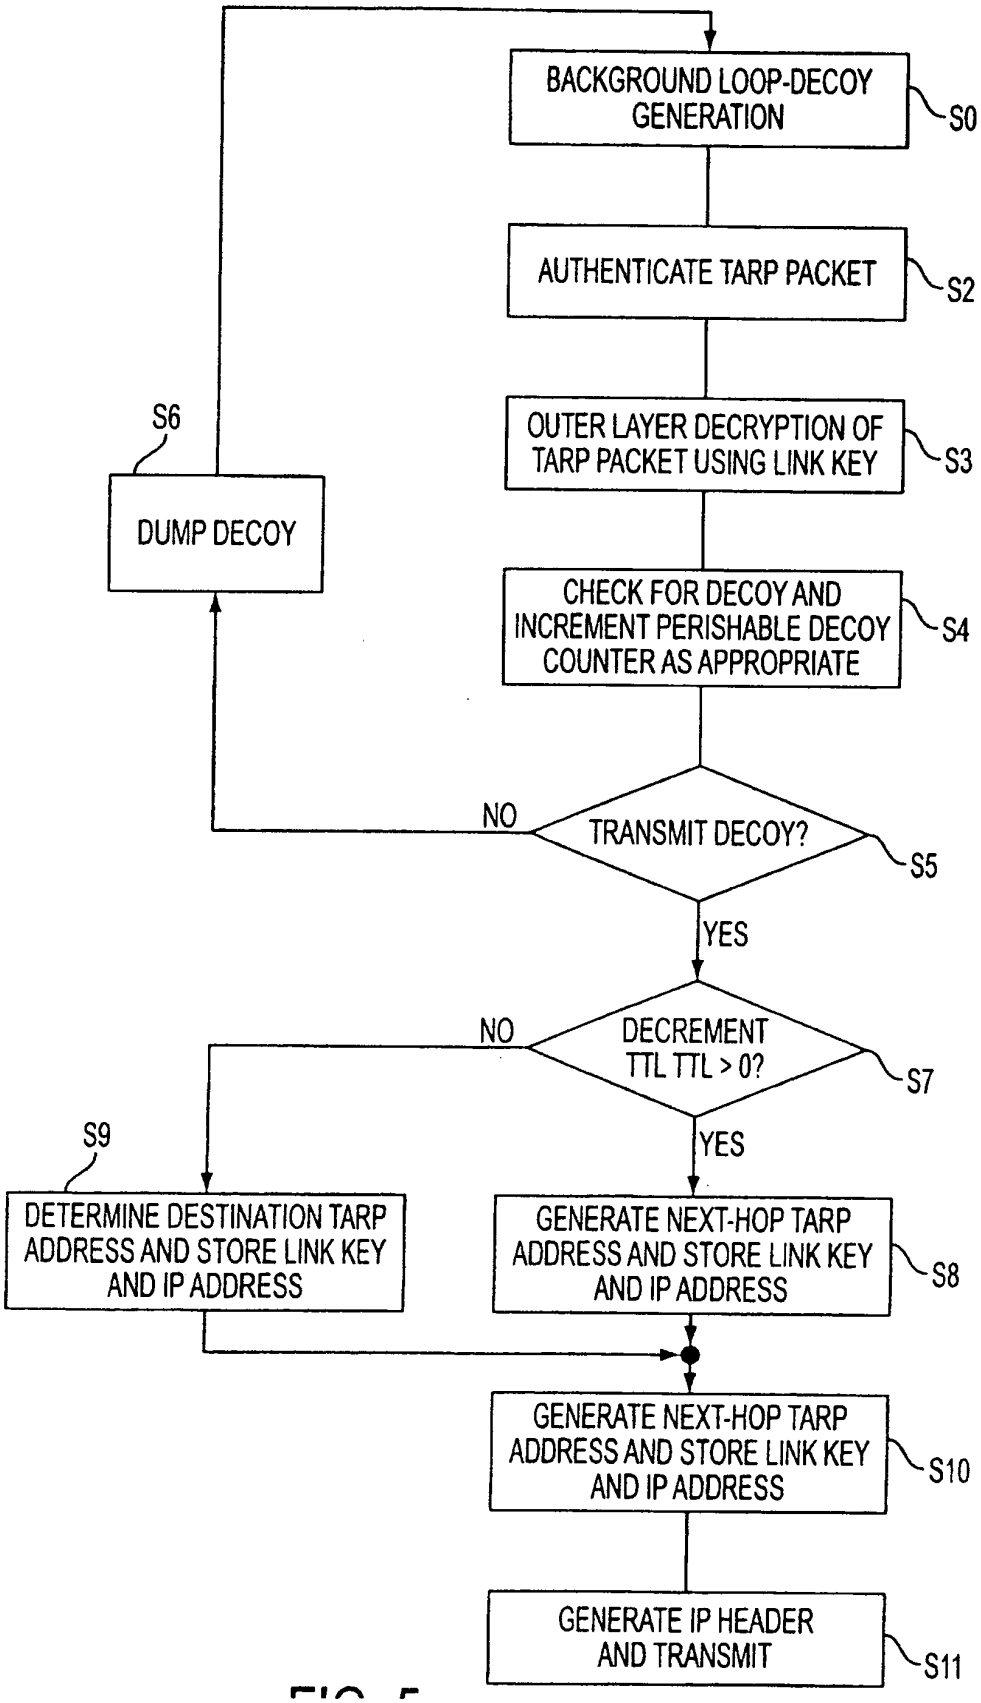


FIG. 5

New Bay Capital, LLC
Ex.1015-Page 2487 of 3151

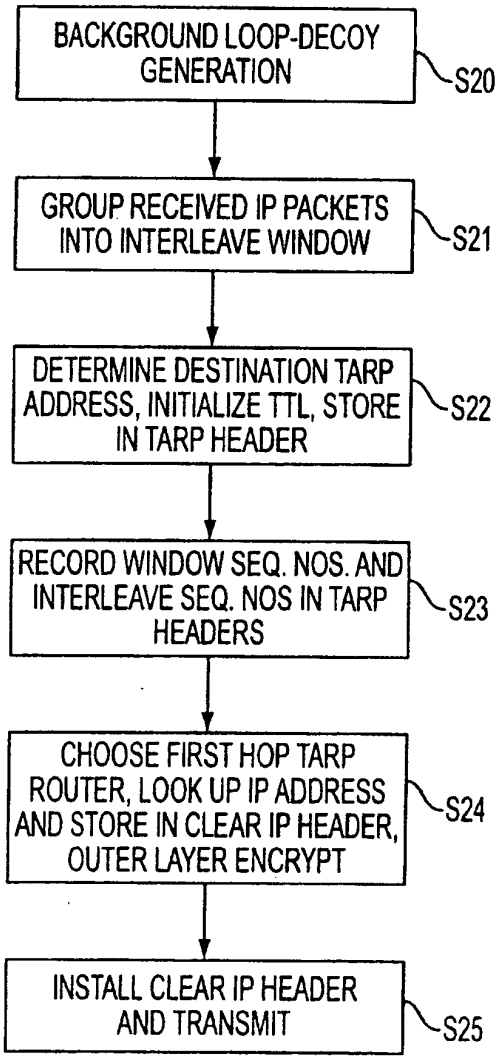


FIG. 6



New Bay Capital, LLC
Ex.1015-Page 2488 of 3151

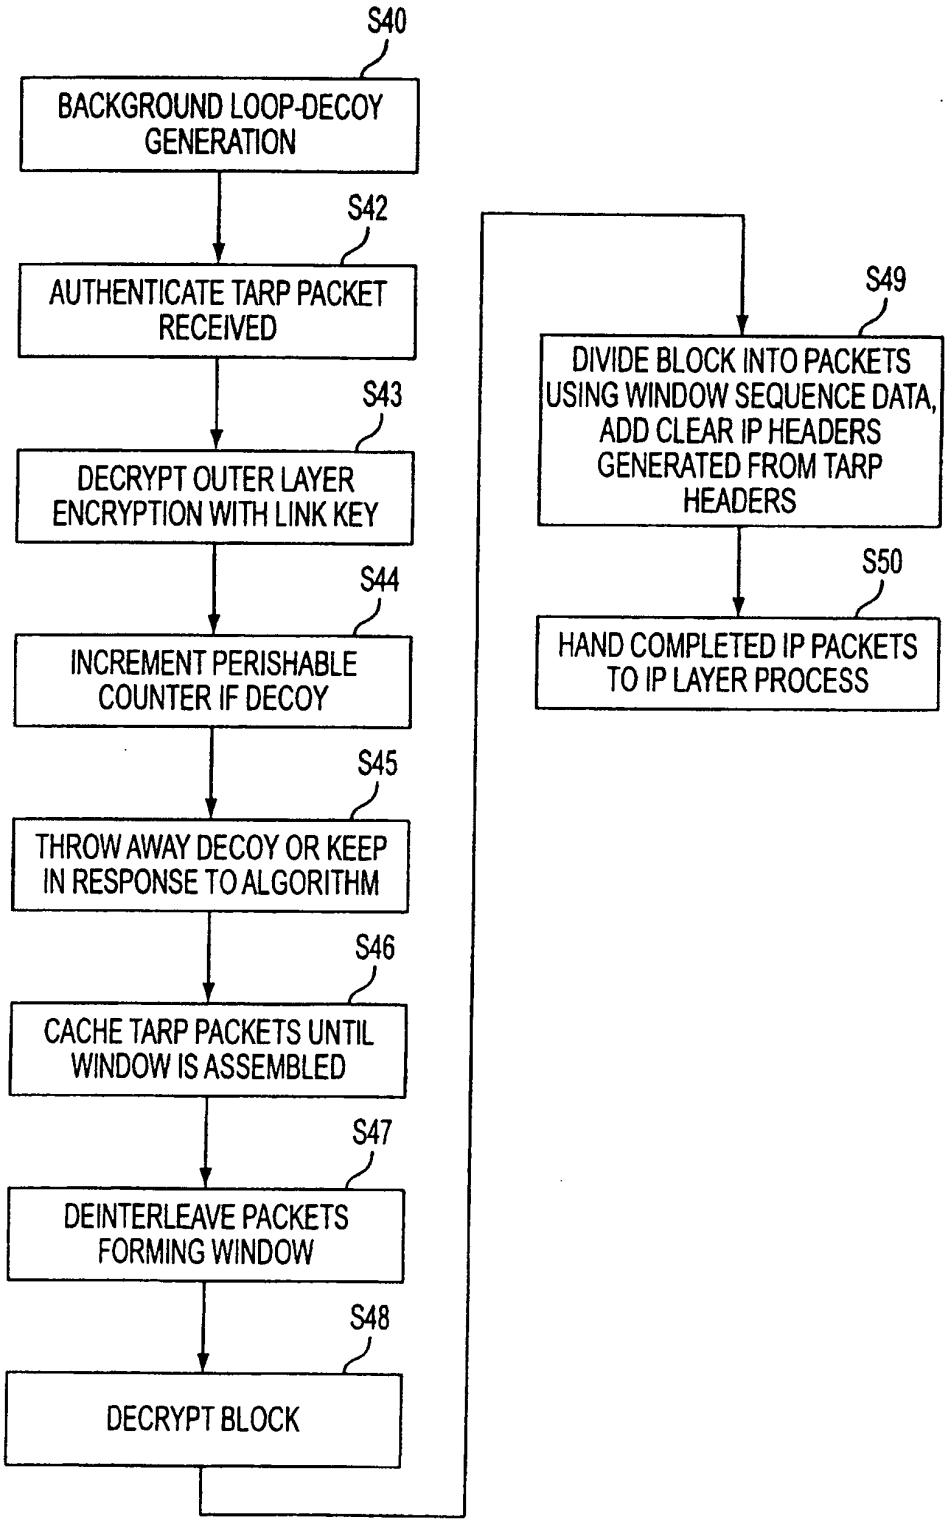


FIG. 7



New Bay Capital, LLC
Ex.1015-Page 2489 of 3151

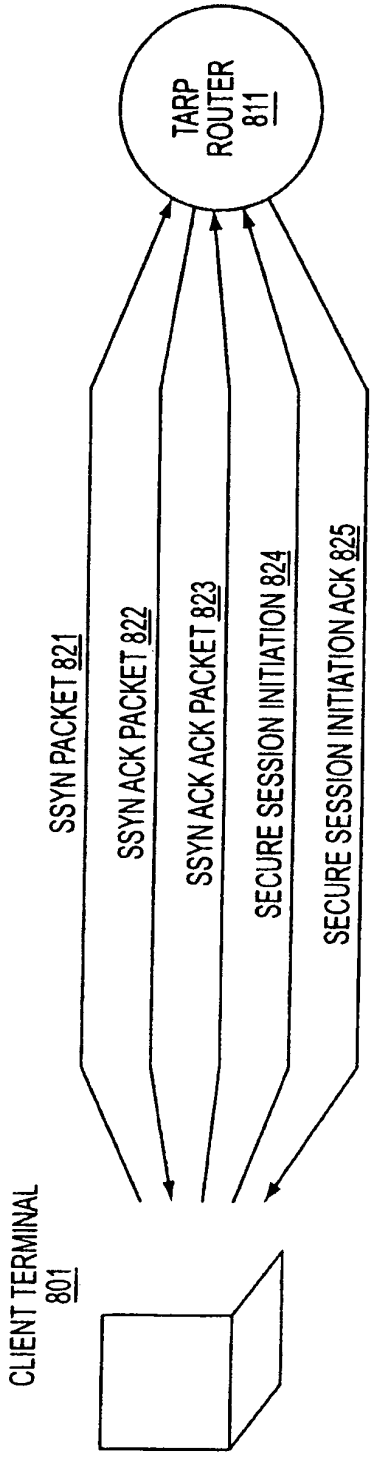


FIG. 8



New Bay Capital, LLC
Ex.1015-Page 2490 of 3151

FIG. 9



New Bay Capital, LLC
Ex.1015-Page 2491 of 3151

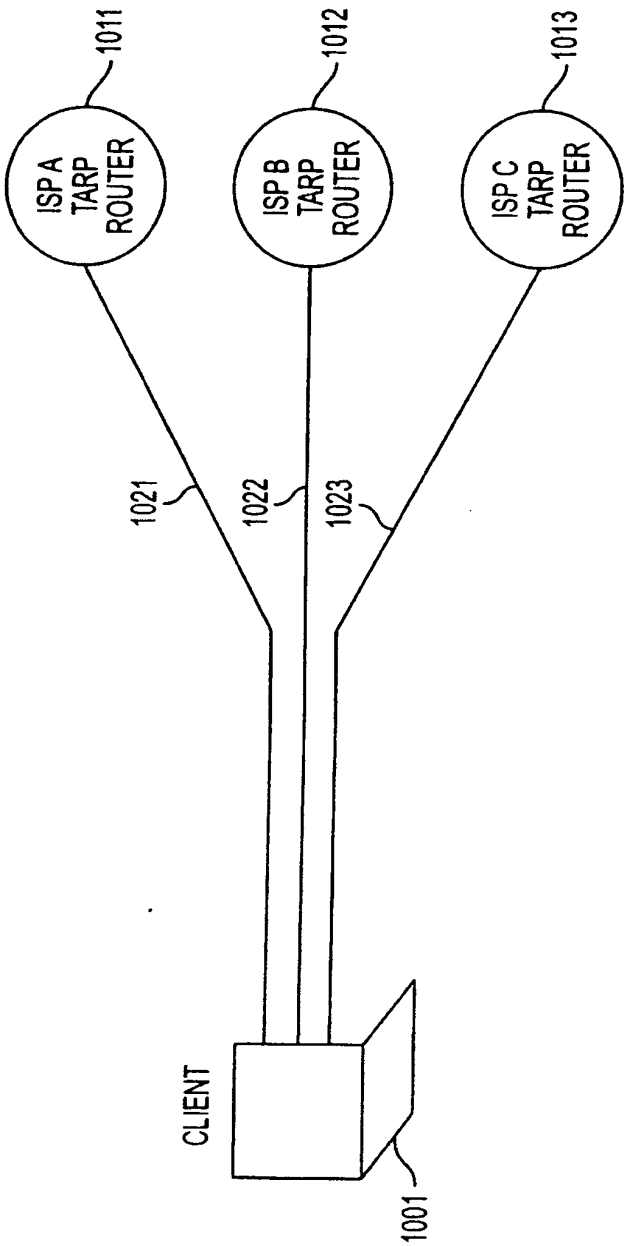


FIG. 10



New Bay Capital, LLC
Ex.1015-Page 2492 of 3151

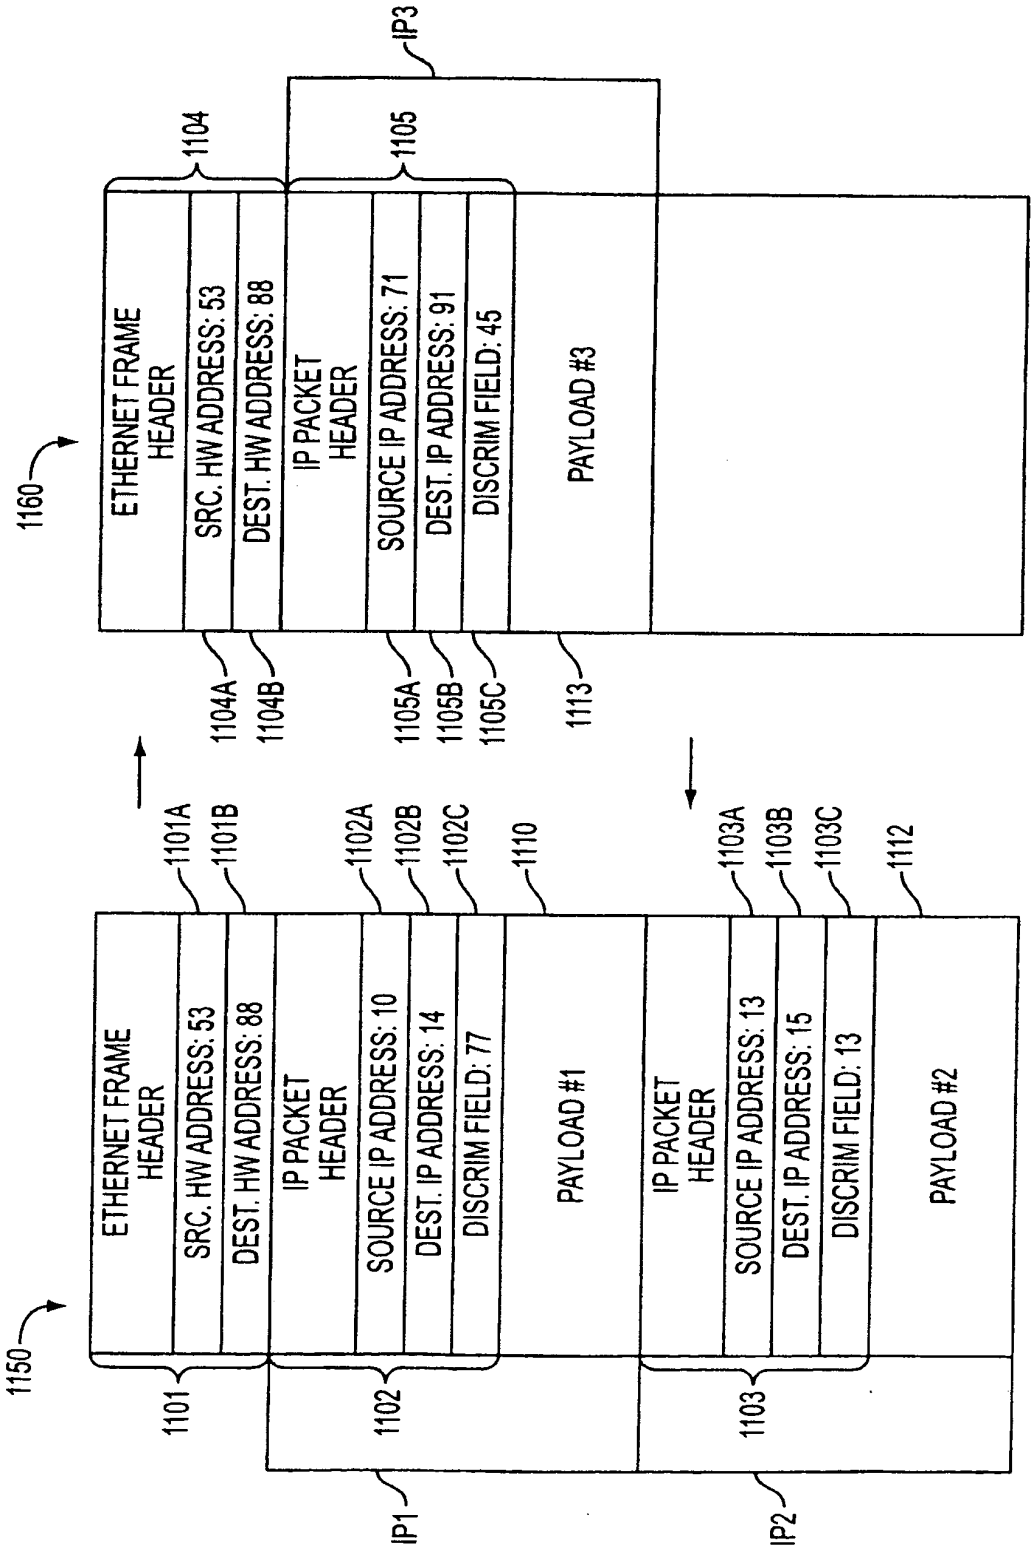
1150
1160
ETHERNET FRAME HEADER
SRC. HW ADDRESS: 53
DEST. HW ADDRESS: 88
IP PACKET HEADER
SOURCE IP ADDRESS: 10
DEST. IP ADDRESS: 14
DISCRIM FIELD: 77
PAYLOAD #1
IP PACKET HEADER
SOURCE IP ADDRESS: 13
DEST. IP ADDRESS: 15
DISCRIM FIELD: 13
PAYLOAD #2
ETHERNET FRAME HEADER
SRC. HW ADDRESS: 53
DEST. HW ADDRESS: 88
IP PACKET HEADER
SOURCE IP ADDRESS: 71
DEST. IP ADDRESS: 91
DISCRIM FIELD: 45
PAYLOAD #3
1101
1102
1103
1104
1105
1101A
1101B
1102A
1102B
1102C
1110
1103A
1103B
1103C
1112
1104A
1104B
1105A
1105B
1105C
1113
IP1
IP2
IP3

FIG. 11

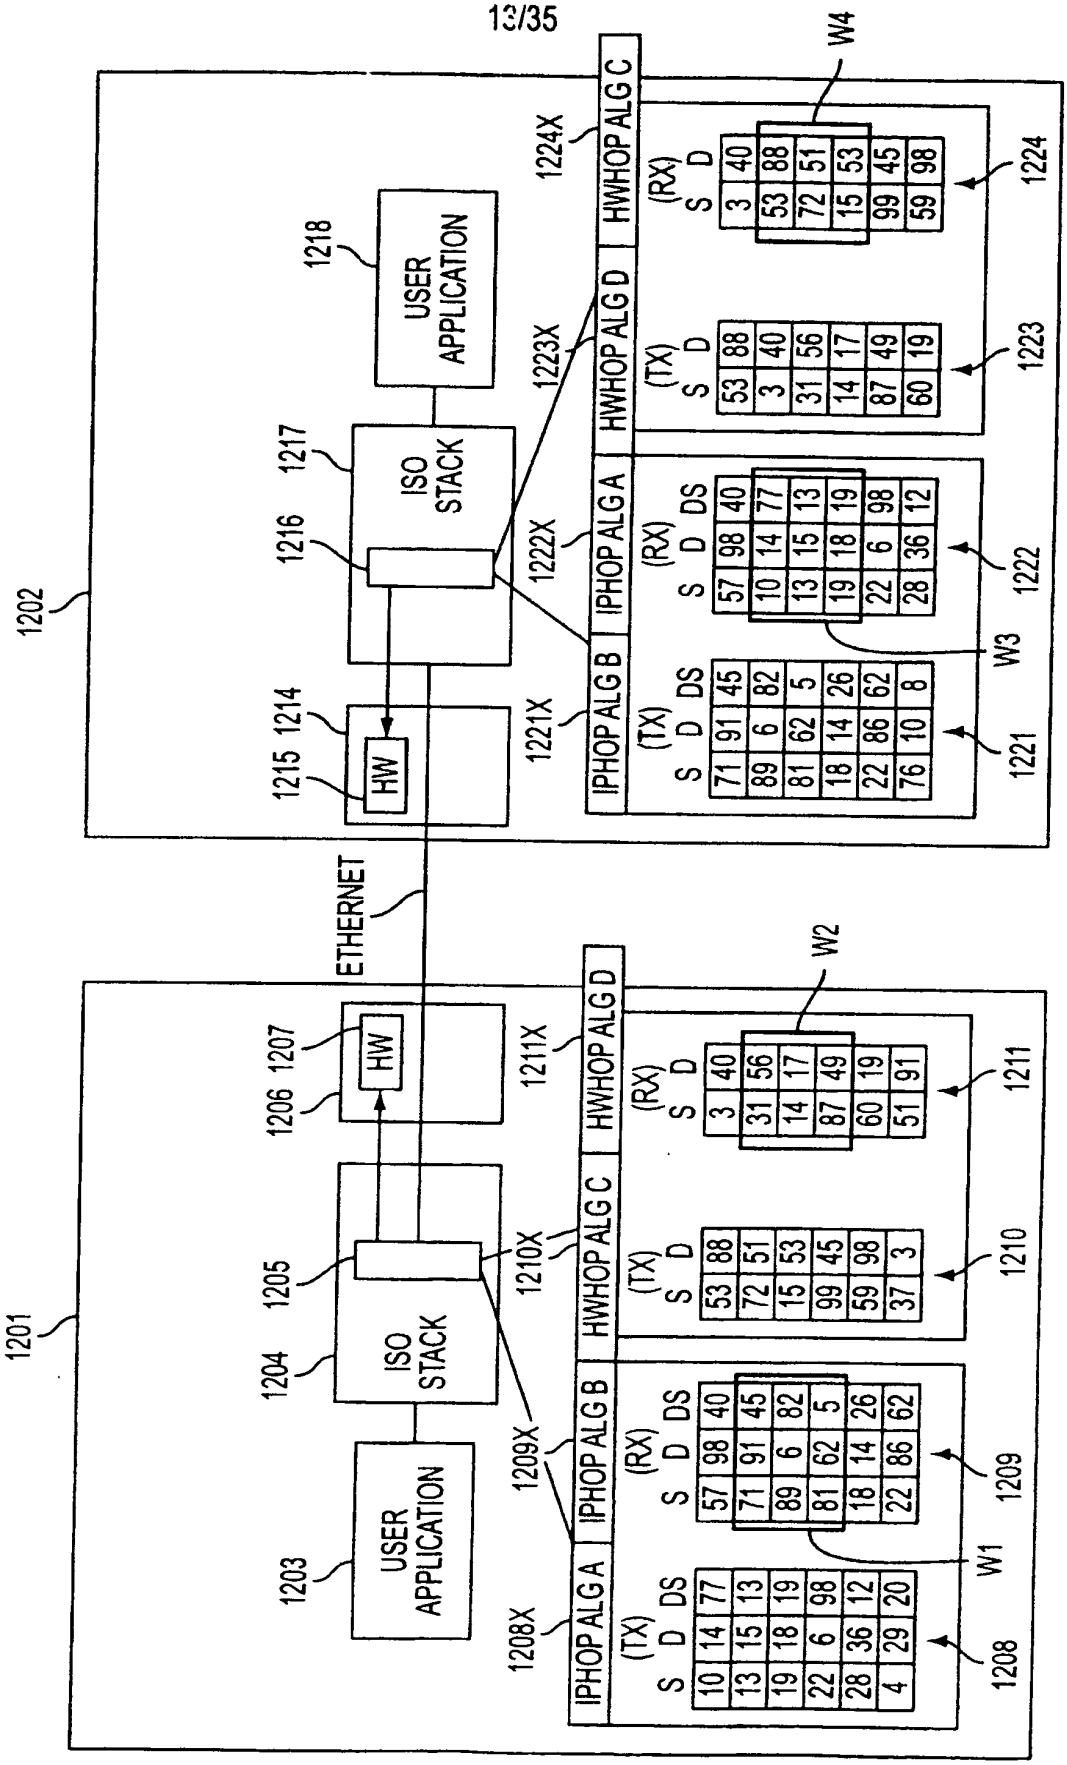
1201
1202
ETHERNET
1203
USER APPLICATION
1204
ISO STACK
1205
1206
1207
HW
1208
1208X
IPHOP ALG B
IPHOP ALG A
1209
1209X
1210
1210X
HWHOP ALG C
1211
1211X
HWHOP ALG D
1214
1215
1216
1217
1218
1221
1221X
1222
1222X
1223
1223X
1224
1224X
W1
W2
W3
W4
(TX)
(RX)
S
D
DS

FIG. 12A

| MODE OR EMBODIMENT | HARDWARE ADDRESSES | IP ADDRESSES | DISCRIMINATOR FIELD VALUES |
|---|---|---|---|
| 1. PROMISCUOUS | SAME FOR ALL NODES OR COMPLETELY RANDOM | CAN BE VARIED IN SYNC | CAN BE VARIED IN SYNC |
| 2. PROMISCUOUS PER VPN | FIXED FOR EACH VPN | CAN BE VARIED IN SYNC | CAN BE VARIED IN SYNC |
| 3. HARDWARE HOPPING | CAN BE VARIED IN SYNC | CAN BE VARIED IN SYNC | CAN BE VARIED IN SYNC |

FIG. 12B



New Bay Capital, LLC
Ex.1015-Page 2495 of 3151

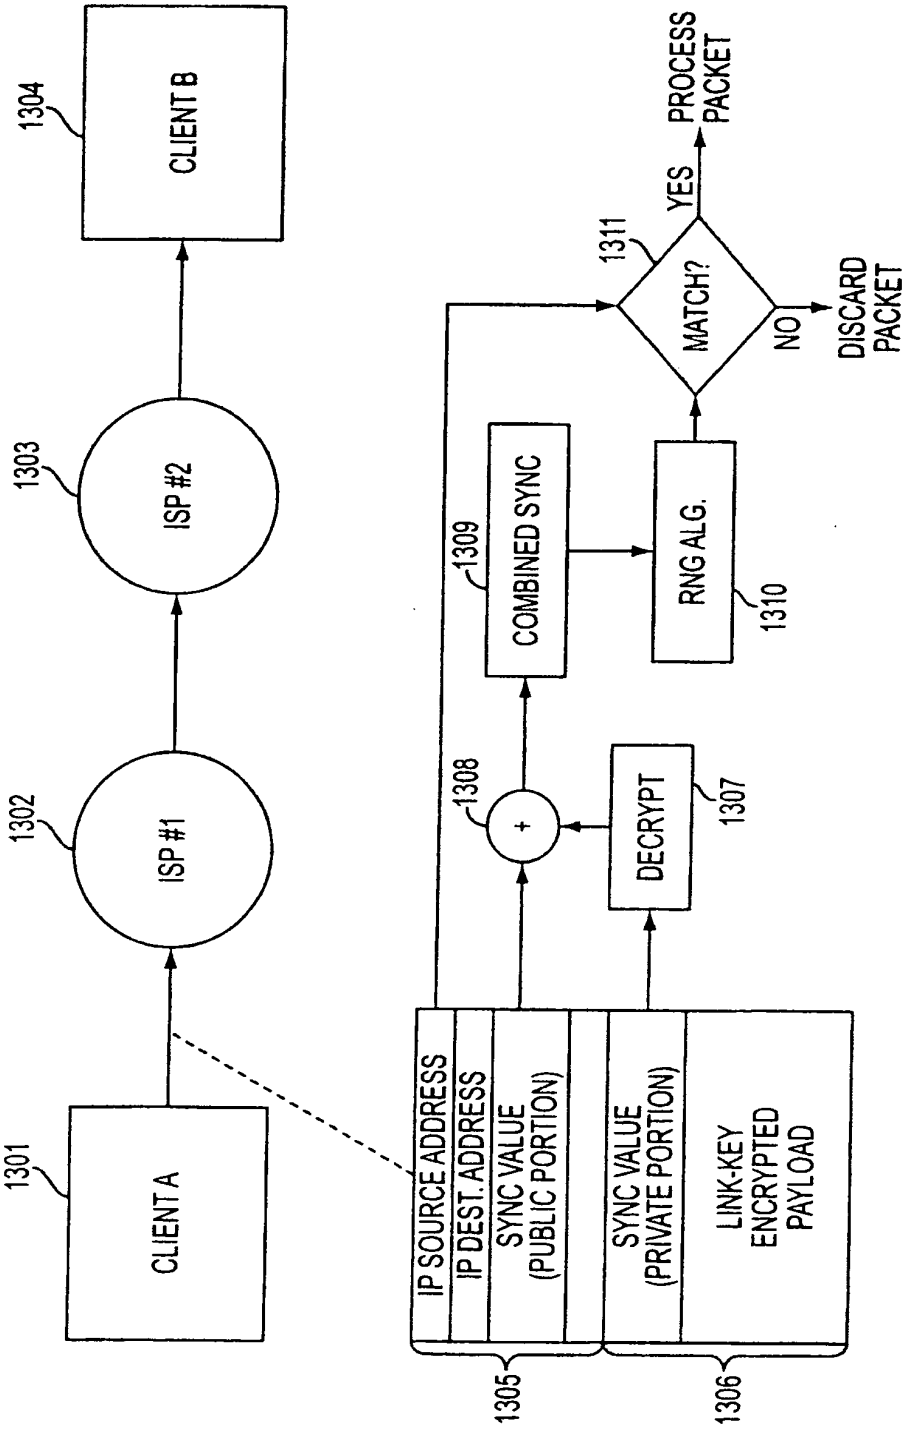


FIG. 13



New Bay Capital, LLC
Ex.1015-Page 2496 of 3151

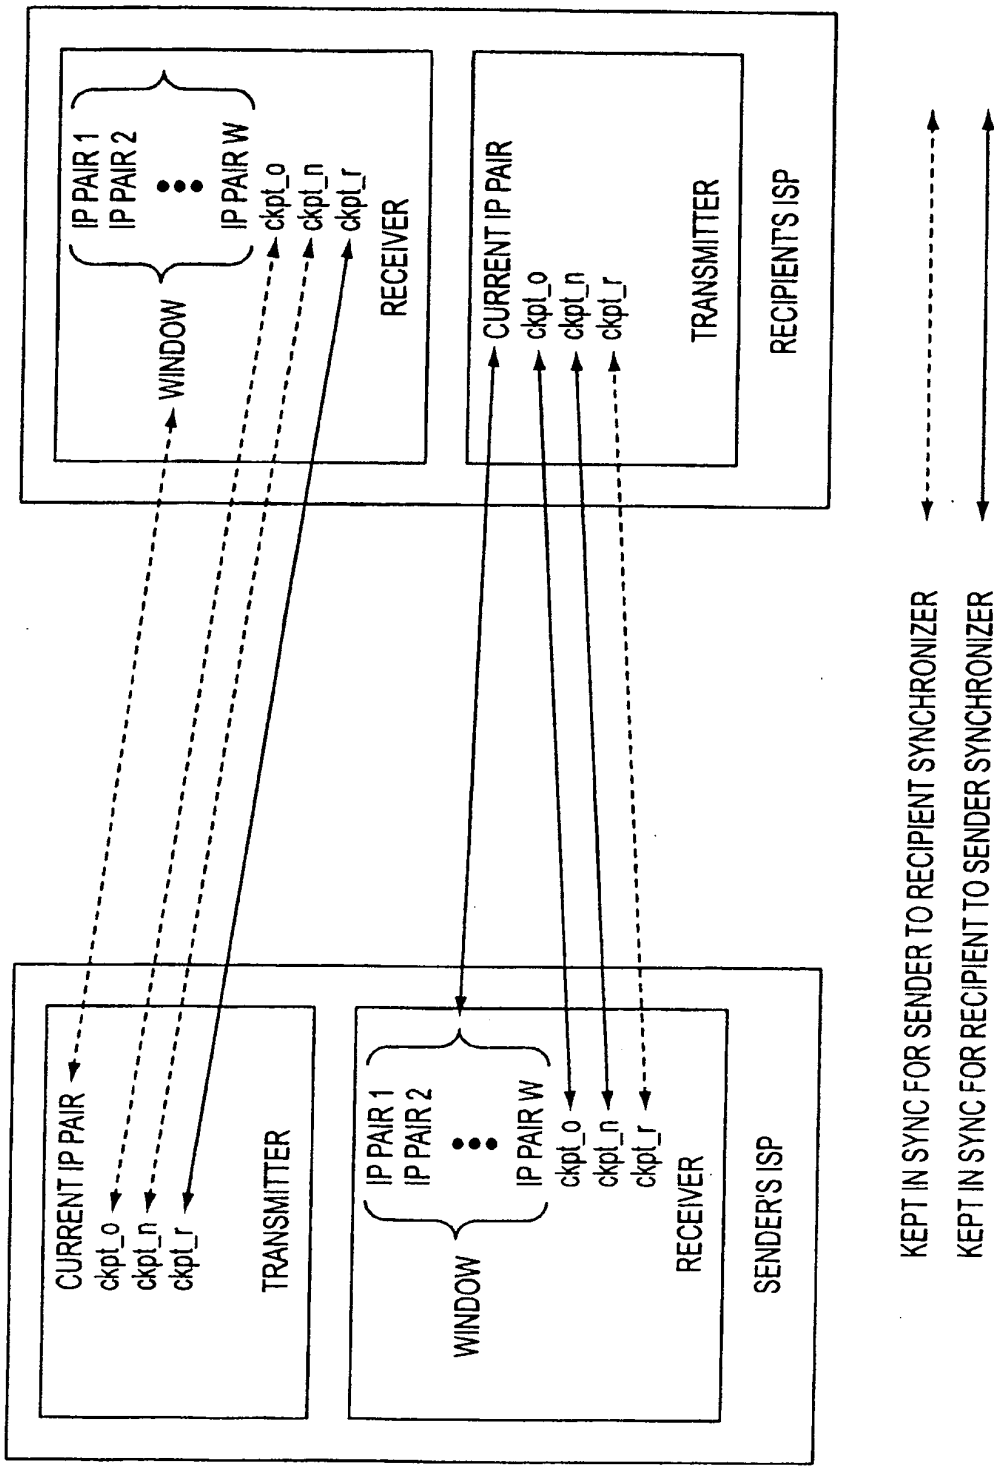


FIG. 14



New Bay Capital, LLC
Ex.1015-Page 2497 of 3151

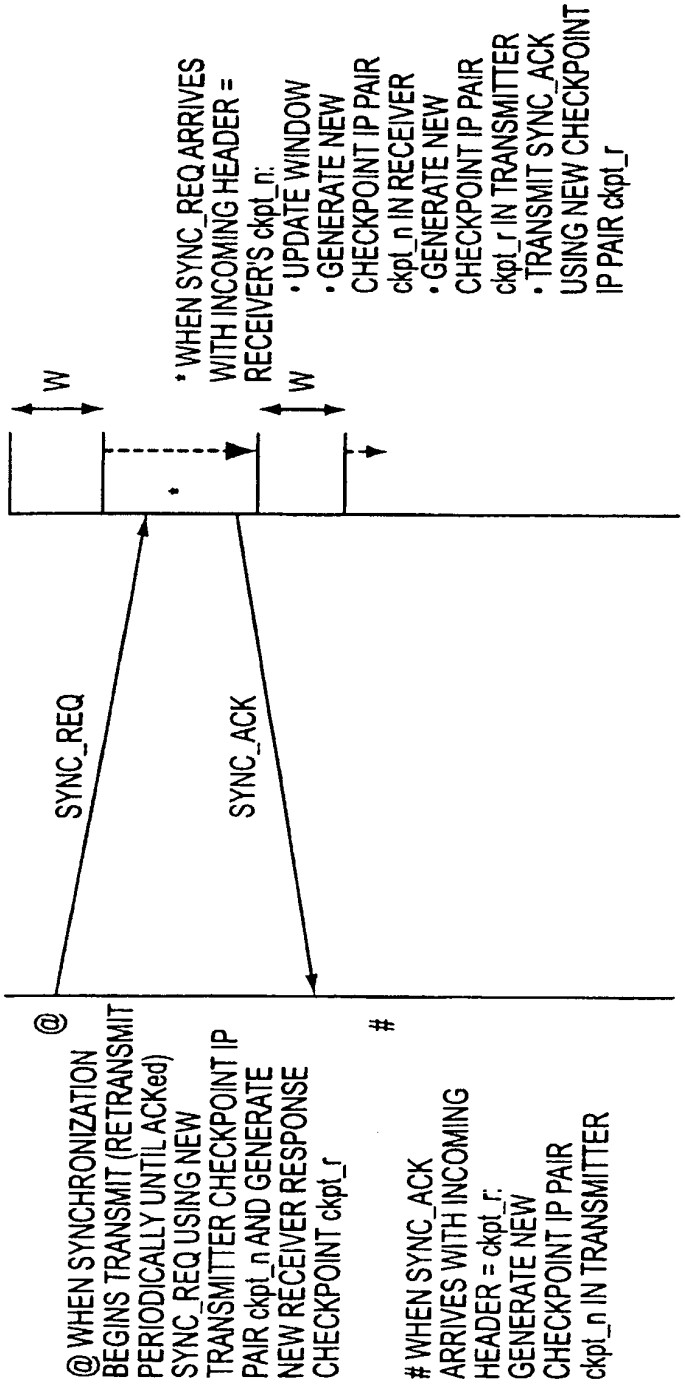


FIG. 15



New Bay Capital, LLC
Ex.1015-Page 2498 of 3151

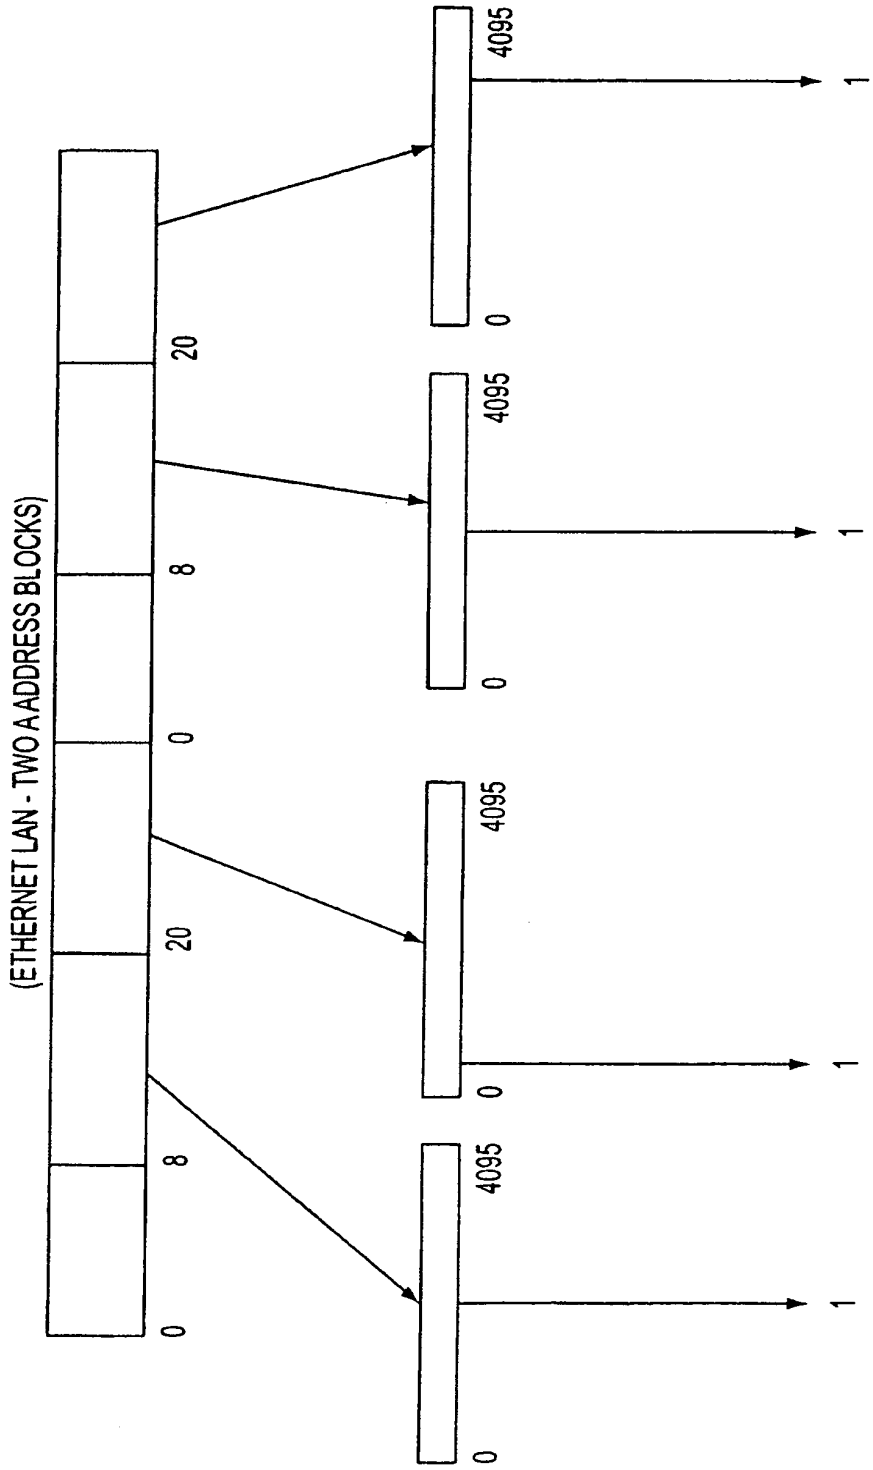


FIG. 16



New Bay Capital, LLC
Ex.1015-Page 2499 of 3151

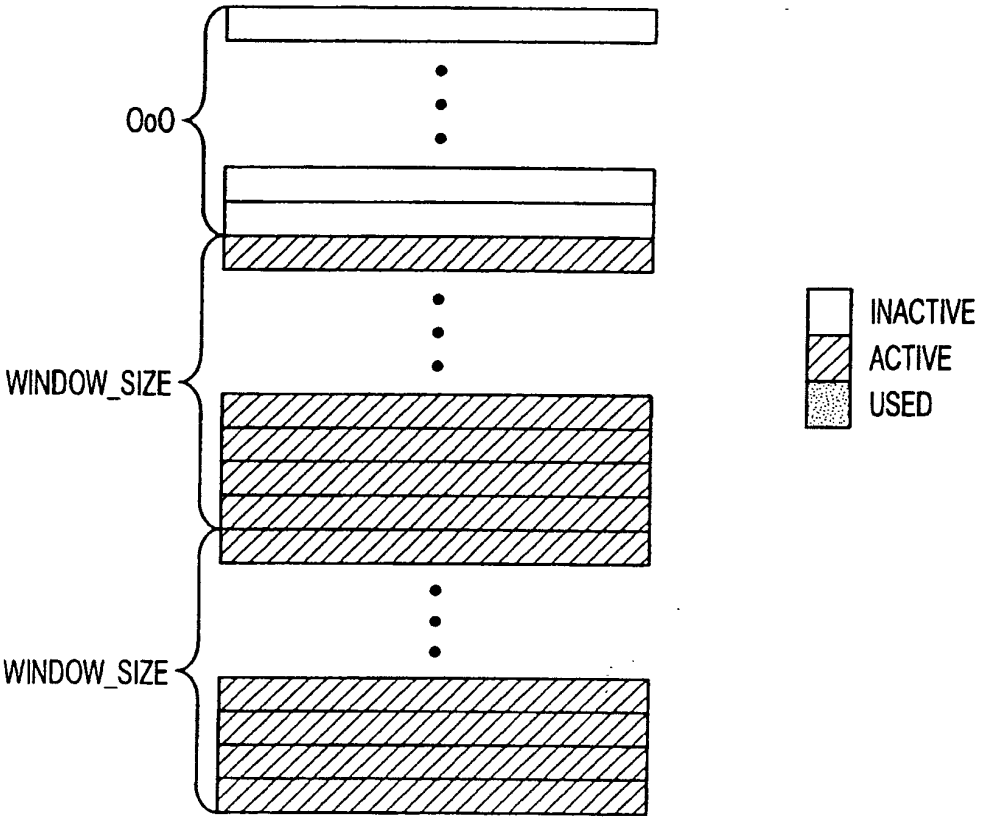


FIG. 17



New Bay Capital, LLC
Ex.1015-Page 2500 of 3151

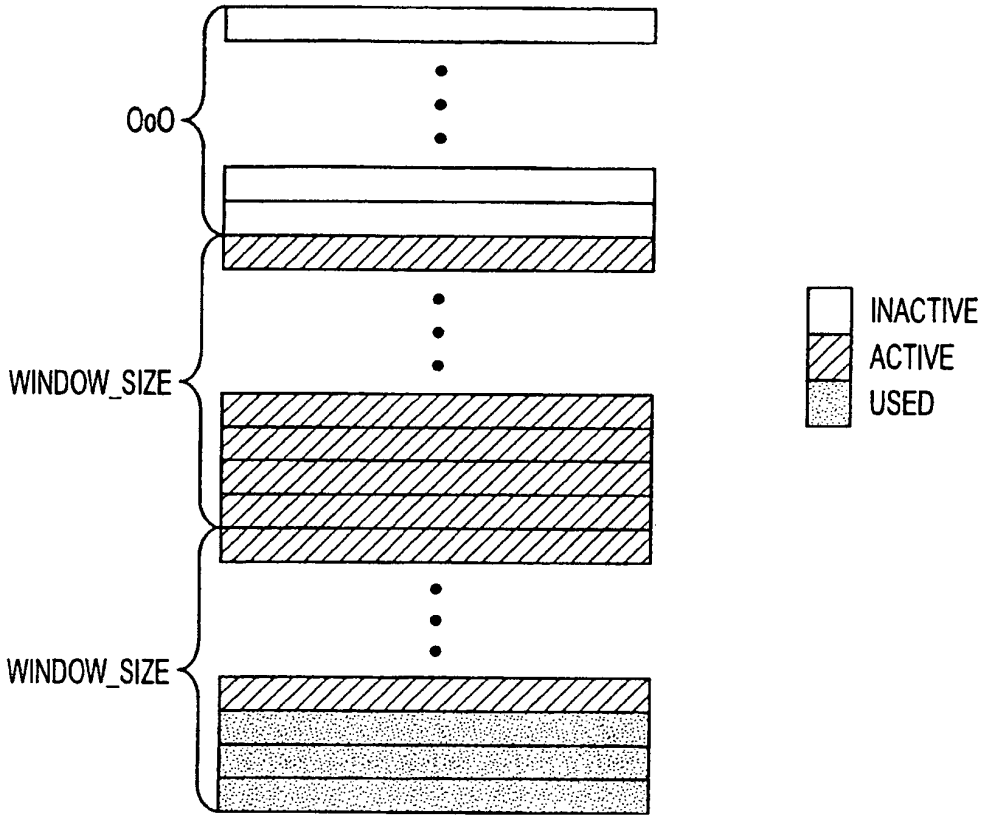


FIG. 18

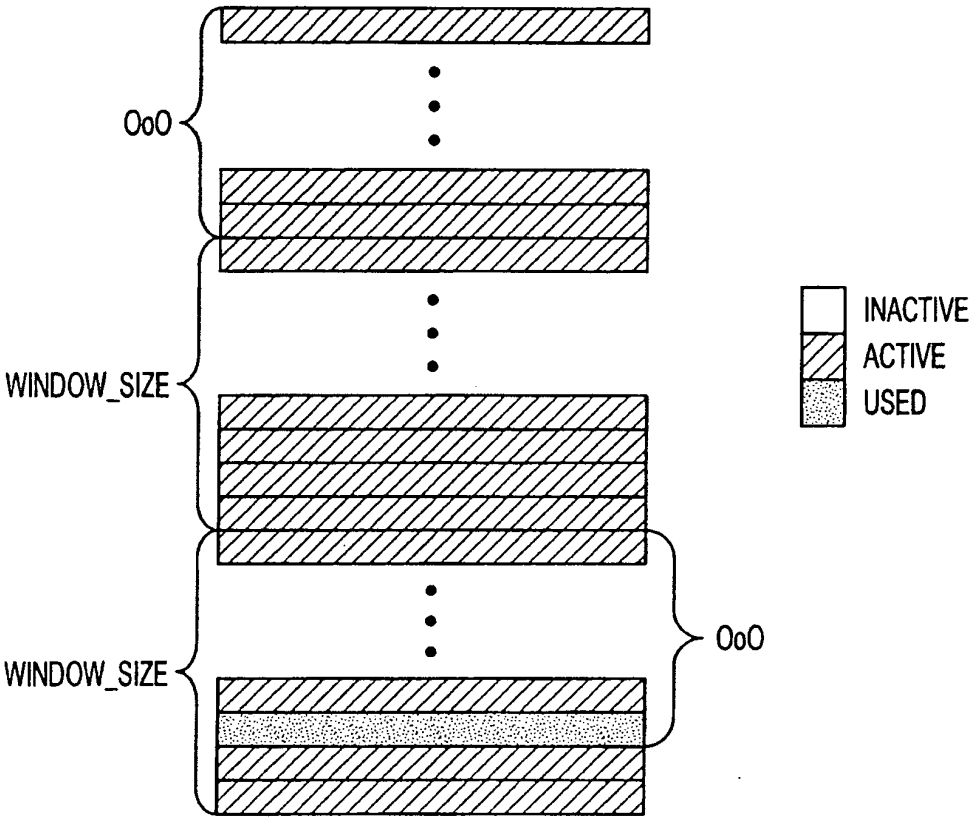


FIG. 19

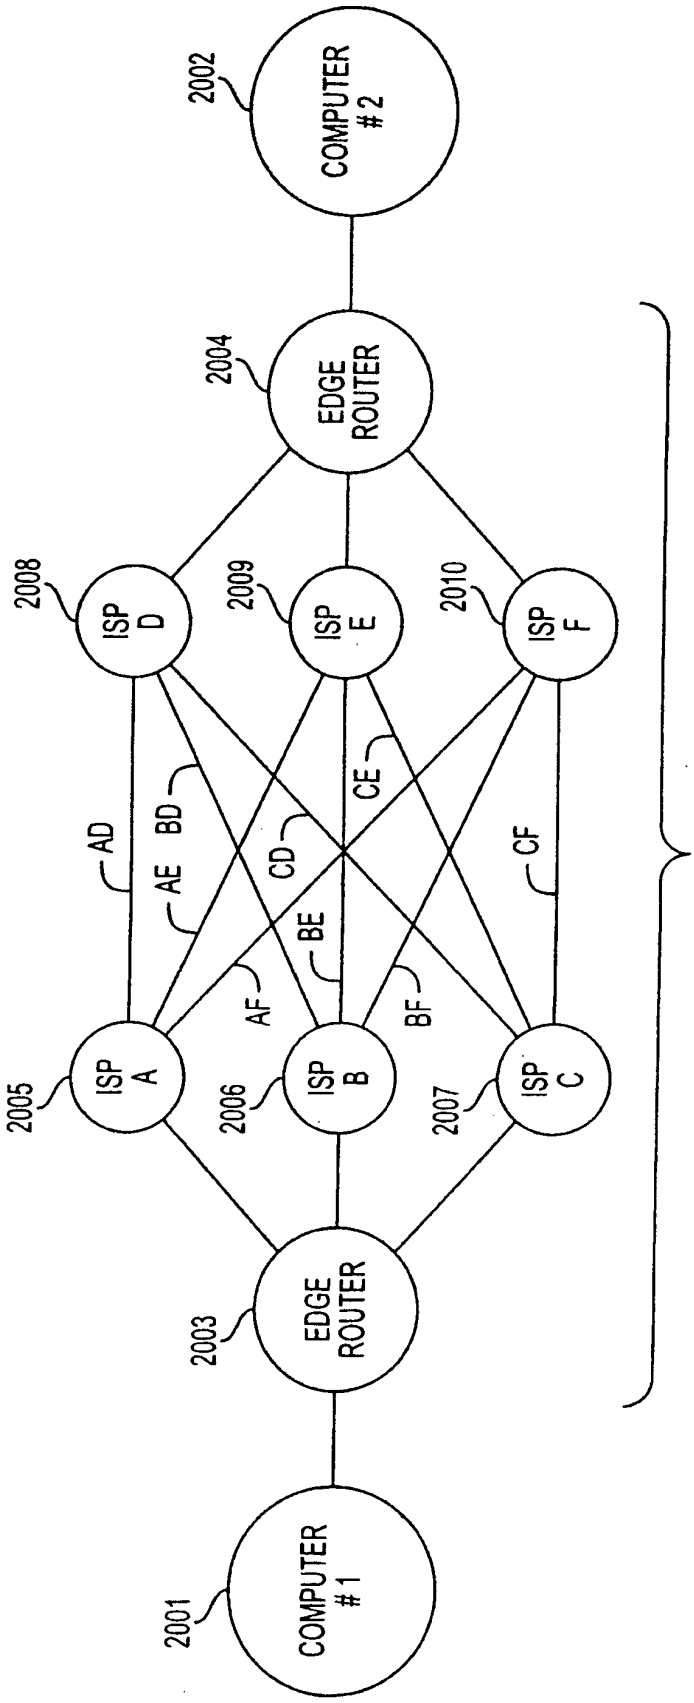
2002
COMPUTER #2
2004
EDGE ROUTER
2008
ISP D
2009
ISP E
2010
ISP F
AD
AE
BD
AF
BE
CD
CE
BF
CF
2005
ISP A
2006
ISP B
2007
ISP C
2003
EDGE ROUTER
2001
COMPUTER #1
2011

FIG. 20

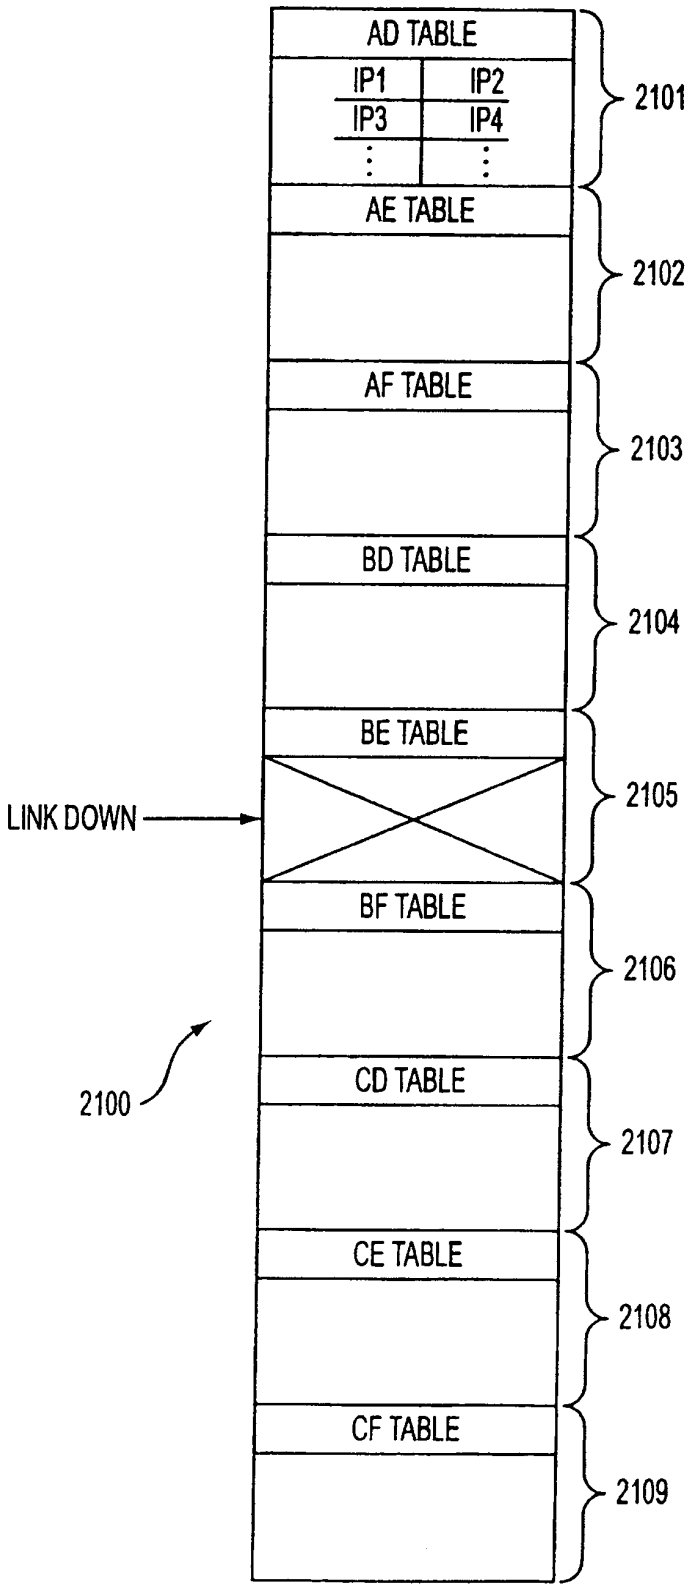
AD TABLE
IP1
IP2
IP3
IP4
2101
AE TABLE
2102
AF TABLE
2103
BD TABLE
2104
BE TABLE
LINK DOWN
2105
BF TABLE
2106
2100
CD TABLE
2107
CE TABLE
2108
CF TABLE
2109

FIG. 21

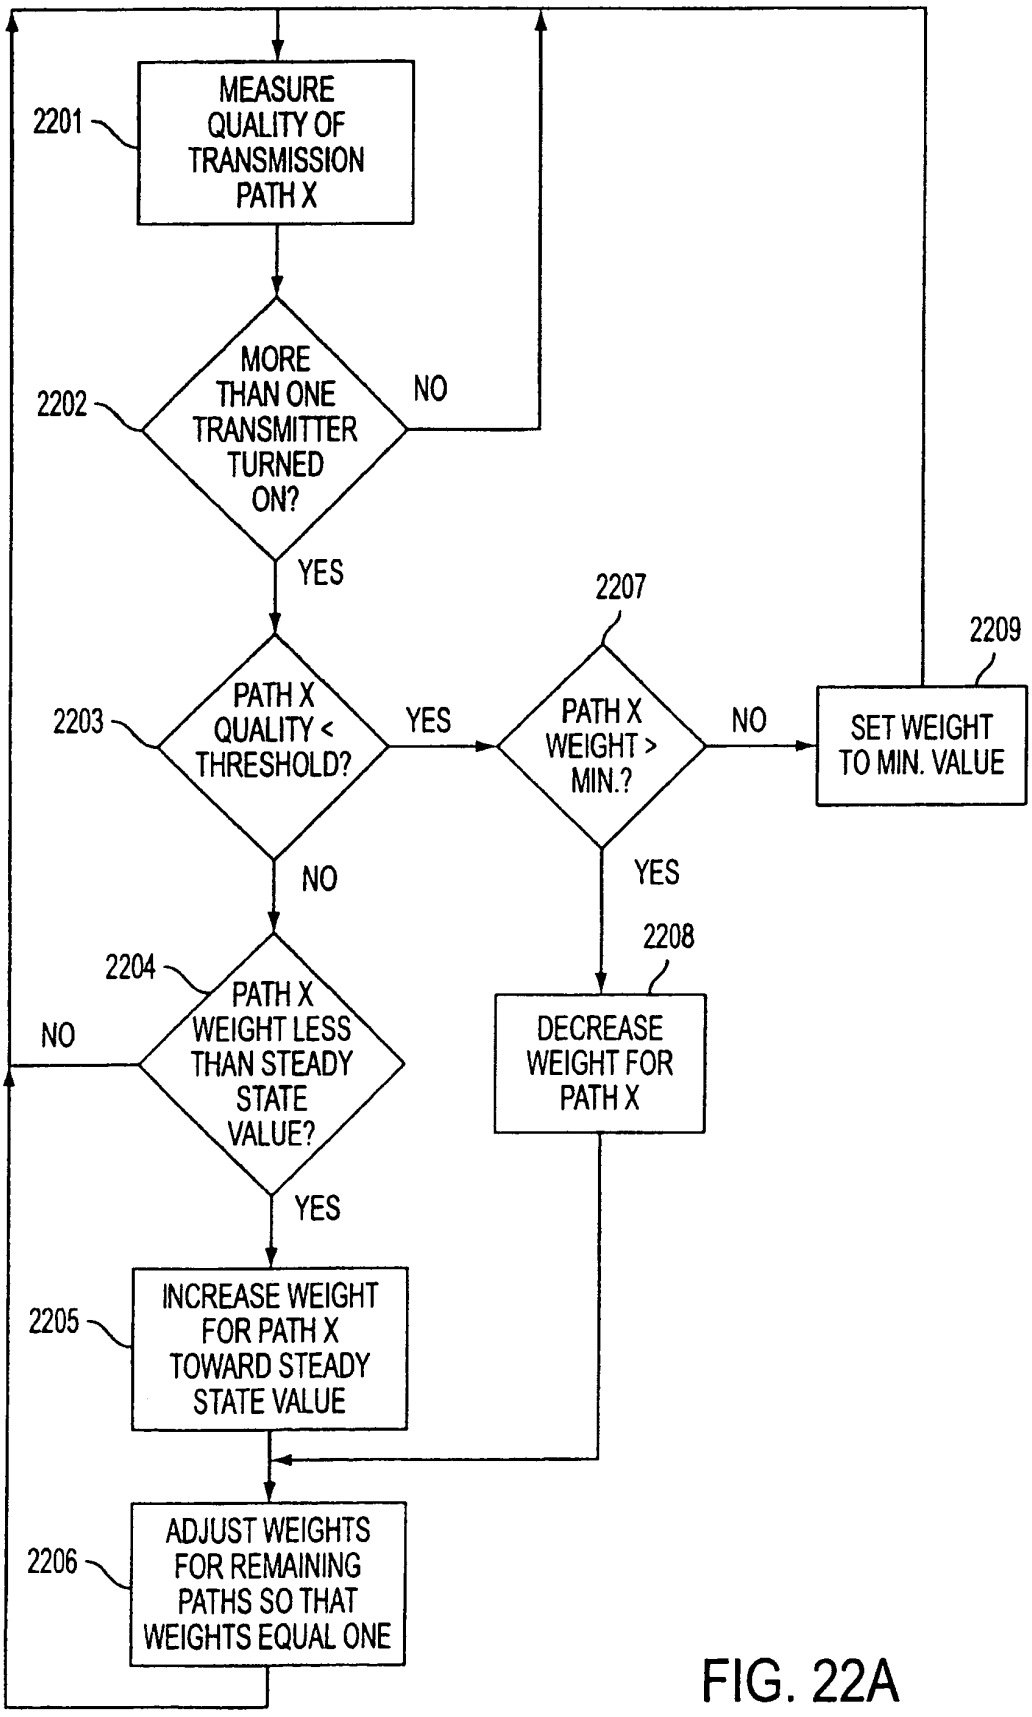


FIG. 22A

New Bay Capital, LLC
Ex.1015-Page 2505 of 3151

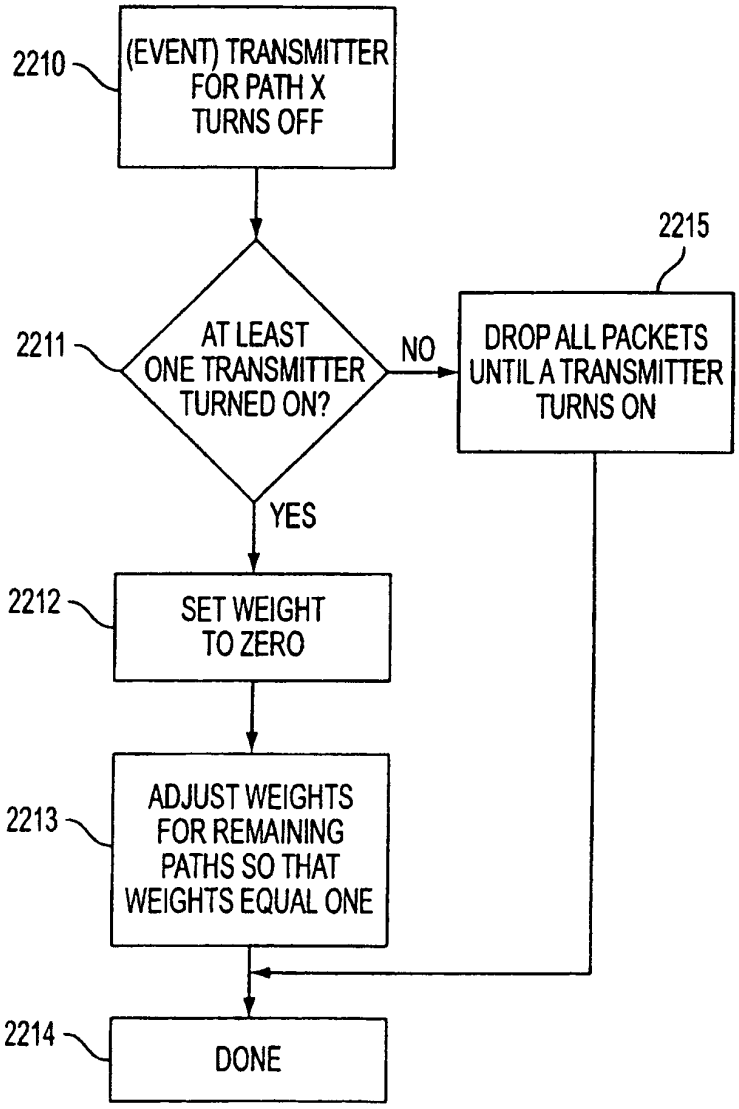
2210
(EVENT) TRANSMITTER FOR PATH X TURNS OFF
2211
AT LEAST ONE TRANSMITTER TURNED ON?
NO
2215
DROP ALL PACKETS UNTIL A TRANSMITTER TURNS ON
YES
2212
SET WEIGHT TO ZERO
2213
ADJUST WEIGHTS FOR REMAINING PATHS SO THAT WEIGHTS EQUAL ONE
2214
DONE

FIG. 22B

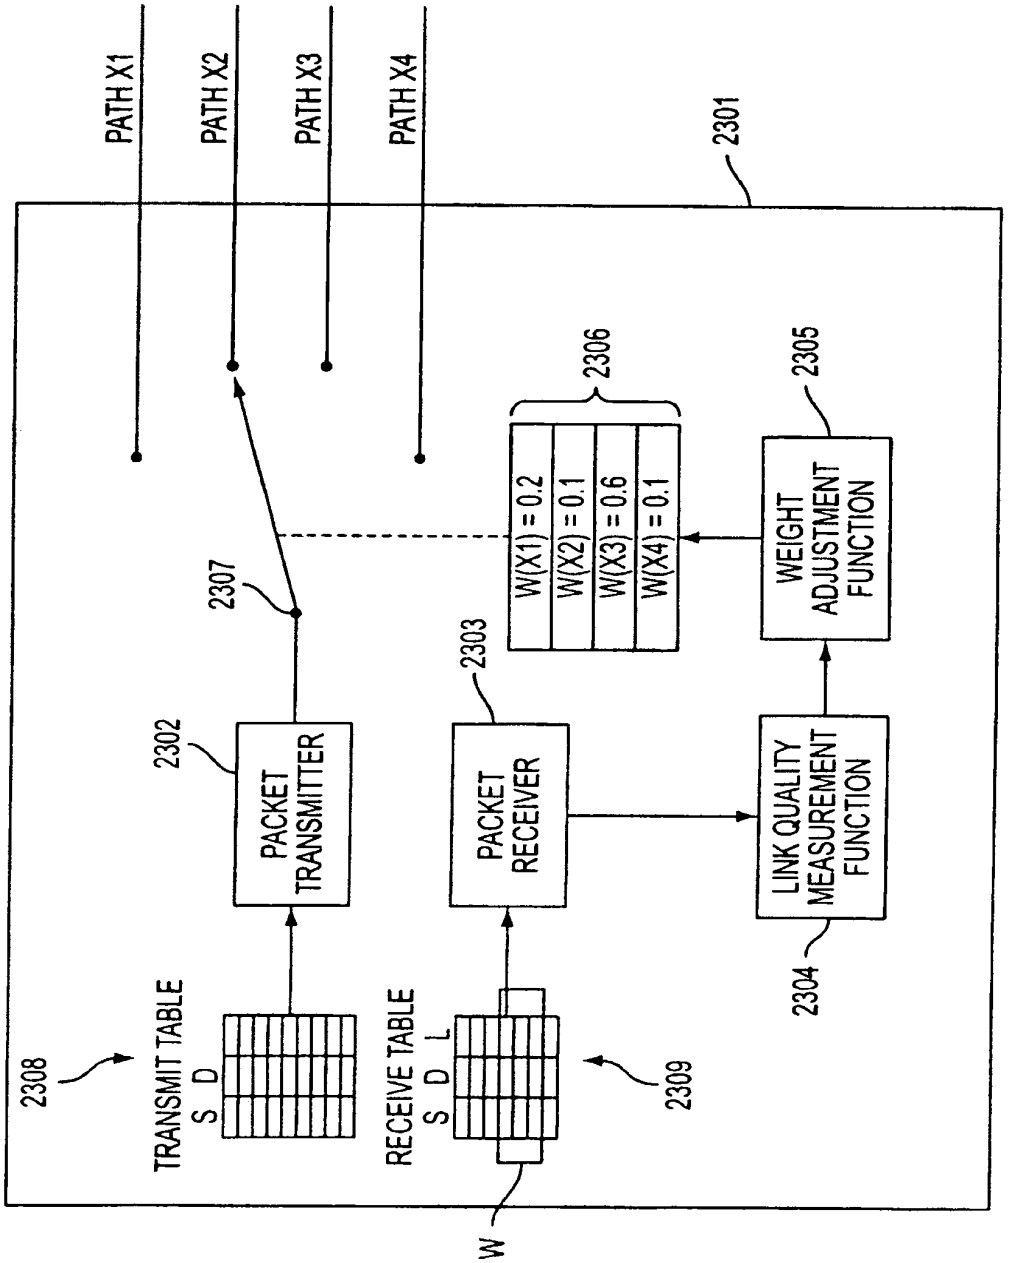


FIG. 23



New Bay Capital, LLC
Ex.1015-Page 2507 of 3151

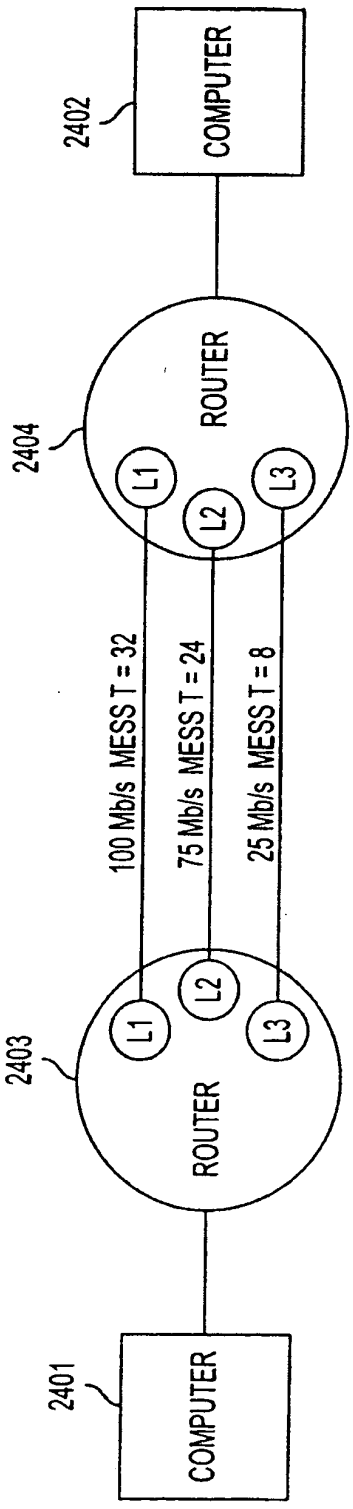
2402
COMPUTER
2404
ROUTER
L1
L2
L3
100 Mb/s MESS T = 32
75 Mb/s MESS T = 24
25 Mb/s MESS T = 8
2403
ROUTER
L1
L2
L3
2401
COMPUTER

FIG. 24

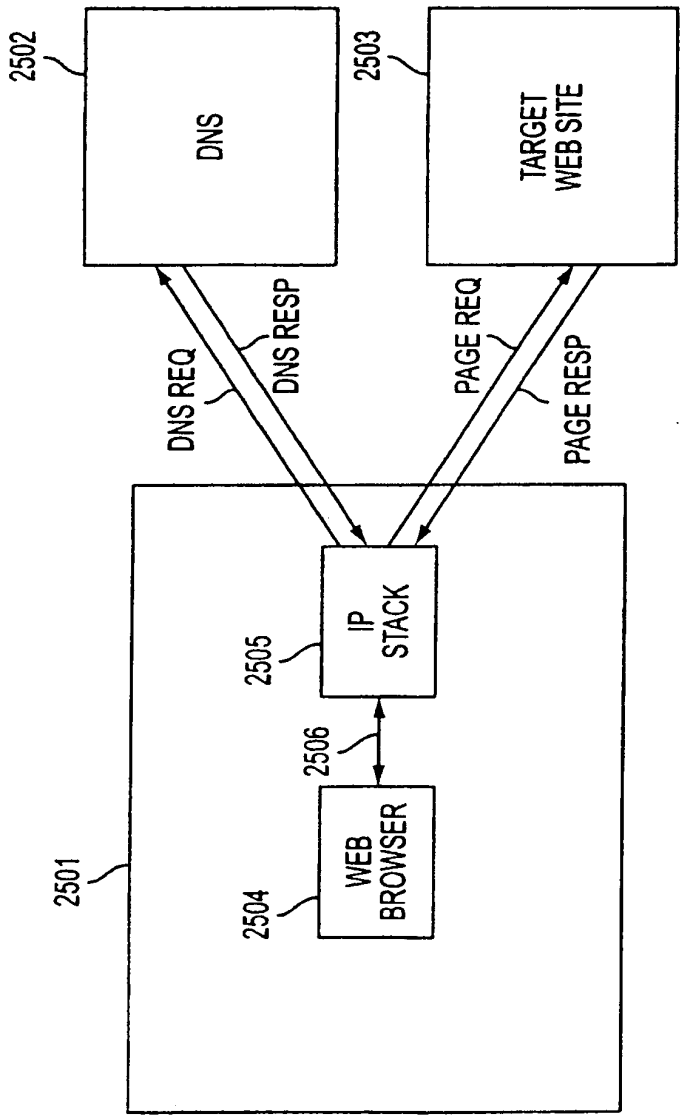


FIG. 25
(PRIOR ART)



New Bay Capital, LLC
Ex.1015-Page 2509 of 3151

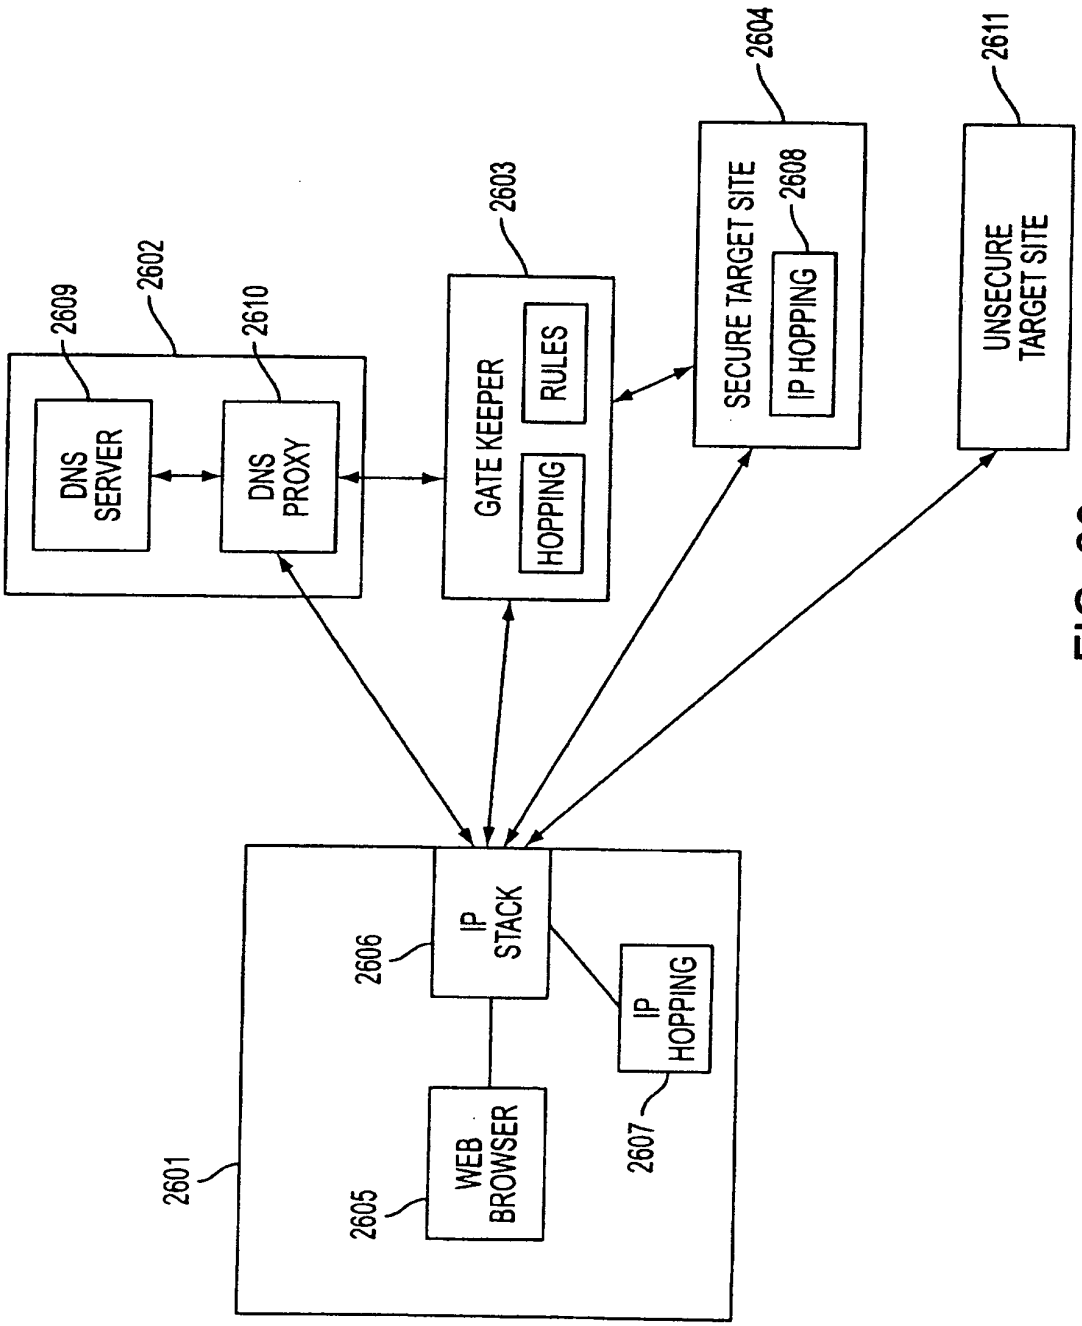


FIG. 26



New Bay Capital, LLC
Ex.1015-Page 2510 of 3151

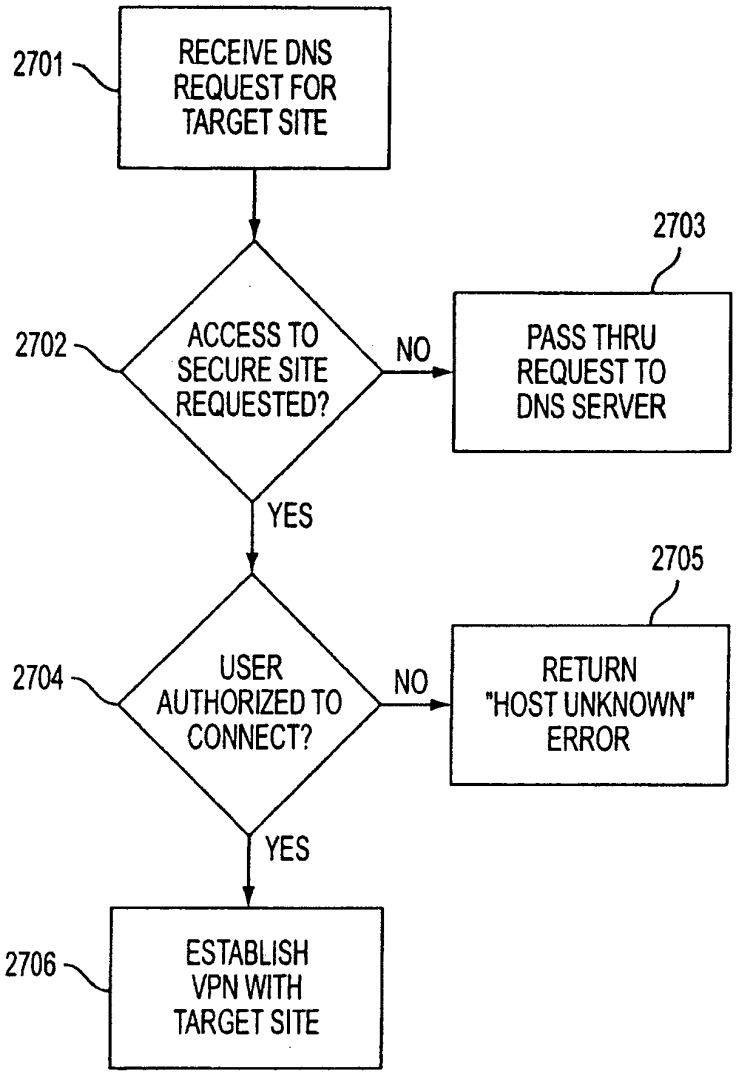


FIG. 27



New Bay Capital, LLC
Ex.1015-Page 2511 of 3151

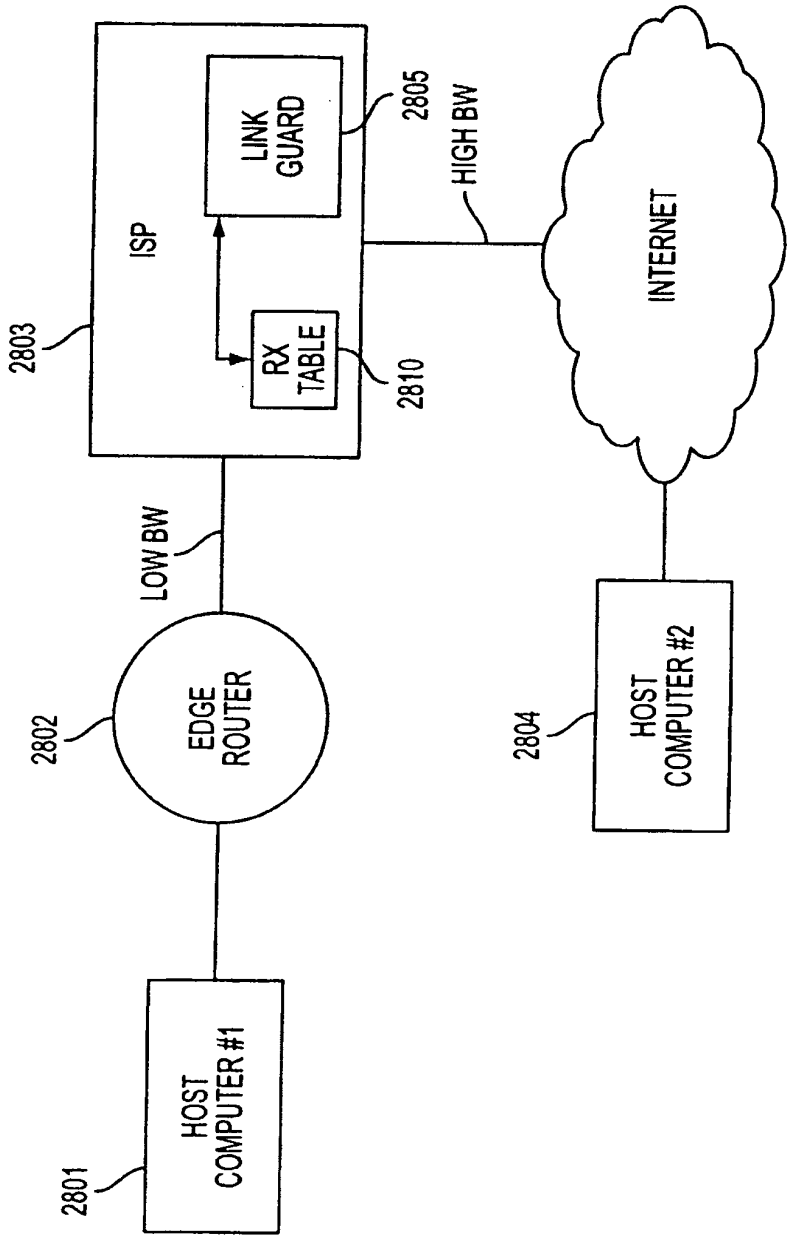


FIG. 28



New Bay Capital, LLC
Ex.1015-Page 2512 of 3151

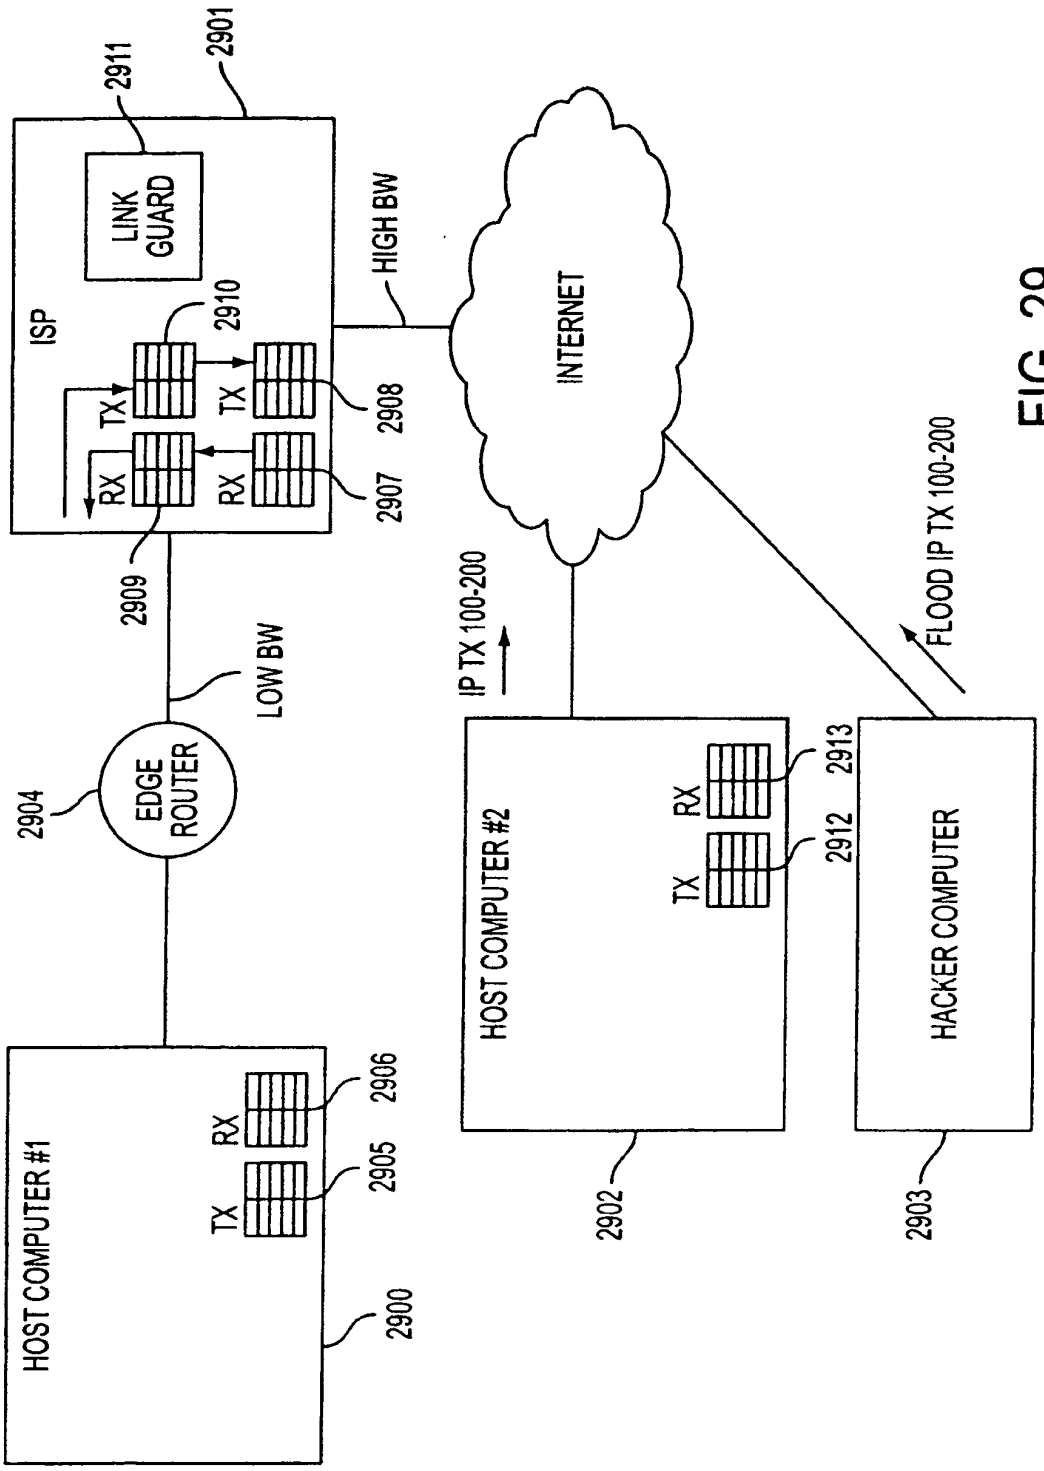


FIG. 29



New Bay Capital, LLC
Ex.1015-Page 2513 of 3151

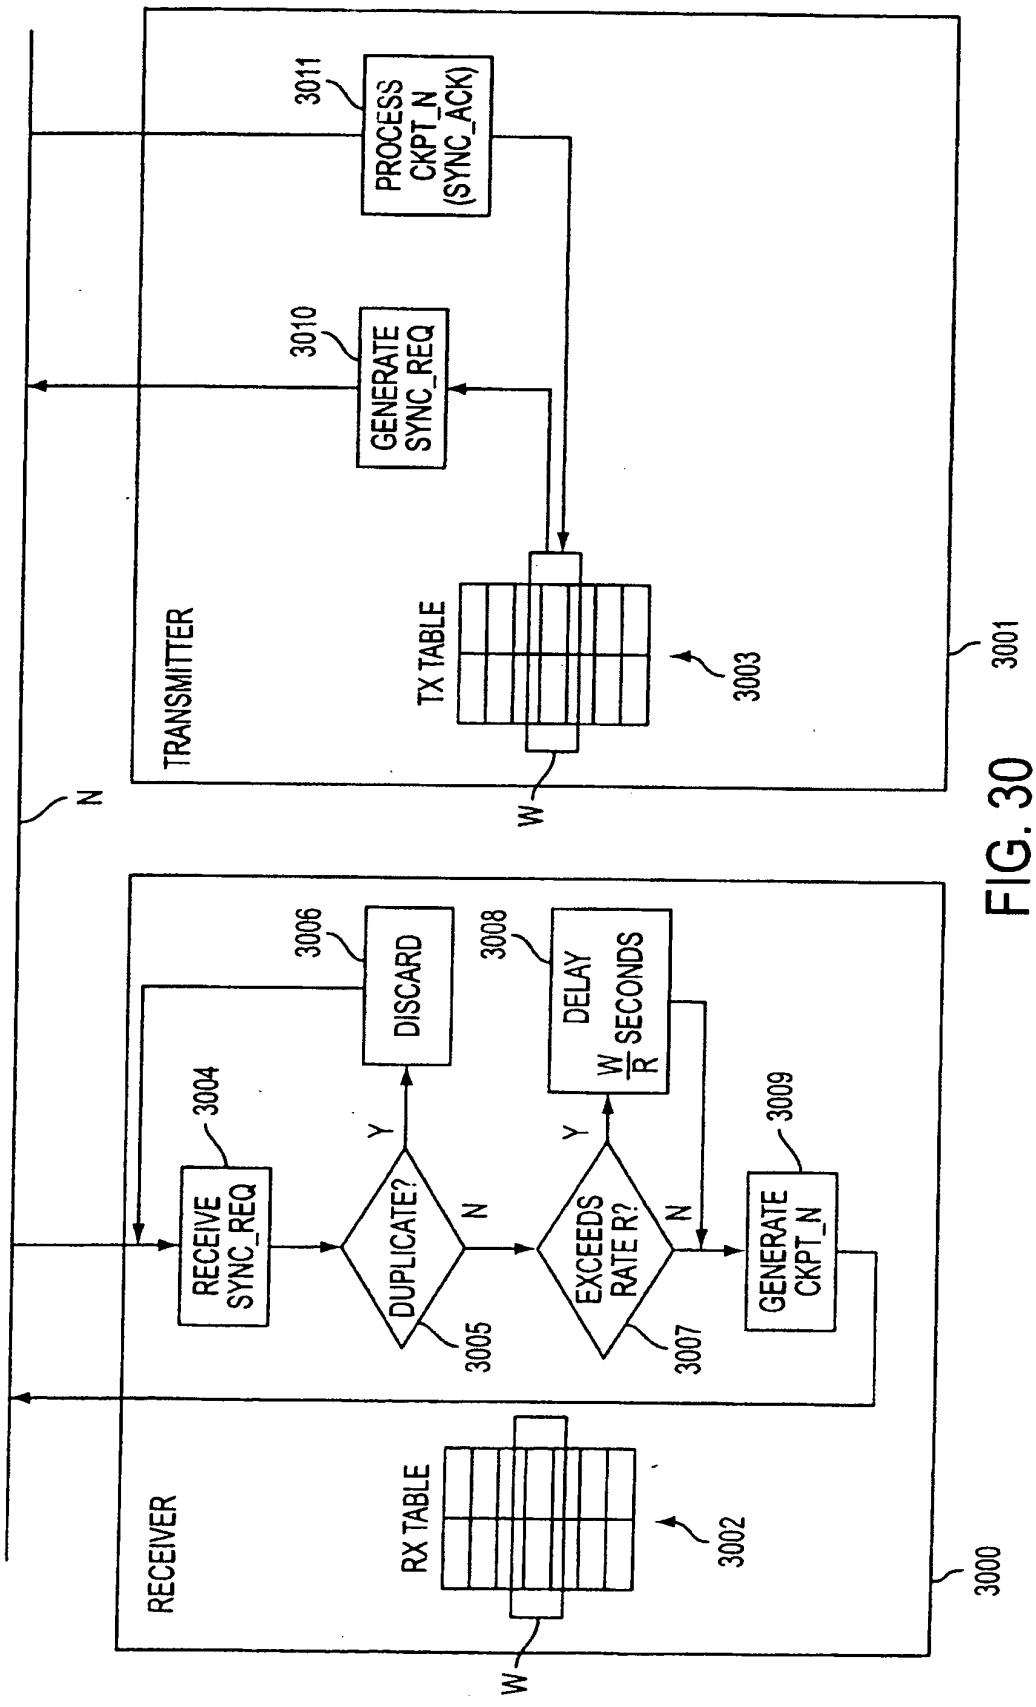
TRANSMITTER
3001
PROCESS CKPT_N (SYNC_ACK)
3011
GENERATE SYNC_REQ
3010
TX TABLE
3003
W
N
RECEIVER
3000
RX TABLE
3002
W
RECEIVE SYNC_REQ
3004
DUPLICATE?
3005
Y
N
DISCARD
3006
EXCEEDS RATE R?
3007
Y
N
DELAY W/R SECONDS
3008
GENERATE CKPT_N
3009

FIG. 30

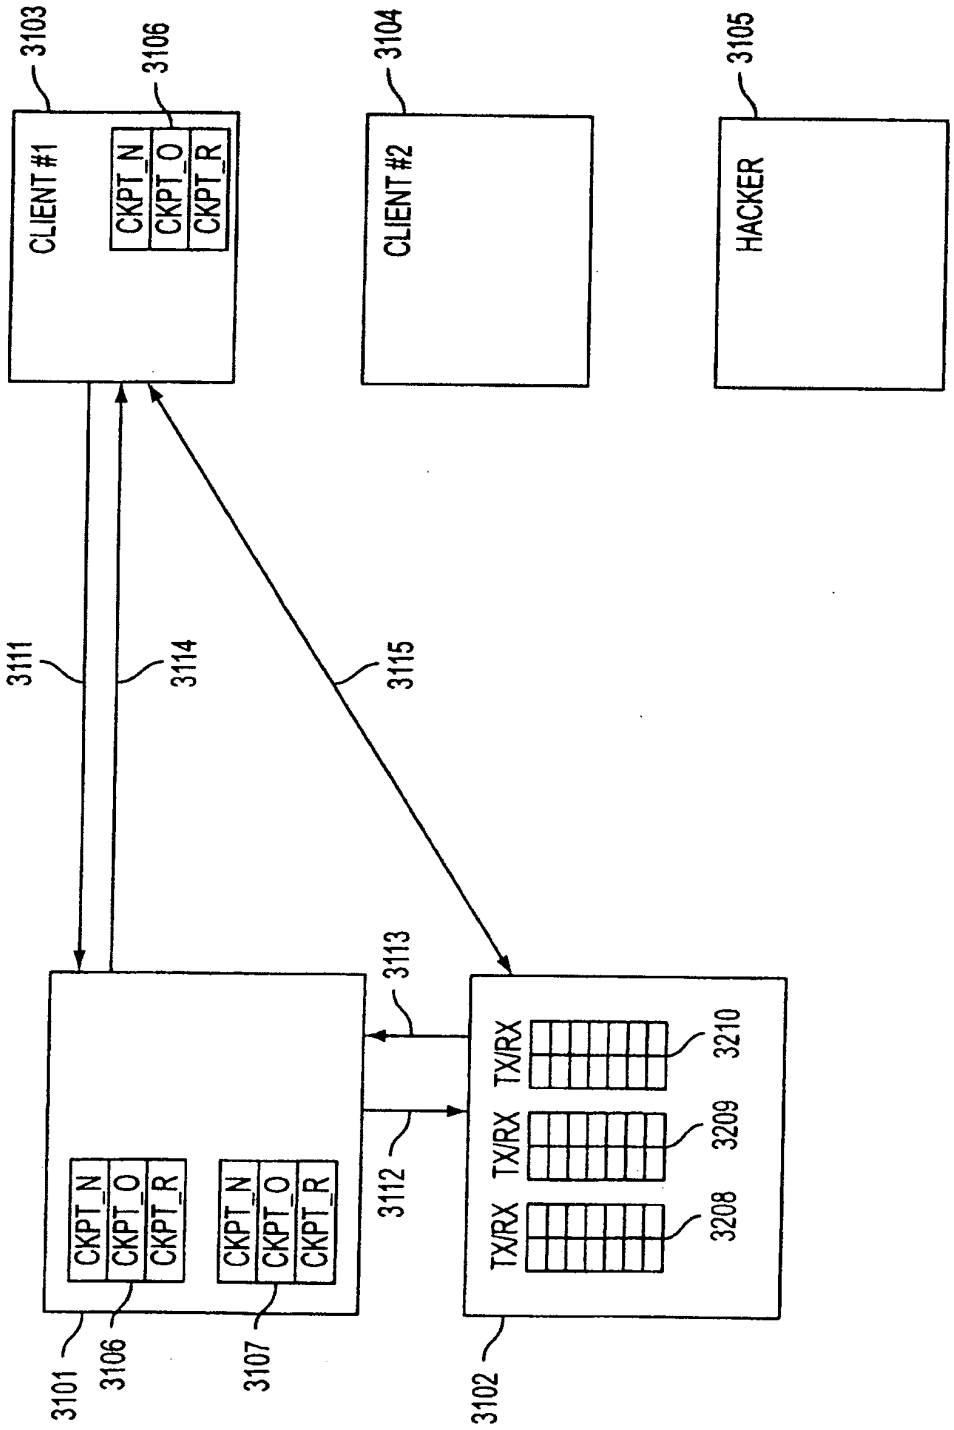


FIG. 31



New Bay Capital, LLC
Ex.1015-Page 2515 of 3151

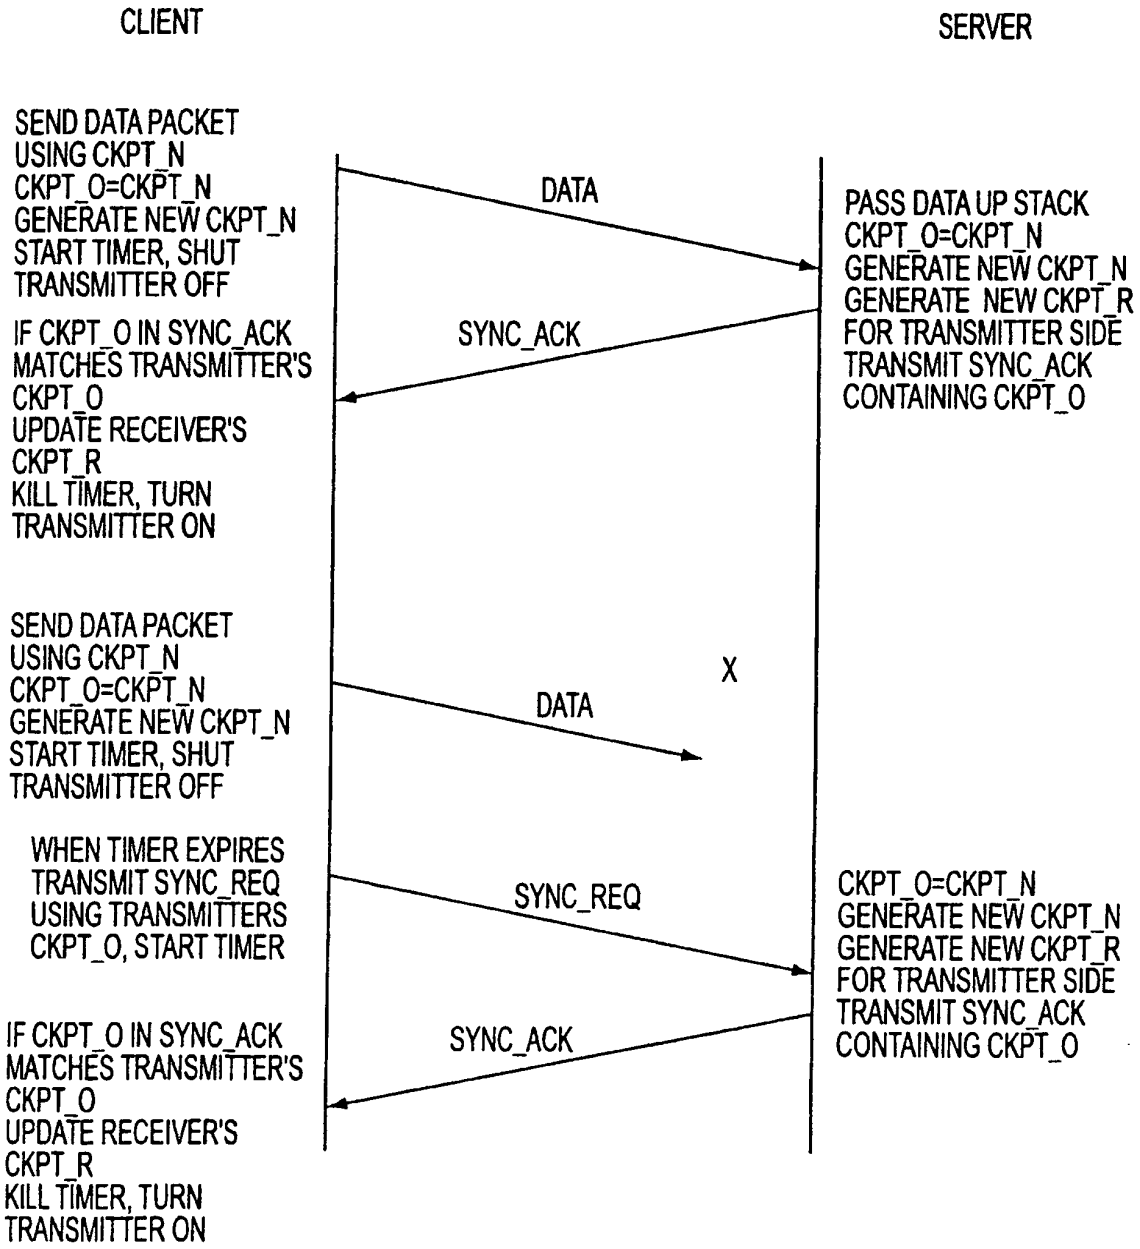


FIG. 32



New Bay Capital, LLC
Ex.1015-Page 2516 of 3151

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
23 August 2001 (23.08.2001)

PCT

(10) International Publication Number
WO 01/061922 A3

(51) International Patent Classification[7]: H04L 12/56, 29/06, 12/46

(21) International Application Number: PCT/US01/04340

(22) International Filing Date: 12 February 2001 (12.02.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/504,783 15 February 2000 (15.02.2000) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 09/504,783 (CON)
Filed on 15 February 2000 (15.02.2000)

(71) Applicant *(for all designated States except US)*: **SCIENCE APPLICATIONS INTERNATIONAL CORPORATION** [US/US]; 10260 Campus Point Drive, San Diego, CA 92121 (US).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **MUNGER, Edmund, Colby** [US/US]; 1101 Opaca Court, Crownsville, MD 21032 (US). **SCHMIDT, Douglas, Charles** [US/US]; 230 Oak Court, Severna Park, MD 21146 (US). **SHORT, Robert, Dunham, III** [US/US]; 38710 Goose Creek Lane, Leesburg, VA 20175 (US). **LARSON, Victor** [US/US]; 12026 Lisa Marie Court, Fairfax, VA 22033 (US). **WILLIAMSON, Michael** [US/US]; 26203 Ocala Circle, South Riding, VA 20152 (US).

(74) Agents: **WRIGHT, Bradley, C.** et al.; Banner & Witcoff, Ltd., 11th Floor, 1001 G Street, N.W., Washington, DC 20001-4597 (US).

*[Continued on next page]*

(54) Title: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS WITH ASSURED SYSTEM AVAILABILITY

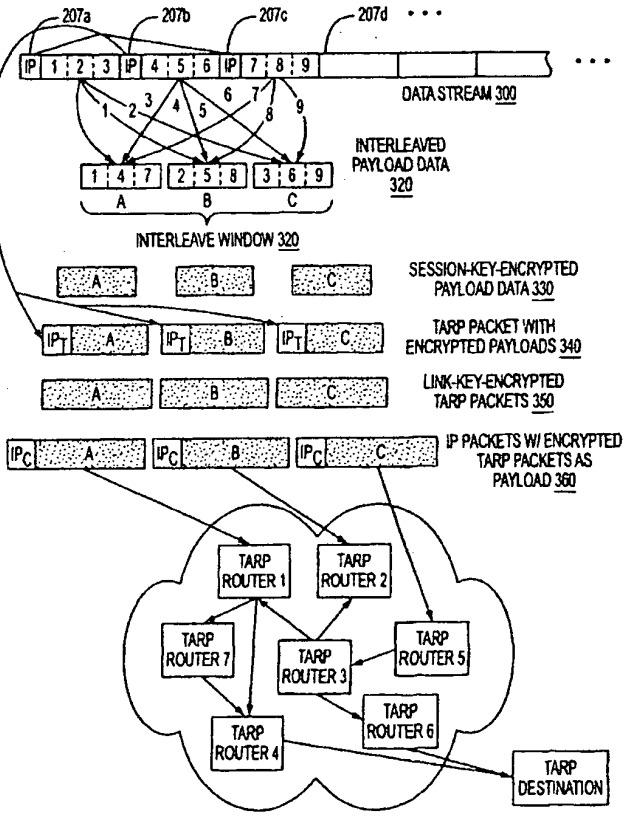


(57) Abstract: A plurality of computer nodes communicate using seemingly random Internet Protocol source and destination addresses. Data packets matching criteria defined by a moving window of valid addresses are accepted for further processing, while those that do not meet the criteria are quickly rejected. Improvements to the basic design include (1) a load balancer that distributes packets across different transmission paths according to transmission path quality; (2) a DNS proxy server that transparently creates a virtual private network in response to a domain name inquiry; (3) a large-to-small link bandwidth management feature that prevents denial-of-service attacks at system chokepoints; (4) a traffic limiter that regulates incoming packets by limiting the rate at which a transmitter can be synchronized with a receiver; and (5) a signaling synchronizer that allows a large number of nodes to communicate with a central node by partitioning the communication function between two separate entities.

WO 01/061922 A3

New Bay Capital, LLC
Ex.1015-Page 2517 of 3151

(81) **Designated States** *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States** *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

**Published:**
— *with international search report*

(88) **Date of publication of the international search report:**
6 March 2003

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56 H04L29/06 H04L12/46

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)
IPC 7 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, WPI Data, PAJ, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category * | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|---|---|---|
| A | EP 0 858 189 A (HITACHI LTD)<br>12 August 1998 (1998-08-12)<br>column 6, line 35 -column 10, line 13<br>---<br>-/-- | 1-27 |

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in 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 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 | Date of mailing of the international search report |
|---|---|
| 6 August 2002 | 20. 08. 2002 |

| Name and mailing address of the ISA | Authorized officer |
|---|---|
| European Patent Office, P.B. 5818 Patentlaan 2<br>NL - 2280 HV Rijswijk<br>Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,<br>Fax: (+31-70) 340-3016 | Ströbeck, A. |

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

| Category ° | Citation of document, with indication,where appropriate, of the relevant passages | Relevant to claim No. |
|---|---|---|
| A | MURTHY ET AL: "Congestion-oriented shortest multipath routing" PROCEEDINGS OF IEEE INFOCOM 1996. CONFERENCE ON COMPUTER COMMUNICATIONS. FIFTEENTH ANNUAL JOINT CONFERENCE OF THE IEEE COMPUTER AND COMMUNICATIONS SOCIETIES. NETWORKING THE NEXT GENERATION. SAN FRANCISCO, MAR. 24 - 28, 1996, PROCEEDINGS OF INFOCOM, L, vol. 2 CONF. 15, 24 March 1996 (1996-03-24), pages 1028-1036, XP010158171 ISBN: 0-8186-7293-5 abstract page 1028, left-hand column, line 38 -right-hand column, line 29 --- | 1-27 |
| E | WO 01 50688 A (TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)) 12 July 2001 (2001-07-12) page 11, line 18 -page 13, line 21 --- | 28,29,34 |
| A | WO 98 59470 A (TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)) 30 December 1998 (1998-12-30) page 4, line 5 -page 5, line 2 --- | 28-39 |
| X | WO 99 48303 A (CISCO TECHNOLOGY, INC.) 23 September 1999 (1999-09-23) page 1, line 8 -page 2, line 5 page 5, line 33 -page 6, line 15 page 7, line 21 - line 33 | 40,50 |
| A | | 41-49, 51-59 |
| | --- | |
| A | JONES JIM ET AL: "Distributed Denial of Service Attacks: Defenses" INTERNET ARTICLE, 'Online! 2000, XP002208785 Retrieved from the Internet: <URL:www.bai.org/pdf/DDOS-defense.pdf > 'retrieved on 2002-08-05! paragraph '0005! --- | 60-66 |
| X | WO 99 38081 A (ASCEND COMMUNICATIONS INC) 29 July 1999 (1999-07-29) page 9, line 13 -page 10, line 17 page 11, line 10 -page 12, line 2 | 67 |
| A | ----- | 68-71 |

# INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 01/04340

## Box I Observations where certain claims were found unsearchable (Continuation of item 1 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 II Observations where unity of invention is lacking (Continuation of item 2 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 fee, this Authority did not invite payment of any additional fee.

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.:

**Remark on Protest**

☐ The additional search fees were accompanied by the applicant's protest.

☒ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

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

1. Claims: 1-27

A system and a method to balance the load between communication paths with varying transmission quality.

2. Claims: 28-39

A system and a method to prevent someone from learning requested IP addresses by intercepting DNS requests.

3. Claims: 40-59

A method to prevent a denial-of-service attack from an unauthenticated user flooding dummy data packets on to a low bandwidth link.

4. Claims: 60-66

A method to prevent an authenticated user residing within a secure system from flooding it with dummy data packets.

5. Claims: 67-71

A method to allocate memory in a central computer communicating with a potentially large number of client computers.

| Patent document cited in search report | | Publication date | Patent family member(s) | | Publication date |
|---|---|---|---|---|---|
| EP 0858189 | A | 12-08-1998 | JP | 10224400 A | 21-08-1998 |
| | | | EP | 0858189 A2 | 12-08-1998 |
| | | | US | 6112248 A | 29-08-2000 |
| WO 0150688 | A | 12-07-2001 | SE | 517217 C2 | 07-05-2002 |
| | | | AU | 2564501 A | 16-07-2001 |
| | | | WO | 0150688 A1 | 12-07-2001 |
| | | | SE | 9904841 A | 30-06-2001 |
| | | | US | 2001006523 A1 | 05-07-2001 |
| WO 9859470 | A | 30-12-1998 | AU | 8052398 A | 04-01-1999 |
| | | | SE | 9702385 A | 24-12-1998 |
| | | | WO | 9859470 A2 | 30-12-1998 |
| WO 9948303 | A | 23-09-1999 | AU | 3098299 A | 11-10-1999 |
| | | | WO | 9948303 A2 | 23-09-1999 |
| WO 9938081 | A | 29-07-1999 | US | 6055575 A | 25-04-2000 |
| | | | AU | 2562599 A | 09-08-1999 |
| | | | CA | 2318267 A1 | 29-07-1999 |
| | | | EP | 1064602 A1 | 03-01-2001 |
| | | | WO | 9938081 A1 | 29-07-1999 |


Reexam $

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | |
|---|---|
| In re *Inter Partes* Reexamination of: | |
| Victor Larson et al. | Control No.: 95/001,851 |
| U. S. Patent No. 7,418,504 | Group Art Unit: 3992 |
| Issued: August 26, 2008 | Examiner: Roland G. Foster |
| For: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | Confirmation No. 1688 |

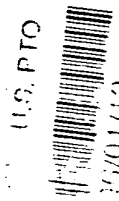


Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Commissioner:

**TRANSMITTAL LETTER**

Enclosed please find the following:

1. Patent Owner's Response to Office Action (71 pages);
2. Declaration of Angelos D. Keromytis, Ph.D. (31 pages) with appended *curriculum vitae*;
3. Declaration of Dr. Robert Dunham Short III in control no. 95/001,788 (5 pages);
4. Appendix - List of Exhibits (1 page);
5. Exhibits Listed on Appendix;
6. Petition Seeking Waiver of 37 C.F.R. § 1.943 for Patent Owner's Response to Office Action of March 1, 2012 (3 pages);
7. Check in the amount of $400 for the Petition Fee; and
8. Certificate of Service (2 pages).

Please grant any extension of time and charge any additional fees to Deposit Account No. 06-0916.

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: June 1, 2012

By: /Joseph E. Palys/
Joseph E. Palys
Reg. No. 46,508



New Bay Capital, LLC
Ex.1015-Page 2525 of 3151

U.S. PTO
03/01/12

PATENT
Customer No. 22,852
Attorney Docket No. 11798.0007

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | | |
|---|---|---|
| In re *Inter Partes* Reexamination of: | ) | |
| | ) | |
| Victor Larson et al. | ) | Control No.: 95/001,851 |
| | ) | |
| U.S. Patent No. 7,418,504 | ) | Group Art Unit: 3992 |
| | ) | |
| Issued: August 26, 2008 | ) | Examiner: Roland Foster |
| | ) | |
| For: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | ) | Confirmation No.: 1688 |

Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

**PATENT OWNER'S RESPONSE TO OFFICE ACTION OF MARCH 1, 2012**

**Table of Contents**

**Page**

I. Introduction ..... 1

II. Background ..... 2

A. Overview of the '504 Patent ..... 2

B. Applicable Legal Standards for Anticipation ..... 4

C. Applicable Legal Standards for Obviousness ..... 4

III. The Rejections Are Improper and Should Be Withdrawn ..... 5

A. The Rejections Based on *Lendenmann* Should Be Withdrawn ..... 5

1. Overview of *Lendenmann* ..... 5

2. Rejection of Claims 1-3, 5, 6, 14-30, 33-54, and 57-60 Under 35 U.S.C. § 102(b) Based on *Lendenmann* (Issue 1) ..... 6

a. Independent Claim 1 ..... 6

(i) *Lendenmann*'s Cell Directory Service Used for "Returning the Network Address Corresponding to a Secure Domain Name" Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 7

(ii) *Lendenmann*'s Cell Directory Service, "Integrated with the Security Services," Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 8

(iii) *Lendenmann*'s Incomplete Binding Handles Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 10

(iv) *Lendenmann*'s "Authentication Challenge" Fails to Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 11

(v) *Lendenmann*'s "Server Status" Counters Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 12



New Bay Capital, LLC
Ex.1015-Page 2527 of 3151

(vi) *Lendenmann*'s "Online Documentation" Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 13

b. Independent Claims 36 and 60 ..... 15

c. Dependent Claims 5, 23, and 47 ..... 15

d. Dependent Claims 16, 17, 27, 33, 40, 41, 51, and 57 ..... 17

e. Dependent Claims 24 and 48 ..... 18

f. Dependent Claims 2, 3, 6, 14, 15, 18-22, 25, 26, 28-30, 34, 35, 37-39, 42-46, 49, 50, 52-54, 58 and 59 ..... 19

3. Rejection of Claims 1-6, 14-30, 33-54, and 57-60 Under 35 U.S.C. § 103(a) Based on *Lendenmann* (Issue 2) ..... 19

4. Rejection of Claim 7 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Wesinger* (Issue 3) ..... 20

5. Rejection of Claims 8 and 9 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* (Issue 4) ..... 21

a. Claim 8 ..... 21

b. Claim 9 ..... 23

6. Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* and *Schneier* (Issue 5) ..... 23

7. Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* and *Martin* (Issue 6) ..... 24

8. Rejection of Claims 12 and 13 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* and RFC 793 (Issue 7) ..... 25

a. Claim 12 ..... 25

b. Claim 13 ..... 26

9. Rejection of Claims 31, 32, 55, and 56 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Ludwig* and RFC 793 (Issue 8) ..... 26

B. The Rejections Based on *Aziz* Should Be Withdrawn ..... 27

1. Overview of *Aziz* ..... 27



New Bay Capital, LLC
Ex.1015-Page 2528 of 3151

2. Rejection of Claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 102(b) Based on *Aziz* (Issue 9) ..... 28

a. Independent Claim 1 ..... 28

(i) *Aziz*'s SX Records Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 29

(ii) *Aziz*'s KEY and SIG Records Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 31

(iii) *Aziz*'s "Information Used for Secure Communications with Protected Hosts" Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 32

(iv) *Aziz*'s Reference to RFC 2065 Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 33

b. Independent Claims 36 and 60 ..... 34

c. Dependent Claims 5, 23, and 47 ..... 34

d. Dependent Claim 8 ..... 35

e. Dependent Claims 17 and 41 ..... 36

f. Dependent Claims 18 and 42 ..... 36

g. Dependent Claims 24 and 48 ..... 37

h. Dependent Claim 50 ..... 38

i. Dependent Claims 2, 6, 7, 14, 15, 19-22, 24, 25, 27, 33-40, 43-46, 49, 51, 52, 58 and 59 ..... 39

3. Rejection of Claims 1, 2, 5-9, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 103(a) Based on *Aziz* (Issue 10) ..... 39

4. Rejection of Claims 3, 4, and 26 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Lawton* (Issue 11) ..... 40

a. Claims 3 and 4 ..... 40

b. Claim 26 ..... 40



New Bay Capital, LLC
Ex.1015-Page 2529 of 3151

5. Rejection of Claim 9 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Franaszek* (Issue 12) ..... 42

a. *Aziz* and *Franaszek* Do Not Disclose or Suggest the Features of Claim 9 ..... 43

b. The Alleged Combination Would Change the Principles of Operation of *Aziz* ..... 43

c. One of Ordinary Skill in the Art Would Not Have Relied on *Franaszek* ..... 45

6. Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Schneier* (Issue 13) ..... 46

7. Rejection of Claims 11-13 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Martin* (Issue 14) ..... 46

a. Claim 11 ..... 46

b. Claim 12 ..... 47

c. Claim 13 ..... 48

8. Rejection of Claims 29-32 and 53-56 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Ludwig* (Issue 15) ..... 48

C. The Rejections Based on *Kiuchi* and *Pfaffenberger* Should Be Withdrawn ..... 49

1. Overview of *Kiuchi* ..... 49

2. Overview of *Pfaffenberger* ..... 50

3. Rejection of Claims 1-4, 6, 8-10, 12-19, 22, 24-30, 33, 34, 36-43, 46, 48-54, and 57-60 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* (Issue 16) ..... 51

a. Independent Claim 1 ..... 51

(i) The C-HTTP Name Server Returning the Public Key of the Server-Side Proxy Is Not an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 51

(ii) The C-HTTP Name Server Returning the IP Address of the Server-Side Proxy Is Not an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link" ..... 52



New Bay Capital, LLC
Ex.1015-Page 2530 of 3151

(iii) *Pfaffenberger* Does Not Remedy the Deficiencies of *Kiuchi* ..... 53

(iv) One of Ordinary Skill in the Art Would Not Have Relied on *Pfaffenberger* ..... 55

b. Independent Claims 36 and 60 ..... 56

c. Dependent Claims 8, 9, 10, 12, and 13 ..... 56

d. Dependent Claim 10 ..... 59

e. Dependent Claim 12 ..... 60

f. Dependent Claim 13 ..... 62

g. Dependent Claims 17 and 41 ..... 63

h. Dependent Claims 24 and 48 ..... 63

i. Dependent Claim 26 ..... 64

j. Dependent Claim 27 ..... 65

k. Dependent Claims 2-4, 6, 7, 11, 14-16, 18, 19, 22, 25, 28-30, 33, 34, 37-40, 42, 43, 46, 49-54, and 57-59 ..... 65

4. Rejection of Claims 5, 23, and 47 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Rivest* (Issue 17) ..... 66

5. Rejection of Claim 7 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Borella* (Issue 18) ..... 66

6. Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Martin* (Issue 19) ..... 67

7. Rejection of Claims 20, 21, 35, 44, and 45 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Broadhurst* (Issue 20) ..... 68

8. Rejection of Claims 31, 33, 35, and 56 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Ludwig* (Issue 21) ..... 68

D. Secondary Considerations Demonstrate Nonobviousness ..... 69

IV. Conclusion ..... 71



New Bay Capital, LLC
Ex.1015-Page 2531 of 3151

## I. Introduction

VirnetX Inc. ("VirnetX"), the owner of U.S. Patent No. 7,418,504 ("the '504 patent"), provides the following remarks in response to the Office Action ("OA") and Order granting reexamination ("Order") mailed March 1, 2012, in the above-identified reexamination proceeding. The U.S. Patent and Trademark Office ("USPTO") issued the Office Action and Order in response to a Request for Reexamination ("Request") filed by Cisco Systems, Inc. ("Cisco" or "Requester") on December 13, 2011.

The patent at issue in this reexamination, the '504 patent, is part of a family of patents ("Munger patent family") that stems from U.S. provisional application nos. 60/106,261 ("the '261 application"), filed on October 30, 1998, and 60/137,704 ("the '704 application"), filed on June 7, 1999. The '504 patent is a continuation of U.S. application no. 09/558,210 ("the '210 application"), filed on April 26, 2000 (now abandoned), which is a continuation-in-part of U.S. application no. 09/504,783 (now U.S. Patent No. 6,502,135, "the '135 patent"). The '135 patent is a continuation-in-part of U.S. application no. 09/429,643 (now U.S. Patent No. 7,010,604, "the '604 patent"), which claims priority to the '261 and '704 applications.

The Munger patent family discloses numerous inventions relating to secure communications. Patents in this family have been subject to several reexamination proceedings and district court actions. For instance, three other patents from the family were asserted in an action against Microsoft Corporation in the Eastern District of Texas.[1] The jury found the asserted claims willfully infringed and not invalid, and awarded VirnetX over one hundred million dollars in damages. (Ex. A-1 at 2.) Microsoft also sought reexamination of two of the patents, but all claims were confirmed during those proceedings. (*See* control nos. 95/001,269 and 95/001,270.) And just recently, the USPTO denied a request for reexamination of one of the patents in the Munger patent family. (Order in control no. 95/001,792.)

Given that the validity of the patents in the Munger patent family has now been tested multiple times, and for the other reasons set forth below, including that the asserted references do not disclose or suggest the combination of features recited in the claims, Patent Owner requests reconsideration and withdrawal of all the rejections in the Office Action and confirmation of the patentability of all of the claims of the '504 patent.

[1] One of these patents, U.S. Patent No. 6,839,759, was asserted initially but was dropped from this case before trial.



New Bay Capital, LLC
Ex.1015-Page 2532 of 3151

This Response is supported by a Declaration of Angelos D. Keromytis, Ph.D. ("Keromytis Decl."), and by a Declaration of Dr. Robert Dunham Short III ("Short Decl.").

## II. Background

### A. Overview of the '504 Patent

The '504 patent discloses several embodiments of a domain name service ("DNS") system for establishing a secure communication link, such as a virtual private network ("VPN") communication link, between devices connected to a network. In one such embodiment, a novel, specialized DNS server receives a traditional DNS request, and the DNS server automatically facilitates the establishment of a secure communication link between a target node and a user. (Keromytis Decl. ¶ 16; '504 patent 39:46-51.) This specialized DNS server is different from a conventional DNS server known at the time of the invention for at least the reason that the specialized DNS server supports the establishment of a secure communication link beyond merely a requested IP address or public key. (Keromytis Decl. ¶ 16.)

For example, in the exemplars of FIGS. 26 and 27 of the '504 patent, reproduced below, a DNS server 2602 including a DNS proxy 2610 supports establishing a VPN link between a computer 2601 and a secure target site 2604. ('504 patent 39:67-41:59; Keromytis Decl. ¶ 17.)

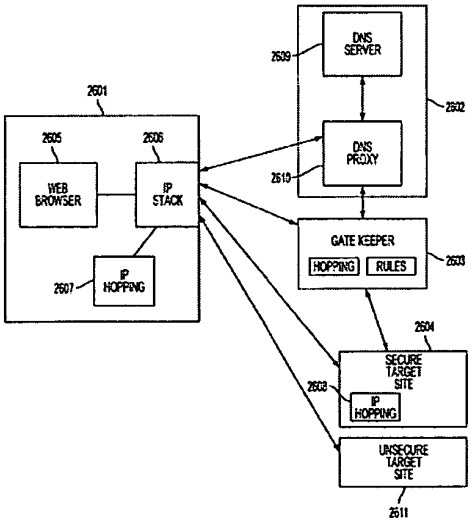


FIG. 26

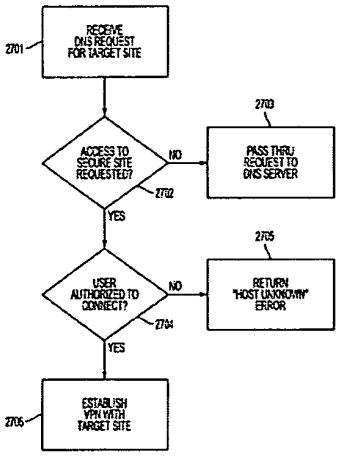


FIG. 27

In one embodiment, the DNS server 2602 receives a DNS request for a target site from computer 2601. ('504 patent 40:49-52; Keromytis Decl. ¶ 18.) The DNS proxy 2610 determines whether the target site is a secure site. ('504 patent 40:6-8, 40:49-56; Keromytis Decl. ¶ 18.) If access to a secure site has been requested, the DNS proxy 2610 determines whether the computer 2601 is authorized to access the site. ('504 patent 40:57-59; Keromytis Decl. ¶ 18.) If so, the DNS proxy



New Bay Capital, LLC
Ex.1015-Page 2533 of 3151

2610 transmits a message to gatekeeper 2603 to facilitate the creation of a VPN link between computer 2601 and secure target site 2604. ('504 patent 40:18-24.) The DNS proxy 2610 then responds to the computer's 2601 DNS request with an address received from the gatekeeper 2604. (*Id.* at 40:19-22; Keromytis Decl. ¶ 18.) A secure VPN link is then established between the computer 2601 and the secure target site 2604. ('504 patent 41:5-8; Keromytis Decl. ¶ 18.) As shown in this example, the specialized DNS server supports creating a secure communication link and does more than a conventional DNS server at the time of the invention. (Keromytis Decl. ¶ 18.)

The '504 patent highlights this distinction between the specialized DNS server disclosed in its specification and a conventional DNS scheme, which merely returns a requested IP address or public key:

> Conventional Domain Name Servers (DNSs) provide a look-up function that returns the IP address of a requested computer or host. For example, when a computer user types in the web name "Yahoo.com," the user's web browser transmits a request to a DNS, which converts the name into a four-part IP address that is returned to the user's browser . . . .
>
> . . .
>
> One conventional scheme that provides secure virtual private networks over the Internet provides the DNS server with the public keys of the machines that the DNS server has the addresses for. This allows hosts to retrieve automatically the public keys of a host that the host is to communicate with so that the host can set up a VPN without having the user enter the public key of the destination host. One implementation of this standard is presently being developed as part of the FreeS/WAN project (RFC 2535).
>
> The conventional scheme suffers from certain drawbacks. For example, any user can perform a DNS request. Moreover, DNS requests resolve to the same value for all users.
>
> According to certain aspects of the invention, a specialized DNS server traps DNS requests and, if the request is from a special type of user (e.g., one for which secure communication services are defined), the server does not return the true IP address of the target node, but instead automatically sets up a virtual private network between the target node and the user.

('504 patent 39:7-51; Keromytis Decl. ¶ 19.) Compared with a conventional DNS known at the time of the filing date of the '504 patent, the specialized DNS disclosed in the '504 patent supports establishing a secure communication link. (Keromytis Decl. ¶ 19.) The claims of the '504 patent are also directed to a domain name service for establishing a secure communication link. (*See, e.g.*, '504 patent 55:49-56, 57:48-58, 60:3-14; Keromytis Decl. ¶ 19.)



New Bay Capital, LLC
Ex.1015-Page 2534 of 3151

**B. Applicable Legal Standards for Anticipation**

To support a rejection under 35 U.S.C. § 102, each and every element of each claim at issue must be found in that single reference. *See* M.P.E.P. § 2131. "The identical invention must be shown in as complete detail as is contained in the . . . claim." *Id.* (quoting *Richardson v. Suzuki Motor Co.*, 868 F.2d 1126, 1236 (Fed. Cir. 1989)). Further, "[t]he elements must be arranged as required by the claim . . . ." *Id.* (citing *In re Bond*, 910 F.2d 831, 832 (Fed. Cir. 1990)). Thus, "unless a reference discloses within the four corners of the document not only all of the limitations claimed but also all of the limitations arranged or combined in the same way as recited in the claim, it . . . cannot anticipate under 35 U.S.C. § 102." *Net MoneyIn, Inc. v. Verisign, Inc.*, 545 F.3d 1359, 1369 (Fed. Cir. 2008). Moreover, "[t]he requirement that the prior art elements themselves be 'arranged as in the claim' means that claims cannot be 'treated . . . as mere catalogs of separate parts, in disregard of the part-to-part relationships set forth in the claims and that give the claims their meaning.'" *Therasense, Inc. v. Becton, Dickinson & Co.*, 593 F.3d 1325, 1332 (Fed. Cir. 2010) (quoting *Lindemann Maschinenfabrik GmbH v. Am. Hoist & Derrick Co.*, 730 F.2d 1452, 1459 (Fed. Cir. 1984)).

**C. Applicable Legal Standards for Obviousness**

Obviousness is a question of law based on underlying factual inquiries that include, inter alia, determining the scope and content of the prior art and ascertaining the differences between the claimed invention and prior art. *See* M.P.E.P. § 2141(II). In order to establish a prima facie case of obviousness, the Examiner must "include[] findings of fact concerning the state of the art and the teachings of the references . . . ." *Id.* Moreover, "[o]nce the findings of fact are articulated, [the Examiner] must provide an explanation to support an obviousness rejection under 35 U.S.C. [§] 103." *Id.*

The reasons why the claimed invention would have been obvious must be clearly articulated and cannot be premised on conclusory statements. *Id.* § 2142. In addition, the references relied on must be enabling, *id.* § 2145, and "[t]he mere fact that references <u>can</u> be combined or modified does not render the resultant combination obvious unless the results would have been predictable to one of ordinary skill in the art" at the time the invention was made, *id.* § 2143.01(III) (citation omitted). "All words in a claim must be considered in judging the patentability of that claim against the prior art." *Id.* § 2143.03 (citation omitted). Also, "[i]n determining the differences between the prior art and the claims, the question under 35 U.S.C. [§] 103 is not whether the differences <u>themselves</u> would

have been obvious, but whether the claimed invention as a whole would have been obvious." *Id.* § 2141.02(I) (citations omitted).

## III. The Rejections Are Improper and Should Be Withdrawn

The Office Action rejects claims 1-60 of the '504 patent as being anticipated or obvious in view of several references.[2] As explained below, however, the combinations of references relied on by the Office Action do not disclose or suggest the combination of features recited in the claims.

### A. The Rejections Based on *Lendenmann* Should Be Withdrawn

The Office Action rejects certain claims under 35 U.S.C. §§ 102 and 103 based on Rolf Lendenmann, "Understanding OSF DCE 1.1 for AIX and OS/2" ("*Lendenmann*") alone and in combination with certain secondary references. As discussed below, however, the rejections under §§ 102 and 103 based on these references should be withdrawn.

#### 1. Overview of *Lendenmann*

*Lendenmann* discloses a distributed computing environment ("DCE"), which "is a layer of services that allows distributed applications to communicate with a collection of computers, operating systems, and networks." (*Lendenmann* 7.) As illustrated in Figure 3, *Lendenmann*'s DCE may include several different components, including security services, time services, and directory services. (*Id.* at 8.)

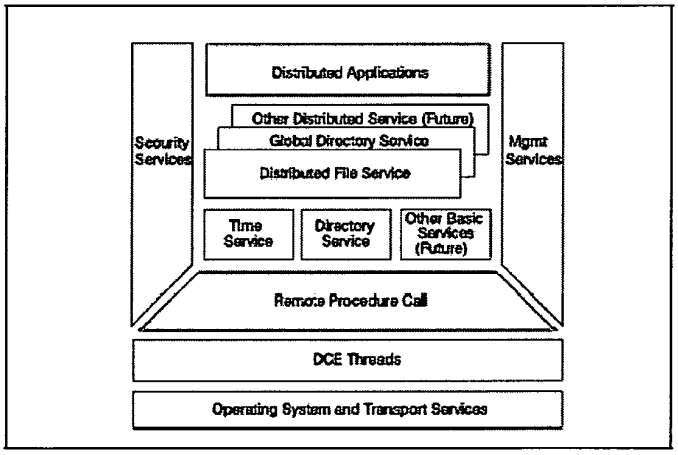


*Figure 3. DCE Architecture*

(*Id.*) It further discloses that a collection of machines, operating systems, and networks managed by a single set of DCE services constitutes a "DCE cell." (*Id.*) At a minimum, a cell must contain a

[2] In formulating the rejections, the Office Action adopts portions of the Request and the three supporting claim chart exhibits of the Request. (*See* OA at 11, 16, 18.) Accordingly, this Response sometimes cites to the Request as support for what is asserted by both the Office Action and the Request.



New Bay Capital, LLC
Ex.1015-Page 2536 of 3151

Security Server, a Cell Directory Server ("CDS"), and Distributed Time Servers. (*Id.* at 9.) These components provide different services for establishing remote procedure calls ("RPCs") between clients and servers.

For example, *Lendenmann* explains that the CDS may receive a request and then "return[] the network address of the named resource." (*Id.* at 21, 29.) Additionally, the CDS may provide other identification-related information to the client in a partially complete "binding handle," for example, the appropriate protocol for communications with the server (e.g., TCP/IP, UDP/IP, etc.) or the object UUID. (*Id.* at 182-85.)

*Lendenmann* also explains that the Security Service provides services to implement security measures for RPCs between clients and servers, although no security is actually necessary for RPCs. (*Id.* at 191-94, 207, explaining that a client may "optionally set up authenticated RPC.") *Lendenmann* first explains that in order to pursue security-enhanced communications between clients and servers, a client needs to "add[] . . . security information to the server binding handle" received from the CDS. (*Id.* at 191.) Afterwards, a client must obtain a "session key" from the Security Service for implementing any authentication or encryption measures. (*Id.* at 192, 194.) The client may then send the session key to the server, which provides the client with a challenge. (*Id.* at 194.) If the client matches the challenge criteria, "everything is set for authenticated RPC," and security-enhanced communications may proceed. (*Id.*)

### 2. Rejection of Claims 1-3, 5, 6, 14-30, 33-54, and 57-60 Under 35 U.S.C. § 102(b) Based on *Lendenmann* (Issue 1)

The Office Action rejects claims 1-3, 5, 6, 14-30, 33-54, and 57-60 under § 102(b) based on *Lendenmann*. (OA at 5.) For the reasons discussed below, this rejection should be withdrawn and the claims should be confirmed.

#### a. Independent Claim 1

Independent claim 1 is directed to, among other things, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." The Request and the Office Action assert that six different features of *Lendenmann* disclose the recited "indication." These assertions are incorrect. Each of the six features in *Lendenmann*, discussed in turn below, fails to disclose a domain name service system configured to comprise an indication that the domain name service supports establishing a secure communication link. Rather, these features disclose nothing more than a conventional domain name service system that is both recognized and distinguished by the '504 patent.

**(i) *Lendenmann*'s Cell Directory Service Used for "Returning the Network Address Corresponding to a Secure Domain Name" Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Office Action and the Request first assert that "by returning the network address corresponding to a secure domain name," the CDS provides "an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (OA at 11; Req. Ex. F-1 at 13.) This is incorrect.

*Lendenmann* discloses a conventional name service function for the CDS: "when given a name, CDS returns the network address of the named resource." (*Lendenmann* 21.) Rather than "comprising an indication that the domain name service system supports establishing a secure communication link," *Lendenmann* instead explains that the CDS merely provides server identification information to a client, as illustrated with Figure 15.

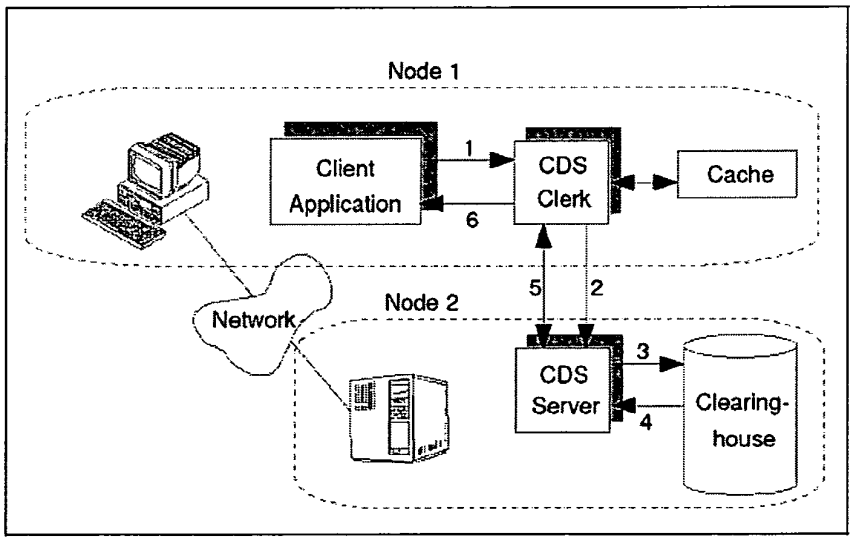


*Figure 15. CDS Components Performing a CDS Look-up*

(*Id.* at 29.) As *Lendenmann* explains, the CDS works as follows: (1) the client sends a lookup request to the CDS clerk; (2) the CDS clerk checks its cache and, not finding the name there, contacts the CDS server; (3) the CDS server checks to see if the name is in the clearinghouse; (4) the CDS server obtains the requested information if the name exists in the clearinghouse; (5) the CDS server then returns the information to the CDS clerk; and (6) the CDS clerk caches the information and passes the requested information to the client. (*Id.* at 29-30.) Thus, *Lendenmann* discloses a CDS lookup process that simply returns a name, taking no measures to provide any indication that the CDS supports establishing a secure communication link. (Keromytis Decl. ¶¶ 27-28.) Instead,



New Bay Capital, LLC
Ex.1015-Page 2538 of 3151

*Lendenmann* discloses a separate Security Service component that handles the aspects of any desired security measures. (*Lendenmann* 191-94.)

By simply returning server identification information to the client, the CDS disclosed in *Lendenmann* is consistent with a conventional domain name service system that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. For instance, the '504 patent explains that "[c]onventional Domain Name Servers (DNSs) provide a look-up function that returns the IP address of a requested computer or host." ('504 patent 39:7-13.) In another example, the '504 patent identifies conventional domain name service systems that store public keys for different machines, allowing hosts to retrieve the keys and then proceed to communicate with the different machines to establish VPNs. (*Id.* at 39:34-42.) The '504 patent recognizes that such conventional domain name systems suffer from certain drawbacks and accordingly discloses various embodiments to address these problems, including "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g., id.* at 39:43-41:59; Keromytis Decl. ¶ 19.) Because the CDS disclosed in *Lendenmann* performs no functions beyond those that the '504 patent distinguishes as characterizing a conventional domain name service, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS to comprise an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶ 28.)

**(ii) *Lendenmann*'s Cell Directory Service, "Integrated with the Security Services," Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Office Action and Request next assert that by only performing operations for users authorized by access control lists ("ACLs"), the CDS provides "an indication that the domain name service system supports establishing a secure communication link," as disclosed in claim 1. (OA at 11; Req. Ex. F-1 at 13-14.) This too is incorrect.

*Lendenmann* discloses that the Security Service plays a gatekeeper role to control access to the information in the CDS. For example, in response to a name request, "ACL management software examines the ACL entry associated with that name or principal name and grants or denies the [CDS] operation." (*Lendenmann* 34.) The Security Service's gatekeeping function, however, has no bearing on the operations of the alleged domain name service system, the CDS. (Keromytis Decl.

¶29.) Even if a client has authorization to receive the information in the CDS, the CDS still merely "returns the network address of the named resource" when given a name in a request, as discussed above and as illustrated in Figure 15 below. (*Lendenmann* 21.)

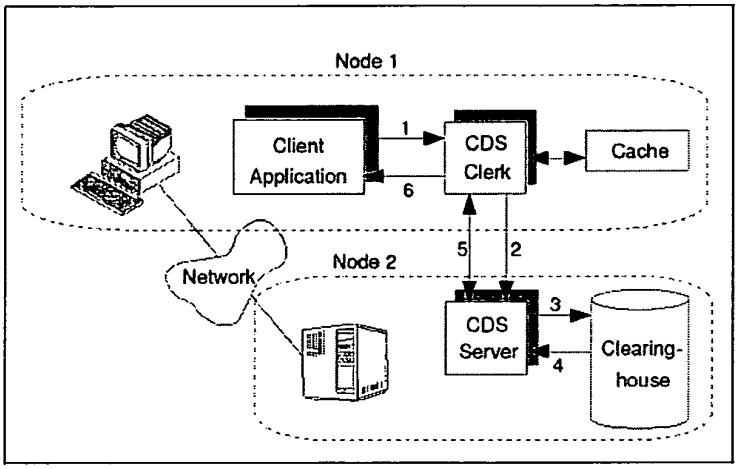


*Figure 15. CDS Components Performing a CDS Look-up*

(*Id.* at 29.)

By contrast, *Lendenmann* discloses a Security Service function, separate from the CDS, to handle any security-related measures for communications after the CDS has returned server identification information to the client. (*Id.* at 191-94.) For example, the Security Service provides for *Lendenmann*'s authentication procedures, employing a "Security Server" to run the authentication process without any action by or reference to a CDS, as illustrated in Figure 21:

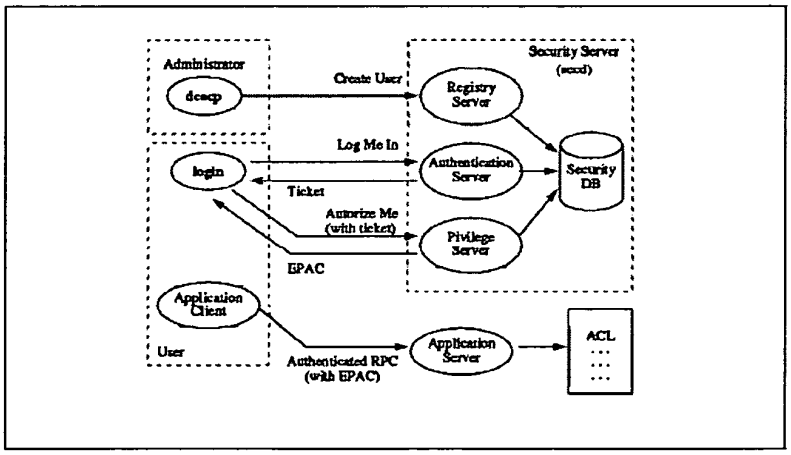


*Figure 21. Authentication Process.*

(*Id.* at 53-55.)

Thus, although the Security Service may permit or deny access to the CDS, the CDS itself merely returns server identification information as illustrated in Figure 15. Accordingly, the CDS

disclosed in *Lendenmann* only performs functions that the '504 patent recognizes and distinguishes as consistent with conventional domain name service systems. (*See, e.g.*, '504 patent 39:43-41:59; Keromytis Decl. ¶ 19.) As a result, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS feature to disclose an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶ 30.)

**(iii) *Lendenmann*'s Incomplete Binding Handles Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Office Action and the Request also assert that the CDS, by providing clients with binding handles containing identification information for servers, discloses "a domain name service system configured . . . to comprise an indication that the domain name service system supports establishing a secure communication link." (OA at 11; Req. Ex. F-1 at 14-16.) This is incorrect.

*Lendenmann* discloses that the CDS may provide binding handles to a client to find a target server for an RPC. (*Lendenmann* 182-85.) These binding handles provide a network address and other identification-related information, for example, the appropriate protocol for communications with the server (e.g., TCP/IP, UDP/IP, etc.) or the object UUID. (*Id.*) The Request argues that because "[t]he binding handles are annotated with security information," this means that the CDS provides these annotated binding handles, and therefore *Lendenmann* discloses "a domain name service system configured . . . to comprise an indication that the domain name service system supports establishing a secure communication link." (*See* Req. Ex. F-1 at 15-16, citing *Lendenmann* 185.) But the Request overlooks the source of the security annotations to the binding handles.

The CDS, in fact, only provides partial binding information. (*Lendenmann* 184.) During the RPC process, the client first sends a request to the CDS, and the CDS returns a "partly bound" binding handle. (*Id.* at 190.) Then, as *Lendenmann* explains, a *client* desiring to establish an RPC with certain security measures must first "specify . . . the authentication service, protection level and authorization service that it wants to use in its communications with a server." (*Id.* at 191.) The CDS does not provide this information or even provide any indication that it supports establishing a secure communication link; rather, the *client* provides the security-related information to the incomplete binding handle received from the CDS. (*Id.*; Keromytis Decl. ¶ 32.) Specifically, "[t]he client [places] a call to rpc#binding#set#auth#info(), which adds this security information to the server binding handle." (*Id.*) Afterwards, "[t]he client then uses this extended binding handle in its further RPC calls." (*Id.*) Accordingly, the alleged domain name service system, the CDS, simply



New Bay Capital, LLC
Ex.1015-Page 2541 of 3151

returns identifying information for the target server, and the client must supplement this incomplete binding information in order to implement any desired security measures for communications.

Thus, the CDS, by providing incomplete binding handles containing only identifying information for a server, performs no functions beyond those that the '504 patent distinguishes as characterizing a conventional domain name service, discussed above. Accordingly, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS to disclose "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (Keromytis Decl. ¶ 32.)

**(iv) *Lendenmann*'s "Authentication Challenge" Fails to Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Office Action and Request assert that the authentication challenge disclosed by *Lendenmann* provides "an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (OA at 11; Req. Ex. F-1 at 16.) This is incorrect.

The *Lendenmann* passage quoted in the Request for the proposition that the CDS "sends an authentication challenge in response to a name service query," Section 10.4.4, does not actually involve the CDS in any fashion. (Keromytis Decl. ¶ 33.) Instead, this section discusses how a *client* and a *server* perform "mutual authentication" for running authenticated RPCs. (*Lendenmann* 193.) This occurs through the use of server keys, service tickets, and session keys, none of which are provided by or involve the CDS, which the Office Action and Request refer to as the alleged domain name service system. (*Id.* at 194.)

During the mutual authentication process, the client's RPC runtime component first requests a service ticket from the Security Service, which "contains the session key for the upcoming client/server communication." (*Id.*) After the client sends the session key to the server, the server challenges the client. (*Id.*) Then, if the client successfully matches the challenge criteria, "everything is set for the authenticated RPC." (*Id.*) Thus, while the Security Service assists the client and the server during this process, the CDS is conspicuously absent. (Keromytis Decl. ¶ 33.)

Accordingly, contrary to the assertions of the Request and the Office Action, the authentication challenge does not occur "in response" to any name service query. (*Id.*) Instead, it occurs in response to the client separately sending the session key to the target server, for which the network address is already known. (*Lendenmann* 194; Keromytis Decl. ¶ 33.) Moreover, the CDS plays no role in the mutual authentication between the client and the target server. As discussed



New Bay Capital, LLC
Ex.1015-Page 2542 of 3151

above, the CDS performs no functions beyond those consistent with the conventional domain name server systems distinguished by the '504 patent from "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." (*See, e.g.*, '504 patent 39:43-41:59; Keromytis Decl. ¶ 19, 28.) Thus, one of ordinary skill in the art would not have understood *Lendenmann*'s authentication challenges to disclose an indication that the alleged domain name service system, the CDS, supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶ 33.)

**(v) *Lendenmann*'s "Server Status" Counters Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Office Action and the Request also assert that the server status counters described in *Lendenmann* disclose "an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (OA at 11; Req. Ex. F-1 at 16-17.) This too is incorrect.

*Lendenmann* discloses that the CDS clerk, CDS server, and CDS clearinghouse all maintain a set of counters "to keep track of the operations performed since it was last started up." (*Lendenmann* 37.) An example display of such counters is shown below:

```
cdscp> show server
                        SHOW
                      SERVER
                          AT   1995-06-01-10:06:33
              Creation Time = 1995-05-31-15:24:20.077
          Future Skew Time = 0
           Read Operations = 6409
          Write Operations = 33
          Skulks Initiated = 7
          Skulks Completed = 7
Times Lookup Paths Broken = 0
          Crucial Replicas = 0
    Child Update Failures = 0
        Security Failures = 0
     Known Clearinghouses = /.../itscl.austin.ibm.com/evl_ch
```

(*Id.*) The Request asserts that the "Security Failures" counter provides an indication that the CDS supports establishing a secure communication link. (Req. Ex. F-1 at 17.)

But nowhere does *Lendenmann* describe the meaning of the "Security Failures" counter, or explain that this counter is related in any manner to a secure communication link. (Keromytis Decl. ¶ 35.) Moreover, this server counter does not provide any "indication" that the CDS itself "supports establishing a secure communication link," as recited in claim 1. (*Id.*) To the contrary, as discussed



New Bay Capital, LLC
Ex.1015-Page 2543 of 3151

above, the CDS itself only returns server identification information. (*See id.; Lendenmann* 21, "when given a name, CDS returns the network address of the named resource.")

Instead, *Lendenmann* provides for a separate Security Service to handle any steps for establishing security measures in communications. (*Lendenmann* 191-94.) These functions are entirely separate from the conventional name-returning features of the CDS. As one example, the authentication process between a client and server is managed by the Security Server, while the CDS plays no role whatsoever:

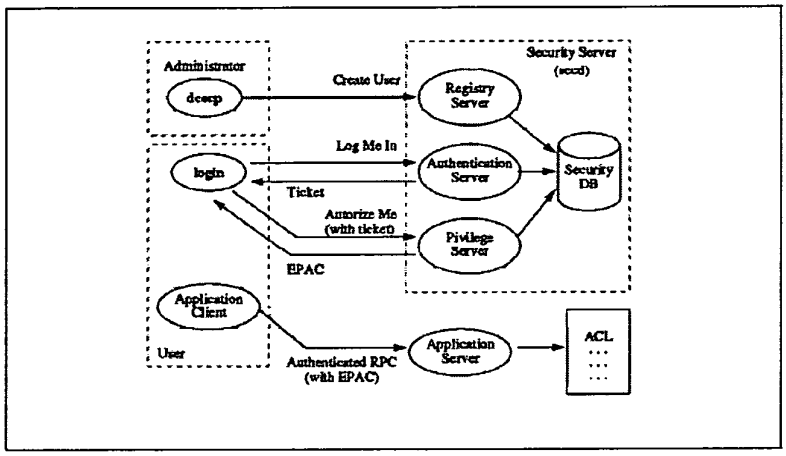


*Figure 21. Authentication Process.*

(*Id.* at 53-55; Keromytis Decl. ¶ 35.) The Security Service similarly handles any steps for any authorization or encryption measures, without CDS involvement. (*Lendenmann* at 191-94.)

Thus, the CDS and the server counters disclosed in *Lendenmann* perform no functions beyond those that the '504 patent distinguishes as characterizing a conventional domain name service. Accordingly, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS and server counters to disclose an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶ 35.)

**(vi) *Lendenmann*'s "Online Documentation" Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Request and the Office Action finally assert that *Lendenmann*'s "online documentation" provides "an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (OA at 11; Req. Ex. F-1 at 17-18.) This is incorrect. *Lendenmann* explains that this "online documentation" is nothing more than a collection of DCE manuals provided in softcopy form. (*Lendenmann* 14.)

### 1.6.4 Online Documentation

> All DCE manuals are provided in softcopy form to be accessed with a graphical viewer. The graphical softcopy files are INF files, which is the standard format for OS/2 online documentation. They can be accessed through the Interactive Presentation Facility (IPF).
>
> The IBM DCE Version 2.1 for AIX Version 4.1 provides an IPF viewer for X-Windows (IPF/X). The xview command that starts IPF/X provides hypertext linking, search and print facilities, inline graphics display, a bookmark function, and online help. Its startup is integrated into InfoExplorer. IBM DCE 2.1 for AIX also provides the documentation in ASCII from which can be viewed from ASCII terminals with an ASCII browser. The dceman command emulates MAN pages for DCE commands.
>
> On IBM DCE 1.3 for AIX, softcopy documentation is in InfoExplorer format.

(*Id.*) Other than the above passage, *Lendenmann* discloses nothing further about this "online documentation," failing to describe the content of these extrinsic and undisclosed manuals in any manner.

The Request and the Office Action fail to make out a prima facie 35 U.S.C. § 102 rejection based on this "online documentation." *See Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236 (Fed. Cir. 1989) ("The identical invention must be shown in as complete detail as is contained in the . . . claim."). The Request simply asserts that *Lendenmann* shows the elements of claim 1 because this "online documentation for the DCE software system would describe the system's services." (Req. Ex. F-1 at 17-18.) But this is not what the claim recites. Claim 1 recites "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." *Lendenmann* does not disclose this indication, much less explain why or how the various undisclosed DCE manuals, which are neither included, identified, nor incorporated in *Lendenmann*, somehow provide the "indication" recited in claim 1. The mere presence of "online documentation" manuals having indeterminate content does not indicate anything about the CDS, much less that the CDS is configured to support establishing a secure communication link. (Keromytis Decl. ¶ 36.) Accordingly, one of ordinary skill in the art would not have understood *Lendenmann*'s "online documentation" to disclose an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (*Id.*)

Moreover, to the extent that the Request and the Office Action rely on the speculative and unspecified content of the "online documentation" manuals, it is improper to rely on this extrinsic "online documentation" to support the rejection under 35 U.S.C. § 102 because the extra references are not cited to (1) prove *Lendenmann* contains an enabled disclosure, (2) explain the meaning of a

term used in *Lendenmann*, or (3) show that a characteristic not disclosed in *Lendenmann* is inherent. *See* M.P.E.P. § 2131.01.

For all of these reasons, the rejection of claim 1 under § 102 should be withdrawn, and its patentability confirmed.

**b. Independent Claims 36 and 60**

Independent claims 36 and 60 include recitations similar to those described above with respect to claim 1. For example, claim 36 recites "instructions executable in a domain name service system, the instructions comprising code for: . . . storing a plurality of domain names and corresponding network addresses; receiving a query for a network address; and supporting an indication that the domain name service system supports establishing a secure communication link." Claim 60 recites, for example, "storing a plurality of domain names and corresponding network address," "receiving a query for a network address," and "the domain name service system comprising an indication that the domain name service system supports establishing a secure communication link." Thus, for reasons similar to those discussed above with respect to independent claim 1, *Lendenmann* fails to anticipate claims 36 and 60. Accordingly, for similar reasons, Patent Owner requests that the rejection of claims 36 and 60 under 35 U.S.C. § 102 be withdrawn, and the patentability of the claims be confirmed.

**c. Dependent Claims 5, 23, and 47**

Dependent claims 5, 23, and 47 depend from independent claim 1 or 36, and include all of its features. Thus, *Lendenmann* does not anticipate these claims, and the rejection of these claims should be withdrawn, at least for the reasons discussed above in connection with independent claims 1 and 36. Claims 5, 23, and 47 also distinguish over *Lendenmann* for additional reasons. For example, dependent claim 5 recites that "the domain name service system is configured to authenticate the query [for a network address] using a cryptographic technique." Similarly, dependent claim 23 recites that "the domain name service system is configured to authenticate the query for the network address," and dependent claim 47 recites that "the instructions [executable in a domain name service system] comprise code for authenticating the query for the network address." Because the portions of *Lendenmann* cited by the Request and the Office Action fail to disclose these features of claims 5, 23, and 47, the rejections of these claims should be withdrawn for this additional reason.

First, *Lendenmann* fails to disclose a "domain name service system . . . configured to authenticate the query [for a network address]," as recited in claim 5 and similarly recited in claims



New Bay Capital, LLC
Ex.1015-Page 2546 of 3151

23 and 47. The Request asserts that by only performing operations for users authorized by access controls lists (ACLs), the CDS provides the authentication feature of claims 5, 23, and 47. (Req. Ex. F-1 at 19-20, 36-37, 55.) This is incorrect.

*Lendenmann* discloses that the Security Service—not the CDS—performs the alleged authentication process in acting as a gatekeeper to the information in the CDS. (Keromytis Decl. ¶ 62.) Indeed, "[t]he CDS server only completes an operation over the clearinghouse if the user is authenticated and authorized *by the Security Service*." (*Lendenmann* 34, emphasis added.) Even if a client has authorization and authentication to access the CDS, the CDS still merely "returns the network address of the named resource" when given a name in a request, as illustrated in Figure 15 below. (*Id.* at 21.)

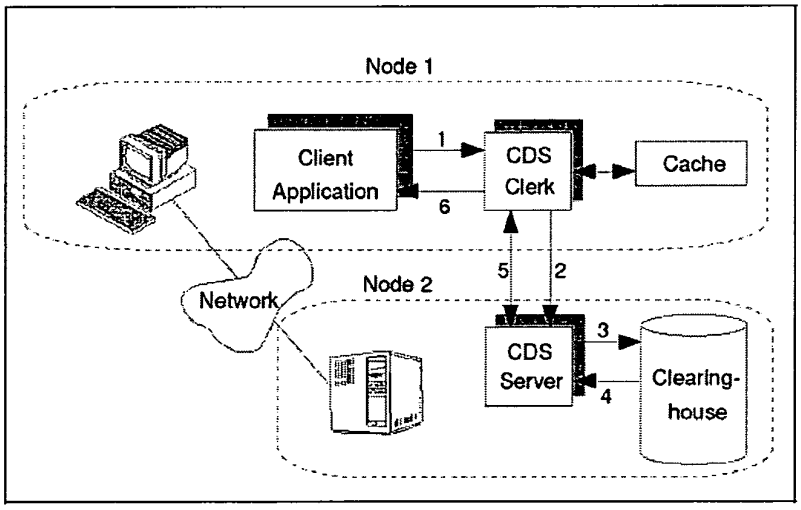


*Figure 15. CDS Components Performing a CDS Look-up*

(*Id.* at 29.) Thus, although the separate Security Service component may permit or deny access to the CDS, the CDS still only returns server identification information, consistent with the conventional domain name servers distinguished by the '504 patent, as discussed above. Accordingly, *Lendenmann* does not disclose that the CDS authenticates queries, and the rejection of claims 5, 23, and 47 under 35 U.S.C. § 102 should be withdrawn, and their patentability confirmed.

Additionally, the Request incorrectly asserts that the encryption features of *Lendenmann* disclose authenticating a query "using a cryptographic technique," as recited in claim 5. (Req. Ex. F-1 at 20-22.) The Request cites to two sections of *Lendenmann* as allegedly disclosing this feature. (*Id.*, citing *Lendenmann* 54-55, 193.) But not only does *Lendenmann* fail to disclose that the CDS authenticates any query as discussed above, neither of these two sections discusses the CDS. (Keromytis Decl. ¶ 63.) Instead, the cited passages on pages 54-55 of *Lendenmann* illustrate the



New Bay Capital, LLC
Ex.1015-Page 2547 of 3151

security measures implemented by the Security Service for RPCs, without having any bearing on network address queries. (*Lendenmann* 54-55; Keromytis Decl. ¶ 63.) This passage fails to discuss the CDS at all: all of these procedures are performed by the separate Security Service component. (*Lendenmann* 54-55; Keromytis Decl. ¶ 63.) Similarly, the cited passage on page 193 also illustrates the security measures performed by the Security Service, once again not only failing to discuss queries for network addresses, but also failing to discuss the CDS. (*Lendenmann* 193; Keromytis Decl. ¶ 63.)

Thus, *Lendenmann* has not been shown to disclose a "domain name service system [that] is configured to authenticate the query [for a network address] using a cryptographic technique," as recited in claim 5 and similarly recited in claims 23 and 47. Accordingly, Patent Owner requests that the rejection of these claims under 35 U.S.C. § 102 be withdrawn, and their patentability confirmed.

**d. Dependent Claims 16, 17, 27, 33, 40, 41, 51, and 57**

As discussed above, *Lendenmann* does not disclose "a domain name service system . . . configured to comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1 and similarly recited in claims 36 and 60. Because *Lendenmann* does not disclose this feature for the reasons discussed above, neither does it disclose a domain name service system configured to either establish or support establishing a secure communication link, as recited in claims 16 and 40. For the same reasons, neither does *Lendenmann* disclose a domain name service system that "comprises an indication that the domain name service system supports establishing a secure communication link," as recited in claim 17 and similarly recited in claim 41, or a domain name service system configured to "enable establishment of a secure communication link," as recited in claims 27, 33, 51, and 57.

The Response points to one or more of the six features of *Lendenmann* that it relied on for the proposed rejections of claim 1 as also allegedly disclosing the features of each of claims 16, 17, 40, and 41. (*See* Req. Ex. F-1 at 27, 29, 50-51.) Accordingly, for reasons similar to those discussed above as to why those alleged features of *Lendenmann* fail to disclose or suggest a domain name service that is configured to "comprise an indication that the domain name service system supports establishing a secure communication link," those features of *Lendenmann* also fail to disclose the recited features of claims 16, 17, 40, and 41.

Meanwhile, the Request cites Figure 68 of *Lendenmann* as disclosing that the CDS is a domain name service system "configured to enable establishment of a secure communication link," as recited in claims 27, 33, 51, and 57. (*See id.* at 40, 44, 56, 57.)



New Bay Capital, LLC
Ex.1015-Page 2548 of 3151

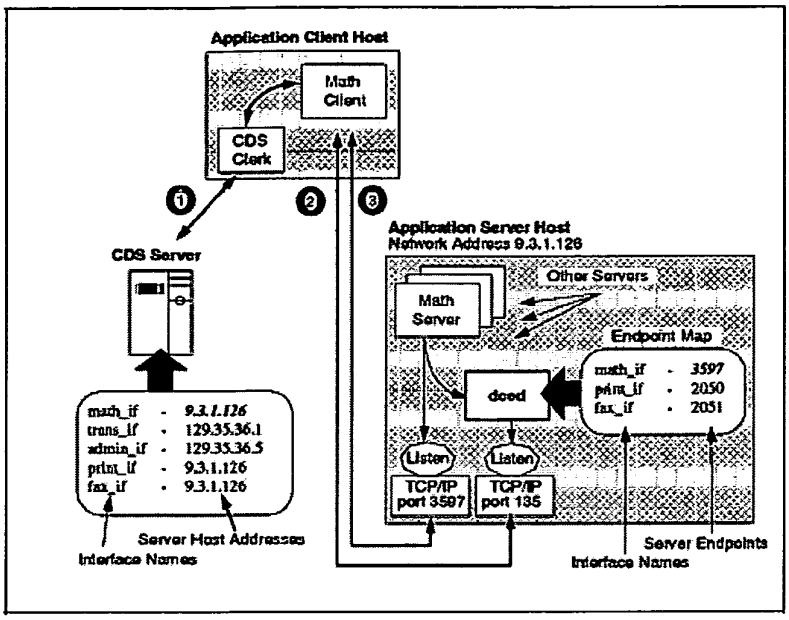


*Figure 68. Steps Involved in Finding a Server*

(*Lendenmann* 190.) But as *Lendenmann* explains with respect to Figure 68, the CDS merely returns incomplete binding handles containing server identification information, reinforcing that *Lendenmann*'s CDS is consistent with the conventional domain name servers distinguished by the '504 patent. (*Id.* at 190-91; *see also id.* at 182-84, explaining the incomplete binding handles.) Rather, the *client* must "add[] . . . security information to the server binding handle" received from the CDS if it desires to establish security measures for an RPC. (*Id.* at 191.) Thus, because the CDS performs no functions beyond those that the '504 patent distinguishes as characterizing a conventional domain name service, *Lendenmann* does not disclose a domain name service system "configured to enable establishment of a secure communication link."

Accordingly, *Lendenmann* does not disclose the features of claims 16, 17, 27, 33, 40, 41, 51, and 57, and the rejection of these claims should be withdrawn and their patentability confirmed.

**e. Dependent Claims 24 and 48**

Dependent claims 24 and 48 depend from independent claims 1 and 36, respectively, and include all of their features. Thus, *Lendenmann* does not anticipate these claims, and the rejection of these claims should be withdrawn, at least for the reasons discussed above in connection with the independent claims. Claims 24 and 48 also distinguish over *Lendenmann* for additional reasons. For example, dependent claim 24 recites, "wherein at least one of the plurality of domain names comprises an indication that the domain name service system supports establishing a secure communication link," and dependent claim 48 similarly recites the elements of independent claim 36,

"wherein at least one of the plurality of domain names includes an indication that the domain name service system supports the establishment of a secure communication link." *Lendenmann* does not disclose these claim features.

The Request contends that DCE names containing a "sec" portion as an abbreviation for "security" discloses these claim features because this "sec" portion "indicates that the server is a security server." (Req. Ex. F-1 at 37-38.) But this is not what the claims recite. Claim 24, for example, recites that "at least one of the plurality of domain names comprises an indication that the *domain name service system* supports establishing a secure communication link" (emphasis added), as similarly recited by claim 48. The Request contends that the "sec" portion "indicates that the server is a security server," but the security status of a particular server indicates nothing regarding the capabilities of the alleged *domain name service system*, the CDS, and certainly would not indicate that *it* supports establishing a secure communication link.

Accordingly, *Lendenmann* does not disclose the features of claims 24 and 48, and the rejection of these claims should be withdrawn and their patentability confirmed.

**f. Dependent Claims 2, 3, 6, 14, 15, 18-22, 25, 26, 28-30, 34, 35, 37-39, 42-46, 49, 50, 52-54, 58 and 59**

Remaining claims 2, 3, 6, 14, 15, 18-22, 25, 26, 28-30, 34, 35, 37-39, 42-46, 49, 50, 52-54, 58 and 59 depend from one of independent claims 1 and 36, and include all of their features. Thus, *Lendenmann* does not anticipate any of these claims for at least the reasons discussed above in conjunction with independent claims 1 and 36. For the reasons set forth above, the rejection of these claims under 35 U.S.C. § 102 based on *Lendenmann* should be withdrawn and their patentability confirmed.

**3. Rejection of Claims 1-6, 14-30, 33-54, and 57-60 Under 35 U.S.C. § 103(a) Based on *Lendenmann* (Issue 2)**

The Office Action rejects claims 1-3, 5, 6, 14-30, 33-54, and 57-60 under 35 U.S.C. § 103(a) as being obvious over *Lendenmann*. (OA at 5.) However, the Request's and the Office Action's analysis of these claims is deficient. "The key to supporting any rejection under 35 U.S.C. 103 is the clear articulation of the reason(s) why the claimed invention would have been obvious." M.P.E.P. § 2142. "[R]ejections on obviousness cannot be sustained with mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness." *Id.* (internal quotations omitted).

The Request and the Office Action fail to provide such a required articulated reasoning. Instead, the Request's and the Office Action's analysis in the § 103(a) rejection is *identical* to its



New Bay Capital, LLC
Ex.1015-Page 2550 of 3151

analysis in the § 102(b) rejection for these claims, with the sole difference being that the word "teaches" has been globally replaced with the words "renders obvious." Patent Owner was unable to identify any other differences in the analysis.

The Request and the Office Action do not explain how *Lendenmann* renders obvious those features that *Lendenmann* fails to disclose, as shown above in the traversal of the § 102(b) rejections based on *Lendenmann* (Issue 1). Thus, Patent Owner incorporates the arguments for why *Lendenmann* does not disclose the recited features of claims 1-3, 5, 6, 14-30, 33-54, and 57-60 that were discussed above with regard to the § 102(b) rejections based on *Lendenmann*. And Patent Owner further asserts that *Lendenmann* does not suggest or otherwise render obvious these missing features. Nor do the Request and the Office Action provide any additional explanation as to why *Lendenmann* allegedly renders these missing features obvious.

The Office Action also rejects claim 4 as being obvious over *Lendenmann*. Claim 4 depends on independent claim 1 and includes all of its features. As explained above, *Lendenmann* fails to disclose or suggest the features of claim 1 and thus does not support the rejection for that claim. Because the additional arguments presented with respect to claim 4 do not attempt to address *Lendenmann*'s shortcomings in disclosing or suggesting the elements of claim 1, *Lendenmann* also does not suggest or render obvious claim 4.

In view of the above, the rejection of claims 1-6, 14-30, 33-54, and 57-60 under § 103(a) as being obvious over *Lendenmann* should be withdrawn and the claims should be confirmed.

#### 4. Rejection of Claim 7 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Wesinger* (Issue 3)

Claim 7 depends from independent claim 1. As explained above, *Lendenmann* does not disclose or suggest the features of claim 1 and thus does not support the rejection of that claim under 35 U.S.C §§ 102 and 103. The Request does not cite to U.S. Patent No. 5,898,830 to Wesinger et al. ("*Wesinger*") to remedy the above noted deficiencies of *Lendenmann* with respect to claim 1, instead employing *Wesinger* to suggest that the CDS of *Lendenmann* may be placed at a firewall. (Req. Ex. F-1 at 120.) But regardless of where the CDS resides, by merely returning server identification information to the client as discussed above, the CDS would simply continue to perform no functions beyond those that the '504 patent recognizes and distinguishes as characterizing a conventional domain name service. (*See Lendenmann* 21, 29; '504 patent 39:7-13, 39:34-42.) *Wesinger* does not disclose or suggest, and the Request and the Office Action do not rely upon *Wesinger* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." As such, the asserted



New Bay Capital, LLC
Ex.1015-Page 2551 of 3151

combination fails to disclose or remotely suggest the limitations of claim 7. Thus, for at least these reasons, the § 103(a) rejection of claim 7 over *Lendenmann* in view of *Wesinger* should be withdrawn.

### 5. Rejection of Claims 8 and 9 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* (Issue 4)

Dependent claim 8 depends from independent claim 1, and dependent claim 9 depends from claim 1 via claim 8. As explained above, *Lendenmann* does not disclose or suggest the features of claim 1 and thus does not support the rejection of claims 8 and 9. The Request does not cite to Jean-Paul Gaspoz, "VPN on DCE: From Reference Configuration to Implementation" ("*Gaspoz*") to remedy the deficiencies of *Lendenmann* with respect to claim 1, instead relying on *Gaspoz* to support its obviousness arguments concerning the additional element "wherein the domain name service system is connectable to a virtual private network through the communication network." (Req. Ex. F-1 at 124-30.) Accordingly, *Gaspoz* does not disclose or suggest, and the Request and the Office Action do not rely upon *Gaspoz* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Accordingly, *Lendenmann* in view of *Gaspoz* fail to disclose or suggest all of the limitations of claims 8 and 9. Consequently, the § 103(a) rejections of claims 8 and 9 over *Lendenmann* in view of *Gaspoz* should be withdrawn.

#### a. Claim 8

Additionally, *Lendenmann* and *Gaspoz*, either alone or in combination, fail to disclose or suggest that "the domain name service system is connectable to a virtual private network through the communication network," as recited in claim 8. The Request asserts that because a client in one DCE cell can engage in shared-secret key authentication with a client in a foreign DCE cell, *Lendenmann* discloses a domain name service system "connectable to a virtual private network." (*Id.* at 125-27.) But this is incorrect.

*Lendenmann* discloses that clients in different cells can communicate with each other while employing the secret key authentication feature. (*Lendenmann* 68.) With this feature, *Lendenmann* illustrates that the Security Server—not the CDS—manages any security measures for subsequent communications, including any authentication, authorization, or encryption. (*Id.* at 53-55; *see also id.* at 191-94.) As one example, the Security Server authenticates users without any action by or reference to a CDS, as illustrated in Figure 21 below:



New Bay Capital, LLC
Ex.1015-Page 2552 of 3151

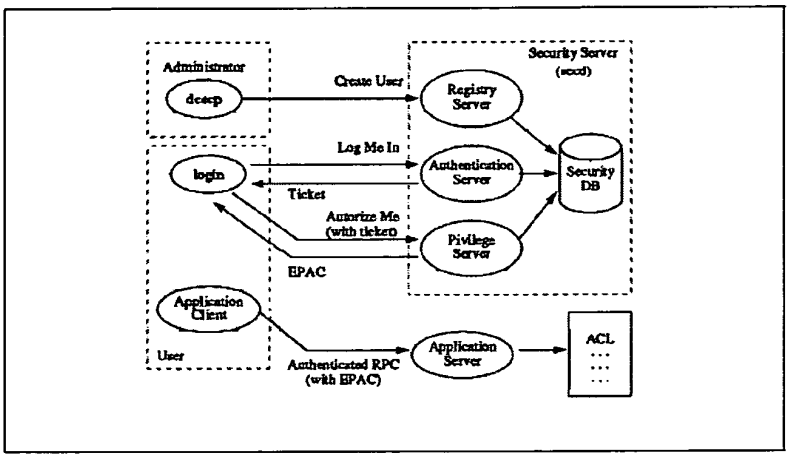


*Figure 21. Authentication Process.*

(*Id.* at 53.) Indeed, throughout this authentication process, *Lendenmann*'s CDS, the alleged domain name server, is conspicuously absent.

Attempting to remedy this shortcoming, the Request asserts that because a CDS is "a required part of a DCE 'cell'," and because clients within one cell can communicate with another cell, this satisfies the claim element of a "domain name service system . . . connectable to a virtual private network through the communication network." (Req. Ex. F-1 at 125-26.) A person of ordinary skill in the art, however, would not have interpreted the entire DCE cell of *Lendenmann* to be the recited "domain name service system." (Keromytis Decl. ¶ 68.) Indeed, *Lendenmann* distinguishes between the subcomponents of a cell and their capabilities, explaining that, "[a]t a minimum, a cell must contain a Security Server, a Cell Directory Server and Distributed Time Servers." (*Lendenmann* 9; Keromytis Decl. ¶ 68.) Simply because a CDS is a component of a cell does not mean that the entire cell is a "domain name service system," as recited in claim 8. Rather, as discussed above, the *Lendenmann* provides that the Security Server facilitates and manages security measures for security systems, while the CDS performs no functions beyond those that the '504 patent distinguishes from conventional domain name servers. (*See Lendenmann* 21; '504 patent 39:7-13, 39:34-42; Keromytis Decl. ¶ 68.) Accordingly, *Lendenmann* fails to disclose or suggest a "domain name service system . . . connectable to a virtual private network through the communication network."

Neither does *Gaspoz* disclose or suggest the additional element of claim 8. The Request asserts that because *Gaspoz* discusses implementing a VPN in conjunction with an OSF DCE platform, "it would have been obvious to connect the virtual private network to the Cell Directory Service (CDS)." (Req. Ex. F-1 at 127.) This is incorrect. *Gaspoz* actually discloses that

implementing a VPN on a DCE was extraordinarily difficult because of a computer language discrepancy. (*Gaspoz* 257-58; Keromytis Decl. ¶ 69.) Although the Request claims that connecting a CDS to a VPN would have been obvious because *Gaspoz* discloses that its objects connected to a VPN "were thoroughly represented as DCE servers," *Gaspoz* admits that the object-oriented computer language (C++) made it difficult to utilize objects in establishing a VPN. (*Gaspoz* 257-58; Keromytis Decl. ¶ 69.) Indeed, *Gaspoz* had to resort to manual file creation to create objects, since "[w]e found no other means for creating persistent objects, because DCE lacks supporting tools in this area." (*Gaspoz* 257-58; Keromytis Decl. ¶ 69.)

Moreover, *Gaspoz* recognized that OSF DCE utilizes a CDS, and instead employed *UserAddr* and *DialList* objects to designate and permit look-ups of the addresses for various other objects within a VPN. (*Gaspoz.* at 256-57.) Thus, *Gaspoz* recognized the limitations of employing a VPN together with a DCE, and disclosed using an alternative to a CDS for assisting in establishing VPNs. (*Id.* at 256-59.) And because the CDS disclosed in *Lendenmann* performs no functions other than returning server identification information, which the '504 patent distinguishes as conventional, there would have been no reason to combine *Lendenmann* and *Gaspoz* because *Gaspoz* had already considered the CDS component of the OSF DCE and developed an alternative with the *UserAddr* and *DialList* objects. (*Id.* at 256-57; *see also Lendenmann* 21; '504 patent 39:7-13, 39:34-42; Keromytis Decl. ¶ 69-70.) Thus, for at least these reasons, the § 103(a) rejection of claim 8 over *Lendenmann* in view of *Gaspoz* should be withdrawn.

**b. Claim 9**

Claim 9 depends from independent claim 1 via claim 8. As explained above, *Lendenmann* does not disclose or suggest the features of claim 1 or 8, and therefore does not support the rejection of those claims. Furthermore, as discussed above, *Gaspoz* does not remedy the deficiencies of *Lendenmann* with respect to claim 8, and the Request does not rely on *Gaspoz* to remedy any deficiencies with respect to claim 1. Thus, *Lendenmann* in view of *Gaspoz* does not disclose or suggest claim 9 for at least the reasons discussed in conjunction with claims 1 and 8, and, accordingly, the rejection of claim 9 under 35 U.S.C. § 103(a) should be withdrawn.

**6. Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* and *Schneier* (Issue 5)**

Claim 10 depends from claim 8, which in turn depends from claim 1. As explained above, *Lendenmann* does not disclose or suggest the features of claim 1, and *Lendenmann* in view of *Gaspoz* does not disclose or suggest the features of claim 8. Furthermore, Bruce Schneier, "Applied Cryptography" ("*Schneier*") does not disclose or suggest, and the Request and the Office Action do



New Bay Capital, LLC
Ex.1015-Page 2554 of 3151

not rely on *Schneier* as allegedly disclosing or suggesting, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1, or that "the domain name service system is connectable to a virtual private network through the communication network," as further recited in claim 8. Accordingly, the rejection of claim 10 over *Lendenmann* in view of *Gaspoz* and *Schneier* should be withdrawn because *Schneier* does not remedy the deficiencies of *Lendenmann* and *Gaspoz* discussed above with respect to claims 1 and 8.

### 7. Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* and *Martin* (Issue 6)

Dependent claim 11 recites that "the virtual private network is based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." *Lendenmann* and *Martin*, either alone or in combination, do not disclose or suggest this feature.

Claim 11 depends from claim 8, which in turn depends from claim 1. As explained above, *Lendenmann* does not disclose or suggest the features of claim 1, and *Lendenmann* in view of *Gaspoz* does not disclose or suggest the features of claim 8. Furthermore, David M. Martin, "A Framework for Local Anonymity in the Internet" ("*Martin*") does not disclose or suggest, and the Request and the Office Action do not rely on *Martin* as allegedly disclosing or suggesting, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1, or that "the domain name service system is connectable to a virtual private network through the communication network," as further recited in claim 8. Accordingly, the rejection of claim 11 over *Lendenmann* in view of *Gaspoz* and *Martin* should be withdrawn because *Martin* does not remedy the deficiencies of *Lendenmann* and *Gaspoz* discussed above with respect to claims 1 and 8.

Additionally, *Lendenmann* does not disclose a "virtual private network . . . based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." Nor do the Request and the Office Action assert that it does. Instead, they point to the binding handles feature of *Lendenmann*, wherein a client may receive multiple binding handles for connecting to an application server because multiple server machines may support the same application servers. (Req. Ex. F-1 at 136-37, citing *Lendenmann* 185-86.) To resolve the situation of having to potentially choose between several binding handles for reaching the same application server, *Lendenmann* indicates that the client may use a random selection. (*Lendenmann* 185.) But this fails to disclose that a "network address hopping regime" is used to perform this random



New Bay Capital, LLC
Ex.1015-Page 2555 of 3151

selection of network addresses, or that a "virtual private network is based on" such a hopping regime using this random selection.

*Martin* does not make up for these deficiencies of *Lendenmann*. In fact, the Request and the Office Action do not even assert that *Martin* discloses the claimed feature. Instead, the Request simply asserts that "[c]hoosing one of the source addresses at random" shows the elements of claim 11. (Req. Ex. F-1 at 137.) But this is not what the claim recites. Claim 11 recites that "the virtual private network is based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." *Martin* does not disclose pseudorandomly changing network addresses in packets, let alone a network address hopping regime that is used to pseudorandomly change network addresses in packets. (*See Martin* 9.) Because the Request and the Office Action do not even assert that *Martin* discloses this feature or explain how the combination of *Lendenmann* and *Martin* would render this missing feature obvious, the rejection is improper on its face and should be withdrawn. Moreover, because *Lendenmann* and *Martin* do not disclose or suggest the features of claim 11, either alone or in combination, the patentability of claim 11 should be confirmed.

### 8. Rejection of Claims 12 and 13 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Gaspoz* and RFC 793 (Issue 7)

The Office Action rejects claims 12 and 13 under 35 U.S.C. § 103(a) based on *Lendenmann* in view of *Gaspoz* and RFC 793. (OA at 6.)

#### a. Claim 12

Dependent claim 12 recites that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values." *Lendenmann* and Information Sciences Institute, "Transmission Control Protocol," RFC 793 ("RFC 793"), either alone or in combination, do not disclose or suggest this feature.

Claim 12 depends from claim 8, which in turn depends from claim 1. As explained above, *Lendenmann* does not disclose or suggest the features of claim 1, and *Lendenmann* in view of *Gaspoz* does not disclose or suggest the features of claim 8. Furthermore, RFC 793 does not disclose or suggest, and the Request and the Office Action do not rely on RFC 793 as allegedly disclosing or suggesting, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1, or that "the domain name service system is connectable to a virtual private network through the communication network," as further recited in claim 8. Accordingly, the rejection of claim 12 over *Lendenmann* in view of *Gaspoz* and RFC 793 should be withdrawn because RFC 793 does not remedy the deficiencies of *Lendenmann* and *Gaspoz* discussed above with respect to claims 1 and 8.



New Bay Capital, LLC
Ex.1015-Page 2556 of 3151

Additionally, *Lendenmann* does not disclose that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values." Nor do the Request and the Office Action assert that it does.

RFC 793 does not remedy the deficiencies of *Lendenmann*. In fact, the Request and the Office Action do not even assert that RFC 793 discloses the claimed feature. Instead, the Request asserts that "*Lendenmann*'s remote procedure call, which operates over TCP, *uses* 'a moving window of valid values' as indicated by RFC 793." (Req. Ex. F-1 at 140, emphasis added.) But merely using a moving window is not what the claim recites. Claim 12 recites that "*the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device* to a moving window of valid values" (emphasis added). Because the Request and the Office Action do not even assert that RFC 793 discloses "a virtual private network . . . based on comparing a value in each data packet . . . to a moving window of valid values," or explain how the combination of *Lendenmann* and RFC 793 would render this missing feature obvious, the rejection is improper on its face and should be withdrawn. Moreover, because *Lendenmann* and RFC 793 do not disclose or suggest the features of claim 12, either alone or in combination, the patentability of claim 12 should be confirmed.

**b. Claim 13**

Claim 13 depends from claim 8, which in turn depends from claim 1. As explained above, *Lendenmann* does not disclose or suggest the features of claims 1 and 8, and thus does not support the rejection of those claims. Also, as discussed above, *Gaspoz* does not remedy the deficiencies of *Lendenmann* with respect to these claims. Thus, the rejection of claim 13 over *Lendenmann* in view of *Gaspoz* and RFC 793 should be withdrawn, because RFC 793 does not make up for the deficiencies of *Lendenmann* and *Gaspoz* with respect to claims 1 and 8. For instance, RFC 793 does not disclose or suggest, and the Request and the Office Action do not rely on RFC 793 as allegedly disclosing or suggesting, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," or that "the domain name service system is connectable to a virtual private network through the communication network." Accordingly, the rejection of claim 13 over *Lendenmann* in view of *Gaspoz* and RFC 793 should be withdrawn, and the patentability of claim 13 should be confirmed.

**9. Rejection of Claims 31, 32, 55, and 56 Under 35 U.S.C. § 103(a) Based on *Lendenmann* in View of *Ludwig* and RFC 793 (Issue 8)**

Claims 31 and 32 depend from independent claim 1, and claims 55 and 56 depend from independent claim 36. As explained above, *Lendenmann* does not disclose or suggest the features of



New Bay Capital, LLC
Ex.1015-Page 2557 of 3151

claims 1 and 36, and thus does not support the rejection of those claims. The rejection of claims 31, 32, 55, and 56 should also be withdrawn because U.S. Patent No. 5,689,641 to Ludwig et al. ("*Ludwig*") and RFC 793, either alone or in combination, do not remedy the deficiencies of *Lendenmann* with respect to independent claims 1 and 36. For instance, neither *Ludwig* nor RFC 793 discloses or suggests, and the Request and the Office Action do not rely upon *Ludwig* or RFC 793 to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," or "instructions executable in a domain name service system, the instructions comprising code for: . . . storing a plurality of domain names and corresponding network addresses; receiving a query for a network address; and supporting an indication that the domain name service system supports establishing a secure communication link." Thus, for at least these reasons, the rejection of claims 31, 32, 55, and 56 over *Lendenmann* in view of *Gaspoz* and RFC 793 should be withdrawn.

**B. The Rejections Based on *Aziz* Should Be Withdrawn**

The Office Action rejects certain claims under 35 U.S.C. §§ 102 and/or 103 based on U.S. Patent No. 6,199,234 to Aziz et al. ("*Aziz*") alone and/or in combination with one or more secondary references. However, as discussed below, the rejections under §§ 102 and/or 103 based on these references should be withdrawn and the claims should be confirmed.

**1. Overview of *Aziz***

*Aziz* discloses a system "for dynamically configuring authorized clients with the address of a protected host and the key and address of an intermediate device (e.g., encrypting firewall, encrypting router, secure gateway) which is protecting a number of hosts on a private network" behind the intermediate device. (*Aziz* 4:3-9.) Fig. 1 of *Aziz*, reproduced below, discloses such a system:

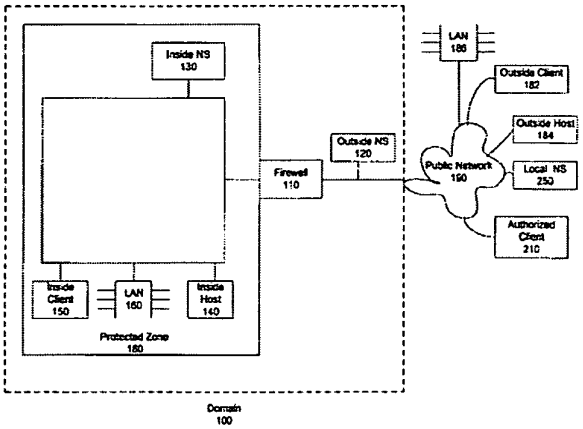




New Bay Capital, LLC
Ex.1015-Page 2558 of 3151

*Aziz* explains that "outside NS" 120 may receive a query for a host address located within domain 100 and may determine whether an SX record exists for that host name. (*Id.* at 9:49-53.) An SX record is a DNS resource record that "contains the identifier (e.g., name or address) of a 'secure exchanger,'" such as firewall 110. (*Id.* at 6:23-40.) If an SX record exists, then outside NS 120 may include the SX record in the response to the requester, which may also include the requested host address, if available. (*Id.* at 9:54-10:5.) *Aziz* also discloses that SIG (signature) and KEY resource records may be included in the response. (*Id.* at 9:35-41.)

*Aziz* also discloses a resolver 225, which is included in the "authorized client" 210 (*id.* at 8:5-50, Figs. 2A-2C), and receives a response to the query for a host address (*id.* at 10:39-41). If the response includes an SX record and the requested host address, then resolver 225 creates a tunnel map entry that provides the information "authorized client" 210 needs to encrypt messages to "inside host" 140. (*Id.* at 11:13-60.) Resolver 225 then returns the requested host address to an application 215, also located in "inside host" 140. (*Id.* at 11:55-60.) According to *Aziz*, "[t]his completes the execution" of the configuration process. (*Id.* at 11:60-62.)

### 2. Rejection of Claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 102(b) Based on *Aziz* (Issue 9)

The Office Action rejects claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 under 35 U.S.C. § 102(b) based on *Aziz*. (OA at 11.) This rejection is deficient and should be withdrawn for at least the reasons discussed below.

#### a. Independent Claim 1

*Aziz* fails to disclose the combination of features recited in claim 1 for at least the reasons discussed below. Independent claim 1 recites, among other things, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." The Request and the Office Action assert that four different features of *Aziz* disclose the recited indication. These assertions are incorrect because each asserted feature in *Aziz*, discussed in turn below, does not disclose a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link. Moreover, these features of *Aziz* disclose nothing more than a conventional domain name service system that is both recognized and distinguished by the '504 patent.

Before discussing the four different features of *Aziz* relied on by the Request, Patent Owner notes that Requester asserts in one portion of the Request that seemingly all of the elements shown in Fig. 1 of *Aziz* are included in the alleged domain name service system. (Req. Ex. F-2 at 5.) However, such an expansive reading is not consistent with the plain and ordinary meaning, and the



New Bay Capital, LLC
Ex.1015-Page 2559 of 3151

broadest reasonable interpretation, of the term "domain name service system." One of ordinary skill in the art would not have understood all of these elements in *Aziz* to be included in the recited domain name service system. (Keromytis Decl. ¶ 40.) Thus, this statement by the Requester is not reasonable. Moreover, the Request, when considered as a whole and, especially with regard to the analysis of the element that the Request identifies as [1.5], appears to assume that NS 120 is the recited domain name service system. (*See, e.g.*, Req. Ex. F-2 at 7-11, relying on resource records such as SX, KEY, and SIG resource records that are stored in *NS 120* as being the recited "indication.")

Further, the Office Action takes the position that the NS 120 of *Aziz* is the recited domain name service system. For example, the Office Action asserts that "[r]egarding the limitation 'domain name service [system] configured for connection [sic] to a communication network' . . . [s]ee also Fig. 1, reproduced below which illustrates the *outside name server 120* (NDS) [sic] connected to public network 190." (OA at 13, emphasis added.) The Office Action continues: "[r]egarding the limitation to provide an 'indication that the domain name service [system] supports establishing a secure communication[] link,' *Aziz* describes configuring the DNS to respond to requests with a special record that includes information needed for secure communications . . . . Thus, the presence of SX records in the response from the *DNS (NS 120)* provides an *indication that the DNS* [supports] establishing a secure communication link." (*Id.* at 15, emphases added.) Because the Request's assertion that essentially all of the elements in Fig. 1 of *Aziz* constitute a domain name service system is unreasonable and is belied by other portions of its analysis, and because the Office Action ultimately takes the position that the NS 120 is the recited domain name service system, Patent Owner's remarks below address the rejections in view of *Aziz* based on the Office Action's position that the NS 120 is the recited domain name service system.

**(i) *Aziz*'s SX Records Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

First, the Request and the Office Action assert that providing "SX records in the response from the DNS (NS 120) provides an indication that the DNS [supports] establishing a secure communication link." (*Id.*; *see also* Req. Ex. F-2 at 7-8.) This is incorrect. The SX record in *Aziz* is not an indication that the alleged domain name service system (NS 120) supports establishing a secure communication link. Instead, the SX record merely "contains the identifier (e.g., name or address) of a 'secure exchanger' [i.e., firewall 110] associated with the owner of the record." (*Aziz* 6:27-38.) Thus, the SX record includes the name or address of firewall 110, which is separate from



New Bay Capital, LLC
Ex.1015-Page 2560 of 3151

the alleged domain name service system, NS 120. (Keromytis Decl. ¶ 42; *see also Aziz* Fig. 1, showing NS 120 separate from firewall 110.) While including the name or address of firewall 110, the SX record includes no indication about the capabilities of the alleged domain name service system itself or about the capabilities of firewall 110, and certainly does not include an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶ 42.)

Indeed, returning an SX record with the name or address of firewall 110 in *Aziz* is a feature of a conventional domain name service system that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*, '504 patent 39:7-42; Keromytis Decl. ¶ 43.) As discussed, the '504 patent indicates that a conventional domain name service system merely returns an IP address or public key that was requested of it. (Keromytis Decl. ¶ 43.) For instance, the '504 patent explains that "[c]onventional Domain Name Servers (DNSs) provide a look-up function that *returns the IP address* of a requested computer or host. For example, when a computer user types in the web name 'Yahoo.com,' the user's web browser transmits a request to a DNS, which converts the name into a four-part IP address that is returned to the user's browser . . . ." ('504 patent 39:7-13, emphasis added; Keromytis Decl. ¶ 43; *see also* '504 patent 39:14-42.) Similar to the conventional domain name systems described by the '504 patent, the NS 120 in *Aziz* merely returns an SX resource record requested for a particular domain name that includes the *name or address* of a secure exchanger associated with that domain name. (*See, e.g.*, *Aziz* 9:49-56, 6:27-38; Keromytis Decl. ¶¶ 42-43.)

The '504 patent recognizes that such conventional domain name systems suffer from certain drawbacks and thus discloses embodiments that address them, including "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*, '504 patent 39:43-41:61; Keromytis Decl. ¶ 19.) And since returning a name or address is a feature of a conventional domain name server of the type distinguished by the '504 patent, one of ordinary skill in the art would not have understood *Aziz*'s SX record to disclose an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶¶ 42-43.)



New Bay Capital, LLC
Ex.1015-Page 2561 of 3151

**(ii) *Aziz*'s KEY and SIG Records Do Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Request and the Office Action assert that *Aziz* discloses a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link "by providing secure DNS service." (OA at 16; Req. Ex. F-2 at 9.) The explanation in the Request and the summary in the Office Action make clear that when referring to "providing secure DNS service," the Request and the Office Action are asserting that the SIG and KEY resource records in *Aziz* are the recited "indication." (*Id.* at 9-10; *see also* OA at 15, stating that "*Aziz* describes automatically adding the KEY and SIG records, which also provides 'an indication' . . . .") This is incorrect because the KEY and SIG records do not provide an indication that *the alleged domain name service system* (*NS 120*) supports establishing a secure communication link. Instead, KEY and SIG records provided in *Aziz* correspond to resource records. (*Aziz* 9:35-41, stating that whenever a resource record is added to a response, the "appropriate SIG and KEY records are also added (i.e., one SIG record for each record type and record owner combination and the KEY record used to generate the SIG record)"; Keromytis Decl. ¶ 45.) But these resource records returned in *Aziz* are A records and SX records that correspond to the inside host 140 and the firewall 110, respectively, which are both separate from the alleged domain name service system, NS 120. (*Aziz* 9:40-67, Fig. 1; Keromytis Decl. ¶ 45.) Thus, these KEY or SIG resource records in *Aziz* include no indication about the capabilities of *the alleged domain name service system itself*, and certainly do not include an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Keromytis Decl. ¶ 45.)

Moreover, returning KEY and SIG resource records in *Aziz* is consistent with a conventional domain name system that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*,'504 patent 39:7-42; Keromytis Decl. ¶ 43, 45.) As discussed, the '504 patent indicates that a conventional domain name system merely stores *public keys* of different machines so that hosts can request and receive those public keys from the domain name service system. ('504 patent 39:34-42; Keromytis Decl. ¶¶ 19, 45.) The '504 patent explains that an example of such a conventional system is disclosed in RFC 2535, which describes the very same KEY and SIG resource records disclosed in *Aziz*. ('504 patent 39:40-42; *see also* Ex. A-6 at 10-23, disclosing the KEY and SIG resource record types.) Thus, the NS 120 in *Aziz* merely returning KEY and SIG resource records is an aspect of a conventional system that the '504



New Bay Capital, LLC
Ex.1015-Page 2562 of 3151

patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*, '504 patent 39:7-42; Keromytis Decl. ¶¶ 19, 45.) Accordingly, the KEY and SIG resource records also are not an indication that the domain name service system supports establishing a secure communication link.

**(iii) *Aziz*'s "Information Used for Secure Communications with Protected Hosts" Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Request and the Office Action also assert that *Aziz* discloses the recited indication "by providing the information needed for secure communications between a client and a protected host." (OA at 16; Req. Ex. F-2 at 8-9.) To support this assertion, the Request cites two portions of *Aziz*. Neither portion of *Aziz* discloses a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link.

The first portion of *Aziz* cited by the Request states that "[t]he data in the SX record is used by a program called a resolver to update information used by a client for secure communications with protected hosts." (*Aziz* 6:57-60.) However, as discussed above, "the data in the SX record" is the name or address of the secure exchanger. (*Id.* at 6:27-38.) And, as discussed above, merely providing the SX record with this name or address in *Aziz* does not disclose an indication that the domain name service system supports establishing a secure communication link and is a conventional feature of a domain name service system that the '504 patent recognizes and distinguishes.

The second portion of *Aziz* cited by the Request states that the resolver "update[s] a data structure on a client containing information used for secure communications with protected hosts .... Such a data structure comprises data sets whose fields typically contain 'tunnel information' (e.g., destination and secure exchanger addresses) and related cryptographic data (e.g., secure exchanger's key or algorithm)." (*Id.* at 7:28-36.) The "data structure on a client containing information used for secure communications" does not disclose a *domain name service system* configured to comprise an indication that the domain name service system supports establishing a secure communication link for two reasons.

First, *Aziz* clearly discloses that "the data structure" is "on a client," which is separate from the alleged domain name service system, NS 120. (*Id.* at 7:30, Fig. 1.) Thus, "the data structure" does not disclose a *domain name service system* configured to comprise any indication. (Keromytis Decl. ¶ 48.) Second, *Aziz* discloses that the data structure includes "destination and secure exchanger



New Bay Capital, LLC
Ex.1015-Page 2563 of 3151

addresses" (i.e., A records and SX records) and secure exchanger keys or algorithms (i.e., SIG and KEY records for the secure exchanger). (*Id.*) As discussed above, however, providing these records in *Aziz* is a conventional domain name service system feature rather than an indication that the domain name service system supports establishing a secure communication link.

**(iv) *Aziz*'s Reference to RFC 2065 Does Not Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The fourth feature in *Aziz* relied on by the Request and the Office Action as allegedly teaching the recited indication is that *Aziz* refers to D. Eastlake and C. Kaufman, "Domain Name System Security Extensions," RFC 2065 ("RFC 2065"), directed to "Domain Name System - Security Extensions." (OA at 16; Req. Ex. F-2 at 10, citing *Aziz* 6:12-18.) According to the Request, this reference to RFC 2065 discloses "indicating that the connection to the domain name server itself can be encrypted," which, in turn, discloses a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. This is incorrect for the reasons discussed below.

First, the portion of *Aziz* relied on by the Request merely states that "[o]ne embodiment of the invention uses the KEY and SIG resource records provided by secure DNS," and that details of secure DNS are well known and described at least in RFC 2065. (*Aziz* 6:11-15.) Thus, the reference to RFC 2065 in *Aziz* is merely for introducing the KEY and SIG resource records, which, as discussed above, are not indications that the domain name service system supports establishing a secure communication link.

Second, the portion of RFC cited by the Request describes a single bit (Bit 8) within the KEY resource record in more detail. (RFC 2065 at 12.) But, because the KEY resource record cannot be the recited indication, as discussed above, a component of the KEY resource record also cannot be the recited indication.

In view of the above, all four portions of *Aziz* relied upon by the Request and the Office Action fail to disclose a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link. Indeed, the alleged domain name service system in *Aziz* (NS 120) merely returns requested resource records such as A, SX, KEY, and SIG resource records, which are not the claimed indication, but rather aspects of a conventional domain name service system distinguished by the '504 patent from the claimed domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link.



New Bay Capital, LLC
Ex.1015-Page 2564 of 3151

Based on the foregoing, *Aziz* fails to disclose the aforementioned features of claim 1, which therefore is not anticipated by *Aziz*. Accordingly, Patent Owner requests that the rejection of claim 1 under 35 U.S.C. § 102 be withdrawn, and the patentability of claim 1 be confirmed.

**b. Independent Claims 36 and 60**

Independent claims 36 and 60 include recitations similar to those described above with respect to claim 1. For example, claim 36 recites "instructions executable in a domain name service system, the instructions comprising code for: . . . storing a plurality of domain names and corresponding network addresses; receiving a query for a network address; and supporting an indication that the domain name service system supports establishing a secure communication link." And claim 60 recites, for example, "storing a plurality of domain names and corresponding network addresses," "receiving a query for a network address," and "the domain name service system comprising an indication that the domain name service system supports establishing a secure communication link." Thus, for reasons similar to those discussed above in connection with independent claim 1, *Aziz* does not anticipate claims 36 and 60. Accordingly, for similar reasons, Patent Owner requests that the rejection of claims 36 and 60 under 35 U.S.C. § 102 be withdrawn, and the patentability of the claims be confirmed.

**c. Dependent Claims 5, 23, and 47**

Dependent claims 5, 23, and 47 depend from independent claim 1 or 36, and include all of its features. Thus, *Aziz* does not anticipate these claims, and the rejection of these claims should be withdrawn, at least for the reasons discussed above in connection with the independent claims. Claims 5, 23, and 47 also distinguish over *Aziz* for additional reasons. For example, dependent claim 5 recites that "the domain name service system is configured to authenticate the query [for a network address] using a cryptographic technique." Similarly, dependent claim 23 recites that "the domain name service system is configured to authenticate the query for the network address," and dependent claim 47 recites that "the instructions [executable in a domain name service system] comprise code for authenticating the query for the network address." The rejection of claims 5, 23, and 47 in view of *Aziz* should be withdrawn for the additional reasons that neither of the two portions of *Aziz* cited by the Request and the Office Action discloses "a domain name service system [that] is configured to authenticate the query [for a network address] using a cryptographic technique," as recited in claim 5 and similarly recited in claims 23 and 47.

The first portion of *Aziz* relied upon by the Request discloses that the SIG resource record "can be used to authenticate data *in other resource records*." (*Aziz* 5:67-6:1 (emphasis added); *see*



New Bay Capital, LLC
Ex.1015-Page 2565 of 3151

Req. Ex. F-2 at 11-13, 41.) But authenticating data in a DNS resource record does not disclose "authenticat[ing] the query [for a network address]," as recited in claim 5. (Keromytis Decl. ¶ 64.) In fact, one of ordinary skill in the art would have understood that a DNS resource record is generally not a query for a network address. (*Id.*) Moreover, merely disclosing that the SIG resource record "*can be used* to authenticate data" does not disclose that "*the domain name service system* is configured to authenticate" anything. (*Id.*) In fact, *Aziz* does not disclose that the alleged domain name service system (NS 120) authenticates queries with the SIG resource record. (*Id.*)

The second portion of *Aziz* relied upon by the Request describes authentication in general terms, stating that "[a]uthentication means that a host is assured that the message is from the client that the message claims," and then lists several standard cryptographic methods. (*Aziz* 3:22-29; *see* Req. Ex. F-2 at 11-13, 28-29, 41.) General statements about the meaning of authentication also do not disclose that "the domain name service system is configured to authenticate the query using a cryptographic technique," as recited in claim 5. (Keromytis Decl. ¶ 65.)

Thus, *Aziz* has not been shown to disclose that "the domain name service system is configured to authenticate the query [for a network address] using a cryptographic technique," as recited in claim 5 and similarly recited in claims 23 and 47. Accordingly, Patent Owner requests that the rejection of claims 5, 23, and 47 under 35 U.S.C. § 102 be withdrawn, and the patentability of claims 5, 23, and 47 be confirmed.

**d. Dependent Claim 8**

Dependent claim 8 depends from independent claim 1 and includes all of its features. Thus, *Aziz* does not anticipate claim 8, and the rejection of this claim should be withdrawn, at least for the reasons discussed above in connection with independent claim 1. Dependent claim 8 also distinguishes over *Aziz* for additional reasons. For example, dependent claim 8 also recites that "the domain name service system is connectable to a virtual private network through the communication network."

The Request cites a portion of *Aziz* that describes inside NS 130 inside a protected zone 180 and asserts that "[h]aving the 'inside NS [130]' (inside Name Server) within the 'protected zone'. . . shows the domain name service system is connectable to a virtual private network . . . ." (Req. Ex. F-2 at 15.) As discussed, however, outside NS 120, and not inside NS 130 of *Aziz*, is the alleged domain name service system. But *Aziz* does not disclose that NS 120, the alleged domain name service system, is connectable to a virtual private network. Because *Aziz* does not disclose that the alleged domain name service system is connectable to a virtual private network, *Aziz* does not



New Bay Capital, LLC
Ex.1015-Page 2566 of 3151

disclose the recitations of claim 8. Accordingly, the rejection of claim 8 should be withdrawn and the claim should be confirmed.

**e. Dependent Claims 17 and 41**

As discussed above with regard to independent claims 1, 36, and 60, *Aziz* does not disclose the recited domain name service system that is configured to comprise an indication that the domain name service system supports establishing a secure communication link. Because *Aziz* does not describe such a domain name service system, it also does not disclose that the domain name service system comprises an indication that the domain name service system supports establishing a secure communication link, as recited in claim 17 and similarly recited in claim 41.

The Request points to the same four alleged features of *Aziz* that were relied on for the proposed rejections of claim 1 as also allegedly disclosing these features of claims 17 and 41. However, for reasons similar to those discussed above as to why those alleged features of *Aziz* do not disclose or suggest a domain name service system that is configured to comprise an indication that the domain name service system supports establishing a secure communication link, those features of *Aziz* also do not disclose or suggest the recited features of claims 17 and 41. Thus, *Aziz* does not anticipate claims 17 and 41, and the rejection of those claims should be withdrawn and their patentability confirmed.

**f. Dependent Claims 18 and 42**

Dependent claims 18 and 42 depend from independent claims 1 and 36, respectively, and include all of their features. Thus, *Aziz* does not anticipate these claims, and the rejection of these claims should be withdrawn, at least for the reasons discussed above in connection with the independent claims. Claims 18 and 42 also distinguish over *Aziz* for additional reasons. For example, dependent claim 18 recites "at least one of the plurality of domain names is reserved for secure communication links," and dependent claim 42 recites "the instructions comprise code for reserving at least one of the plurality of domain names for secure communication links." *Aziz* does not disclose these features of claims 18 and 42.

The Request and the Office Action assert that "[h]aving a domain name 'zone' configured to 'ensure that . . . communications are secure' shows that at least one of the plurality of domain names—the domain name associated with the protected zone—is reserved for secure communication links." (OA at 16; Req. Ex. F-2 at 24-25.) But *Aziz* does not disclose reserving any domain names for secure communication links, and the portions of *Aziz* relied on by the Request do not support the Request's assertions to the contrary. (Keromytis Decl. ¶ 88.)



New Bay Capital, LLC
Ex.1015-Page 2567 of 3151

The Request cites to three portions in the Background section of *Aziz* as allegedly disclosing the claimed feature. (*Id.*, citing *Aziz* 1:58-68, 2:9-22, 3:14-24.) The first portion discloses that organizations may divide a domain into zones and that each zone may have a database that contains names, addresses, and other information for that zone. (*Aziz* 1:58-68; Keromytis Decl. ¶ 89.) The second portion discloses that the organization can configure the registered name servers for the zone to either contain information regarding a machine (i.e., in a publicly visible zone) or not contain information regarding the machine (i.e., in a visibility-limited or "protected" zone). (*Id.* at 2:9-22; Keromytis Decl. ¶ 89.) The Request asserts that this portion teaches that "*machines in one zone can be reserved for secure communication links* by being put in a 'protected zone.'" (Req. Ex. F-2 at 24-25, emphasis added.) But this portion of *Aziz* is silent regarding reserving machines for secure communication links, let alone that *a domain name* is reserved for secure communication links. (Keromytis Decl. ¶ 89.) The third portion of *Aziz* relied on by the Request merely discloses that when a "client and host want to ensure that their communications are 'secure,'" the client may require information in addition to the address of the host. (*Aziz* 3:14-24.) This portion also does not disclose that a domain name is reserved for secure communication links. (Keromytis Decl. ¶ 89.)

When taken together, the above portions of *Aziz* relied on by the Request at best disclose that an organization can divide a domain into zones, may then selectively hide addresses of certain machines within a protected zone, and that if a client wants to ensure that communications with a host are secure, it may require additional information other than the host's address. These portions are silent regarding reserving domain names, let alone that at least one of a plurality of domain names is reserved for secure communication links. In fact, *Aziz* as a whole simply does not disclose that a domain name is reserved for secure communication links, as recited in claim 18 and similarly recited in claim 42. Accordingly, the rejection of these claims should be withdrawn and the patentability of these claims should be confirmed.

**g. Dependent Claims 24 and 48**

Dependent claims 24 and 48 depend from independent claims 1 and 36, respectively, and include all of their features. Thus, *Aziz* does not anticipate these claims, and the rejection of these claims should be withdrawn, at least for the reasons discussed above in connection with the independent claims. Claims 24 and 48 also distinguish over *Aziz* for additional reasons. For example, dependent claim 24 recites "at least one of the plurality of domain names comprises an indication that the domain name service system supports establishing a secure communication link," and dependent claim 48 recites "at least one of the plurality of domain names includes an indication



New Bay Capital, LLC
Ex.1015-Page 2568 of 3151

that the domain name service system supports the establishment of a secure communication link." *Aziz* does not disclose at least these features.

The Request asserts that *Aziz* anticipates these claims because "[h]aving a machine be in the visibility-limited 'protected zone' is an indication that the domain name service system supports establishing a secure communication link . . . ." (Req. Ex. F-2 at 30, citing *Aziz* 2:9-22.) But claim 24 recites "*at least one of the plurality of domain names* comprises an indication that the domain name service system supports establishing a secure communication link" (emphasis added). The portion of *Aziz* relied on by the Request merely discloses that if a machine is in a "visibility-limited" or "protected zone," the registered name servers may not have any information about the zone name servers for that zone. (*Aziz* 2:9-22; Keromytis Decl. ¶ 91.) It is silent regarding what the domain names of machines within that zone include, and certainly does not disclose that "at least one of a plurality of domain names comprises an indication that the domain name service supports establishing a secure communication link." (Keromytis Decl. ¶ 91.) Even the Request's assertion that "[h]aving a machine be in the visibility-limited 'protected zone' is an indication . . ." says nothing about what a *domain name* comprises. (Req. Ex. F-2 at 30, citing *Aziz* 2:9-22; Keromytis Decl. ¶ 91.)

Indeed, just because a machine is in a protected zone does not mean that the domain name of the machine comprises an indication that a domain name service system supports establishing a secure communication link. (Keromytis Decl. ¶ 91.) And *Aziz* does not disclose a domain name that includes any indication of the capabilities of the alleged domain name service system, let alone an indication that the alleged domain name service system supports establishing a secure communication link.

Accordingly, *Aziz* does not disclose the features of claims 24 and 48, and the rejection of these claims should be withdrawn.

**h. Dependent Claim 50**

Dependent claim 50 depends from independent claim 36, and includes its features. Thus, *Aziz* does not anticipate this claim, and the rejection should be withdrawn, at least for the reasons discussed above in connection with claim 36. Claim 50 also distinguishes over *Aziz* for additional reasons. For example, dependent claim 50 recites "at least one of the plurality of domain names is configured so as to enable establishment of a secure communication link." In the rejection of claim 50, the Office Action, incorporating the Request, refers to the analysis of claim 26 without providing any additional analysis. (OA at 16; Req. Ex. F-2 at 42, "See analysis of portion [26.1].") However,



New Bay Capital, LLC
Ex.1015-Page 2569 of 3151

claim 26 is not rejected under § 102 in view of *Aziz*, but is only rejected under § 103 over *Aziz* in view of George Lawton, "New Top-Level Domains Promise Descriptive Names" ("*Lawton*") (Issue 11). Thus, the Request and the Office Action have not shown that claim 50 is anticipated by *Aziz*. Moreover, for reasons similar to those discussed below with regard to the rejection of claim 26 (Issue 11), *Aziz* does not disclose or suggest at least one of the plurality of domain names is configured so as to enable establishment of a secure communication link, as recited in claim 50. In view of the above, the rejection of claim 50 should be withdrawn.

**i. Dependent Claims 2, 6, 7, 14, 15, 19-22, 24, 25, 27, 33-40, 43-46, 49, 51, 52, 58 and 59**

Remaining claims 2, 6, 7, 14, 15, 19-22, 24, 25, 27, 33-40, 43-46, 49, 51, 52, 58, and 59 depend from one of independent claims 1 and 36 and include all of their features. Thus, *Aziz* does not anticipate any of these claims for at least the reasons discussed above in connection with independent claims 1 and 36. For the reasons set forth above, the rejection of claims 2, 6, 7, 14, 15, 19-22, 24, 25, 27, 33-40, 43-46, 49, 51, 52, 58, and 59 based on *Aziz* should be withdrawn and the claims should be found patentable.

**3. Rejection of Claims 1, 2, 5-9, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 103(a) Based on *Aziz* (Issue 10)**

The Office Action rejects claims 1, 2, 5-9, 14-25, 27, 28, 33-52, and 57-60 under 35 U.S.C. § 103(a) as being obvious over *Aziz*. (OA at 12.) However, the Request's and the Office Action's analysis of these claims is deficient. "The key to supporting any rejection under 35 U.S.C. 103 is the clear articulation of the reason(s) why the claimed invention would have been obvious." M.P.E.P. § 2142. "[R]ejections on obviousness cannot be sustained with mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness." *Id.* (internal quotations omitted).

The Request and the Office Action fail to provide such a required articulated reasoning. Instead, the Request's and the Office Action's analysis in the § 103(a) rejection is *identical* to its analysis in the § 102(b) rejection, with the sole difference being that the word "teaches" has been globally replaced with the words "renders obvious." (*Compare* Req. Ex. F-2 at 3-44 *with id.* at 45-85.) Patent Owner was unable to identify any other differences in the analysis.

In particular, the Request and the Office Action do not explain how *Aziz* renders obvious those features that *Aziz* fails to disclose, as shown above in the traversal of the § 102(b) rejections based on *Aziz* (Issue 9). Thus, Patent Owner incorporates the arguments for why *Aziz* does not disclose the recited features of claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 that were discussed



New Bay Capital, LLC
Ex.1015-Page 2570 of 3151

above with regard to the § 102(b) rejections based on *Aziz*. And Patent Owner further asserts that *Aziz* does not suggest or otherwise render obvious these missing features. Nor do the Request and the Office Action provide any additional explanation as to why *Aziz* allegedly renders these missing features obvious. In view of the above, the rejection of claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 under § 103(a) as being obvious over *Aziz* should be withdrawn and the claims should be confirmed.

While the Office Action also purports to reject claim 9 under § 103(a) in view of *Aziz*, the Request and the Office Action are devoid of any analysis regarding the alleged obviousness of claim 9 over *Aziz*. (*See id.* at 56-57.) Thus, the rejection of this claim in view of *Aziz* should also be withdrawn and the claim should be confirmed.

#### 4. Rejection of Claims 3, 4, and 26 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Lawton* (Issue 11)

The Office Action rejects claims 3, 4, and 26 under 35 U.S.C. § 103(a) based on *Aziz* in view of *Lawton*. (OA at 12.) For the reasons discussed below, the rejection should be withdrawn.

##### a. Claims 3 and 4

Claims 3 and 4 depend from independent claim 1. As explained above, *Aziz* does not disclose or suggest the features of claim 1 and thus does not support the rejection of those claims. The rejection of claims 3 and 4 should also be withdrawn because *Lawton* does not remedy the deficiencies of *Aziz* with respect to independent claim 1. For instance, *Lawton* does not disclose or suggest, and the Request and the Office Action do not rely upon *Lawton* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claims 3 and 4 over *Aziz* in view of *Lawton* should be withdrawn.

##### b. Claim 26

Dependent claim 26 recites "at least one of the plurality of domain names enables establishment of a secure communication link." *Aziz* and *Lawton*, alone or in combination, do not disclose or suggest this feature.

*Aziz* does not disclose or suggest that at least one of the plurality of domain names enables establishment of a secure communication link. The portion of *Aziz* relied on by the Request merely discloses that domain 100 in Fig. 1 includes a protected zone 180 that includes a number of machines. (*Id.* at 88, citing *Aziz* 5:11-18.) In the background section of *Aziz*, it explains that the network topology of a "protected zone" may be hidden by configuring registered name servers to not have information about the zone name servers, and to only allow machines within the protected zone



New Bay Capital, LLC
Ex.1015-Page 2571 of 3151

to direct queries to the zone name servers. (*Aziz* 2:9-22.) But hiding the network topology of a protected zone discloses nothing about the capability of domain names, let alone that at least one of a plurality of domain names enables establishment of a secure communication link.

*Lawton* does not make up for the deficiencies of *Aziz* because *Lawton* also does not disclose or suggest that at least one of the plurality of domain names enables establishment of a secure communication link. *Lawton* discloses that "[t]here are efforts underway . . . to create alternative domain name server (DNS) extensions" such as ".med," ".xxx," and ".ltd" instead of the top-level domains (TLDs) such as ".com," ".org," etc. (*Lawton* 1.) *Lawton* also discloses that one can "download the alternate list of TLDs for use on your DNS server" to give people access to the alternate TLDs. (*Id.* at 2.) But the mere existence of alternative top-level domains also does not disclose or suggest that at least one of the plurality of domain names enables establishment of a secure communication link.

Apparently conceding that neither *Aziz* nor *Lawton* discloses the recitation of claim 26, the Request, attempting to combine *Aziz* and *Lawton*, asserts that

> [h]aving a non-standard domain name such as that taught by *Lawton* would require a connection to the appropriate "inside NS" (inside Name Server) within the "protected zone," enabling a secure communication link as recited by the claim. For example, in view of the teachings of *Lawton*, it would have been obvious to register the ".secure" or ".encrypted" domain names to identify secure or encrypted servers.

(Req. Ex. F-2 at 89.) These assertions are incorrect for at least two reasons and thus fail to demonstrate that the combination of *Aziz* and *Lawton* discloses or suggests the features of claim 26.

First, the assertion that "[h]aving a non-standard domain name such as that taught by *Lawton* would require a connection to the appropriate 'inside NS' (inside Name Server) within the 'protected zone'" is simply unsupported by the applied references. In particular, *Lawton* does not disclose or suggest that "having a non-standard domain name" requires a connection to an inside name server within a protected zone. In fact, *Lawton* suggests the opposite by showing different ways to "mak[e] [a non-standard domain name] *accessible to everyone*." (*Lawton* 2, emphasis added.) Thus, contrary to the Request's assertion, *Lawton* does not disclose that "having a non-standard domain name such as that taught by *Lawton*" would require a connection to an inside NS within a protected zone. And *Aziz* is silent regarding non-standard domain names. The Request asserts that this "requirement" teaches "enabling a secure communication link as recited in the claim." Because *Lawton* and *Aziz* do not disclose or suggest such a requirement, the Request has failed to show that the combination of *Lawton* and *Aziz* discloses or suggests that at least one of the plurality of domain names enables establishment of a secure communication link.



New Bay Capital, LLC
Ex.1015-Page 2572 of 3151

Second, the assertion that "in view of the teachings of *Lawton*, it would have been obvious to register the '.secure' or '.encrypted' domain names to identify secure or encrypted servers," is also incorrect. First, *Lawton* does not disclose the ".secure" or ".encrpyted" domain names. Second, *Lawton* does not disclose or suggest any connection between the alternative DNS extensions and the attributes of the servers that they identify. In particular, contrary to the Request's assertions, *Lawton* does not disclose or suggest "registering the '.secure' or '.encrypted' domain names to identify secure or encrypted servers." Moreover, the Request does not point out how "registering the '.secure' or '.encrypted' domain names to identify secure or encrypted servers" discloses the recited claim feature of at least one of the plurality of domain names enabling establishment of a secure communication link.

In view of the above, the Request and the Office Action have failed to demonstrate that *Aziz* and *Lawton*, alone or in combination, disclose or suggest "at least one of the plurality of domain names enables establishment of a secure communication link." Accordingly, the rejection of claim 26 in view of *Aziz* and *Lawton* should be withdrawn and the claim should be confirmed.

**5. Rejection of Claim 9 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Franaszek* (Issue 12)**

Claim 9 depends from claim 8, which in turn depends from claim 1. As explained above, *Aziz* does not disclose or suggest the features of claims 1 and 8, and thus does not support the rejection of those claims. The rejection of claim 9 over *Aziz* in view of U.S. Patent No. 4,952,930 to Franaszek et al. ("*Franaszek*") should be withdrawn because *Franaszek* does not make up for the deficiencies of *Aziz* discussed above with respect to claims 1 and 8. For instance, *Franaszek* does not disclose or suggest, and the Request and the Office Action do not rely on *Franaszek* as allegedly disclosing or suggesting, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," or that "the domain name service system is connectable to a virtual private network through the communication network."

As discussed below, the rejection of claim 9 over *Aziz* in view of *Franaszek* should also be withdrawn because *Aziz* and *Franaszek*, alone or in combination, do not disclose or suggest that "the virtual private network is one of a plurality of secure communication links in a hierarchy of secure communication links," as recited in claim 9 and because one of ordinary skill in the art would not have looked to *Franaszek* to modify the features disclosed in *Aziz* in the manner suggested in the Request..



New Bay Capital, LLC
Ex.1015-Page 2573 of 3151

**a. *Aziz* and *Franaszek* Do Not Disclose or Suggest the Features of Claim 9**

Claim 9 recites that "the virtual private network is one of a plurality of secure communication links in a hierarchy of secure communication links." *Aziz* and *Franaszek*, alone or in combination, do not disclose or suggest this feature.

*Aziz* does not disclose or suggest that the virtual private network is *one of a plurality of secure communication links in a hierarchy of secure communication links*, because *Aziz* does not disclose a hierarchy of secure communication links. In its analysis of dependent claim 8, from which claim 9 depends, the Request asserts that the connection to inside NS 130 in *Aziz* is "the virtual private network" recited in the claim. But *Aziz* does not disclose that this connection to inside NS 130 is one of a plurality of secure communication links in a hierarchy of secure communication links. Nor do the Request and the Office Action assert that it does. Instead, the Request merely asserts that *Aziz* "contemplates imposing some kind of organization on the network structure." (Req. Ex. F-2 at 91.) But this also does not disclose or suggest a hierarchy of secure communication links. (Keromytis Decl. ¶ 71.) Thus, *Aziz* does not disclose or suggest that the virtual private network is *one of a plurality of secure communication links in a hierarchy of secure communication links.*

*Franaszek* does not make up for the deficiencies of *Aziz*. The Request asserts that *Franaszek* discloses "[a] hierarchy of multipath networks." (*Id.* at 91-92, quoting *Franaszek* Abstract.) Merely disclosing a hierarchy of multipath networks does not disclose or suggest a hierarchy of *secure communication links*. In fact, *Franaszek* does not disclose that any of the multipath networks in the "hierarchy of multipath networks" are secure communication links. (Keromytis Decl. ¶ 71.)

Thus, *Aziz* and *Franaszek*, alone or in combination, do not disclose or suggest a hierarchy of secure communication links, let alone that "the virtual private network is one of a plurality of secure communication links in a hierarchy of secure communication links," as recited in claim 9.

**b. The Alleged Combination Would Change the Principles of Operation of *Aziz***

One of ordinary skill in the art would not have combined *Franaszek* with *Aziz* as suggested in the Request because *Franaszek* would change the principles of operation of *Aziz*. M.P.E.P. § 2143.01(VI). The Request and the Office Action assert that "[i]t would have been obvious to organize the secure communication links into a hierarchy as taught by *Franaszek*." (Req. Ex. F-2 at 91.) This is incorrect because incorporating the "hierarchy of multipath networks" of *Franaszek* into the system disclosed in *Aziz* would change the principles of operation of *Aziz*. (Keromytis Decl. ¶ 91.)



New Bay Capital, LLC
Ex.1015-Page 2574 of 3151

*Franaszek* is directed to an interconnection system that can be used in a shared-memory computer system, such as a multiple instruction multiple data (MIMD) computer system, that includes multiple processors interconnected to multiple memory systems. (*See, e.g.*, *Franaszek* 1:20-28, 3:7-31.) *Franaszek* shows an exemplary system in Fig. 1, reproduced below:

FIG. 1

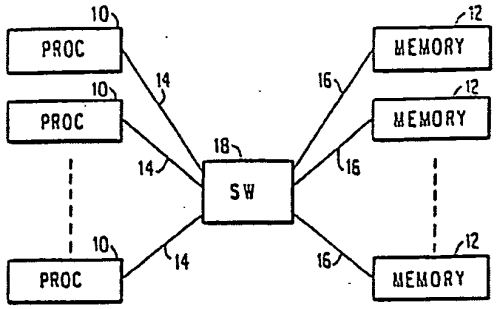


The particular interconnection system disclosed by *Franaszek* includes "a hierarchical network comprising multiple levels of multistage networks." (*Id.* at 2:3-5.) In particular, *Franaszek* explains that the "hierarchy of multipath networks" relied on by the Request is a "network hierarchy . . . configured with multiple levels of multistage networks, each one with different message transfer latencies." (*Id.* at 2:13-19.) *Franaszek* admits that these "multistage networks" are known and discloses an exemplary architecture of one multistage network, the Delta network, in Fig. 2 reproduced below. (*Id.* at 3:32-34.) *Franaszek* explains that the Delta network is "one component of" the hierarchical network included in the disclosed interconnection system. (*Id.* at 2:59-60.)

FIG. 2 PRIOR ART

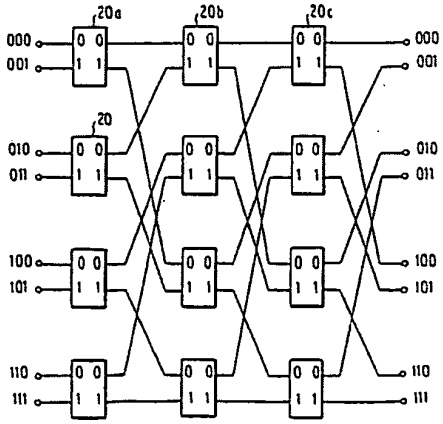




New Bay Capital, LLC
Ex.1015-Page 2575 of 3151

The Request proposes combining *Aziz* and *Franaszek* to "organize the secure communication links into a hierarchy as taught by *Franaszek*." (Req. Ex. F-2 at 91.) But the hierarchy taught by *Franaszek* includes a hierarchy of multistage networks, such as the one shown above, that are used for communicating between multiple processors and memories within shared-memory computer systems. (Keromytis Decl. ¶ 74.) The Request does not explain how at the time of the inventions one of ordinary skill in the art would have incorporated a hierarchy of these multistage networks into the system shown in Fig. 1 of *Aziz*. In contrast, one of ordinary skill in the art would have realized that connecting the hierarchy of multistage networks of *Franaszek* would change the principle of operation of *Aziz*, which discloses that the devices are connected to each other via traditional computer networking architectures such as LANs, WANs, the Internet, etc. (*Aziz* 4:50-55; Keromytis Decl. ¶ 74.) Attempting to incorporate the multistage network architecture described in *Franaszek* with the system of *Aziz* would change the principle of operation of *Aziz* because it would require the computer networks of *Aziz* to be modified to accommodate hardware components traditionally used for connecting processors and memories within a shared-memory computer. (Keromytis Decl. ¶ 74.)

**c. One of Ordinary Skill in the Art Would Not Have Relied on *Franaszek***

The rejection of claim 9 is further deficient because one of ordinary skill in the art would not have relied on *Franaszek* because *Franaszek* is nonanalogous art. One of the problems solved by embodiments of the '504 patent is that it makes it easy and convenient to enable secure communications. (Short Decl. in control no. 95/001788 ¶ 3.) For example, independent claim 1 recites "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Claims 8 and 9 further specify that the domain name service system is connectable to a virtual private network and that the virtual private network is one of a plurality of secure communication links in a hierarchy of secure communication links.

*Franaszek* is nonanalogous art because it is not reasonably pertinent to solving this problem. As discussed in the section above, *Franaszek* is directed to a hierarchy of multistage networks that can be used in an MIMD computer system. (*See e.g., Franaszek*, Fig. 1.) *Franaszek* does not provide any information whatsoever about making it easy and convenient to enable secure communications. Instead, *Franaszek* is concerned with developing a hierarchy technology that allows for an increased number of nodes without increased transfer delays. (*Franaszek* 2:6-12.) In contrast, one of ordinary skill in the art, faced with the problems solved by the '504 patent, would have attempted to find ways to enable secure communications easily and conveniently. (*See*



New Bay Capital, LLC
Ex.1015-Page 2576 of 3151

Keromytis Decl. ¶ 75.) The hierarchy of multistage networks disclosed in *Franaszek* would have been useless in addressing the problems solved by the features recited in claim 9 and disclosed by the '504 patent. Accordingly, *Franaszek* would not "have commended itself to an inventor's attention in considering *his problem*," and is thus nonanalogous art to the '504 patent. *In re Clay*, 966 F.2d 656, 659 (Fed. Cir. 1992) (emphasis added).

For at least the reasons set forth above, the rejection of claim 9 under § 103 should be withdrawn and the claim should be confirmed.

**6. Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Schneier* (Issue 13)**

Claim 10 depends from claim 8, which in turn depends from claim 1. As explained above, *Aziz* does not disclose or suggest the features of claims 1 and 8, and thus does not support the rejection of those claims. The rejection of claim 10 over *Aziz* in view of *Schneier* should be withdrawn because *Schneier* does not make up for the deficiencies of *Aziz* discussed above with respect to claims 1 and 8. For instance, *Schneier* does not disclose or suggest, and the Request and the Office Action do not rely on *Schneier* as allegedly disclosing or suggesting, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," or that "the domain name service system is connectable to a virtual private network through the communication network." Accordingly, the rejection of claim 10 over *Aziz* in view of *Schneier* should be withdrawn and the claim should be confirmed.

**7. Rejection of Claims 11-13 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Martin* (Issue 14)**

The Office Action rejects claims 11-13 under 35 U.S.C. § 103(a) based on *Aziz* in view of *Martin*. (OA at 12.) For the reasons discussed below, the rejection should be withdrawn.

**a. Claim 11**

Claim 11 depends from independent claim 1. As explained above, *Aziz* does not disclose or suggest the features of claim 1 and thus does not support the rejection of those claims. The rejection of claim 11 should also be withdrawn because *Martin* does not remedy the deficiencies of *Aziz* with respect to independent claim 1. For instance, *Martin* does not disclose or suggest, and the Request and the Office Action do not rely upon *Martin* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claim 11 over *Aziz* in view of *Martin* should be withdrawn.



New Bay Capital, LLC
Ex.1015-Page 2577 of 3151

Additionally, dependent claim 11 recites that "the virtual private network is based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." *Lendenmann* does not disclose these features—nor do the Request and the Office Action assert that it does. They also do not even assert that *Martin* discloses the claimed feature. Instead, the Request asserts that "[c]hoosing one of the source addresses at random" shows the elements of claim 11. (Req. Ex. F-2 at 99.) But this is not what the claim recites. Claim 11 recites that "the virtual private network is based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." *Martin* does not disclose pseudorandomly changing network addresses in packets, let alone a network address hopping regime that is used to pseudorandomly change network addresses in packets. (*See Martin* 9.) Because the Request and the Office Action do not even assert that *Martin* discloses this feature or explain how the combination of *Aziz* and *Martin* would render this missing feature obvious, the rejection is improper on its face and should be withdrawn. Moreover, because *Aziz* and *Martin* do not disclose or suggest the features of claim 11, either alone or in combination, the patentability of claim 11 should be confirmed.

**b. Claim 12**

Dependent claim 12 recites that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values." *Aziz* and *Martin*, alone or in combination, do not disclose or suggest this feature.

*Aziz* does not disclose that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values." Nor do the Request and the Office Action assert that it does.

*Martin* does not make up for the deficiencies of *Aziz*. In fact, the Request and the Office Action do not even assert that *Martin* discloses the claimed feature. Instead, the Request asserts that "it would be obvious to one of skill in the art to check the source of the data packet as coming from one of of [sic] the range of valid values as recited in the claim." (Req. Ex. F-2 at 99.) But this is not what the claim recites. Claim 12 recites that "the virtual private network is based on *comparing a value in each data packet* transmitted between a first device and a second device *to a moving window of valid values*" (emphases added). *Martin* does not disclose a moving window of valid values, let alone comparing a value in each data packet to a moving window of valid values. Because the Request and the Office Action do not even assert that *Martin* discloses this feature or explain how the combination of *Aziz* and *Martin* would render this missing feature obvious, the rejection is improper



New Bay Capital, LLC
Ex.1015-Page 2578 of 3151

on its face and should be withdrawn. Moreover, because *Aziz* and *Martin*, alone or in combination, do not disclose or suggest the features of claim 12, claim 12 is patentable over *Aziz* and *Martin*.

**c. Claim 13**

Dependent claim 13 recites that "the virtual private network is based on a comparison of a discriminator field in a header of each data packet to a table of valid discriminator fields maintained for a first device." *Aziz* and *Martin*, alone or in combination, do not disclose or suggest this feature.

*Aziz* does not disclose that "the virtual private network is based on a comparison of a discriminator field in a header of each data packet to a table of valid discriminator fields maintained for a first device." Nor do the Request and the Office Action assert that it does.

*Martin* does not make up for the deficiencies of *Aziz*. In fact, the Request and the Office Action do not even assert that *Martin* discloses the claimed feature. Instead, the Request asserts that "it would be obvious to one of skill in the art to check the source of the data packet as coming from one of the range of valid port numbers and addresses and discriminate based on those fields as recited in the claim." (Req. Ex. F-2 at 100.) But this is not what the claim recites. Claim 13 recites that "the virtual private network is based on a comparison of a discriminator field in a header of each data packet *to a table of valid discriminator fields maintained for a first device*" (emphasis added). *Martin* does not disclose a table of valid discriminator fields, let alone comparing a discriminator field in a header of each data packet to a table of valid discriminator fields. Because the Request and the Office Action do not even assert that *Martin* discloses this feature or explain how the combination of *Aziz* and *Martin* would render this missing feature obvious, the rejection is improper on its face and should be withdrawn. Moreover, because *Aziz* and *Martin*, alone or in combination, do not disclose or suggest the features of claim 13, claim 13 is patentable over *Aziz* and *Martin*.

In view of at least the above, the rejection of claims 11-13 in view of *Aziz* and *Martin* should be withdrawn and the patentability of the claims should be confirmed.

**8. Rejection of Claims 29-32 and 53-56 Under 35 U.S.C. § 103(a) Based on *Aziz* in View of *Ludwig* (Issue 15)**

Claims 29-32 depend from independent claim 1, and claims 53-56 depend from independent claim 36. As explained above, *Aziz* does not disclose or suggest the features of claims 1 and 36, and thus does not support the rejection of those claims. The rejection of claims 29-32 and 53-56 should also be withdrawn because *Ludwig* does not remedy the deficiencies of *Aziz* with respect to independent claims 1 and 36. For instance, *Ludwig* does not disclose or suggest, and the Request and the Office Action do not rely upon *Ludwig* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a



New Bay Capital, LLC
Ex.1015-Page 2579 of 3151

secure communication link," or "instructions executable in a domain name service system, the instructions comprising code for: . . . storing a plurality of domain names and corresponding network addresses; receiving a query for a network address; and supporting an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claims 29-32 and 53-56 over *Aziz* in view of *Ludwig* should be withdrawn.

### C. The Rejections Based on *Kiuchi* and *Pfaffenberger* Should Be Withdrawn

The Office Action rejects certain claims under 35 U.S.C. § 103 based on Takahiro Kiuchi and Shigekoto Kaihara, "C-HTTP — The Development of a Secure, Closed HTTP-based Network on the Internet" ("*Kiuchi*") in view of Bryan Pfaffenberger, "Netscape Navigator 3.0: Surfing the Web and Exploring the Internet" ("*Pfaffenberger*"), and based on *Kiuchi* in view of *Pfaffenberger* in combination with one or more additional references. However, as discussed below, the rejections under § 103 based on these references should be withdrawn.

#### 1. Overview of *Kiuchi*

*Kiuchi* proposes a technique called "closed HTTP" (C-HTTP) for providing secure HTTP communications "within a closed group of institutions on the Internet, where each member is protected by its own firewall." (*Kiuchi* 64.) According to *Kiuchi*, C-HTTP is useful in the medical community, where "there is a strong need for closed networks among hospitals and related institutions" to handle patient data and other sensitive medical information. (*Id.*)

C-HTTP requires three main components: "1) a client-side proxy on the firewall of one institution, 2) a server-side proxy on the firewall of another institution, and 3) a C-HTTP name server, which manages a given C-HTTP-based network and the information for [all of its] proxies." (*Id.*) When an institution wants to participate in a C-HTTP network, it must, among other things, install a client-side and/or server-side proxy on its firewall, register an IP address and a hostname for its proxy, and give the proxy's public key to the C-HTTP name server. (*Id.* at 65.) During C-HTTP communications, "[a] client-side proxy and server-side proxy communicate with each other using a secure, encrypted protocol (C-HTTP)." (*Id.* at 64.)

When a user agent computer behind a client-side proxy wants to establish a C-HTTP session with a server behind a server-side proxy, the following C-HTTP setup process occurs:

> (1) The client-side proxy asks the C-HTTP name server whether it can communicate with the server.



New Bay Capital, LLC
Ex.1015-Page 2580 of 3151

(2) The C-HTTP name server determines whether the server-side proxy is in the closed network and whether the connection is permitted.

(3) If so, the C-HTTP name server sends the IP address and public key of the server-side proxy, as well as request and response Nonce values, to the client-side proxy.

(4) The client-side proxy sends a connection request to the server-side proxy, encrypted with the server-side proxy's public key.

(5) The server-side proxy asks the C-HTTP name server whether the client-side proxy is also in the closed network and whether the connection is permitted.

(6) If so, the C-HTTP name server sends to the server-side proxy the IP address and public key of the client-side proxy, as well as the same request and response Nonce values previously sent to the client-side proxy.

(7) The server-side proxy then authenticates the client-side proxy, generates a connection ID, generates a second symmetric key for C-HTTP response encryption, and sends this information to the client-side proxy. When the client-side proxy accepts and checks this information, the connection is established.

(8) Once the connection is established, a client-side proxy forwards requests from the user agent in encrypted form using C-HTTP format.

(*Id.* at 65-66.) *Kiuchi* explains that "[t]he [C-HTTP] session is finished when the client accesses another C-HTTP server." (*Id.* at 65.)

**2. Overview of *Pfaffenberger***

*Pfaffenberger* is a guide that describes how to surf the World Wide Web using the Netscape Navigator® web browser, including how to use its various features. (*See generally Pfaffenberger.*) One of the features of Netscape Navigator that *Pfaffenberger* describes is a "doorkey icon," which is displayed in the lower left corner of the web browser to show when the web browser is accessing a secure server. (*See id.* at 9, 13, FIG. 1.2.) While *Pfaffenberger* focuses on the functionality of Netscape Navigator from the standpoint of the end-user, it is not concerned with the inner workings of domain name service systems.



New Bay Capital, LLC
Ex.1015-Page 2581 of 3151

### 3. Rejection of Claims 1-4, 6, 8-10, 12-19, 22, 24-30, 33, 34, 36-43, 46, 48-54, and 57-60 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* (Issue 16)

The Office Action rejects claims 1-4, 6, 8-10, 12-19, 22, 24-30, 33, 34, 36-43, 46, 48-54, and 57-60 under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger*. (OA at 16.) This rejection is deficient and should be withdrawn for at least the reasons discussed below.

#### a. Independent Claim 1

*Kiuchi* fails to disclose the combination of features recited in claim 1 for at least the reasons discussed below. Independent claim 1 recites, among other things, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." The Request and the Office Action assert that two different features of *Kiuchi* disclose the recited indication. The Request and the Office Action are incorrect because each asserted feature in *Kiuchi*, discussed in turn below, does not disclose a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link. Moreover, these features of *Kiuchi* are nothing more than features of a conventional domain name service system recognized and distinguished by the '504 patent.

##### (i) The C-HTTP Name Server Returning the Public Key of the Server-Side Proxy Is Not an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"

Incorporating portions of the Request by reference, the Office Action contends that "[t]he sending of the 'public key' [of the server-side proxy] is an indication that the domain name service [system] (C-HTTP name server) supports the establishment of [a] subsequent, secure communication link . . . ." (OA at 18; *see also* Req. Ex. F-3 at 11, quoting *Kiuchi* 65.) This is incorrect.

The public key of the server-side proxy returned by *Kiuchi*'s C-HTTP name server is not an indication that the alleged *domain name service system*, the C-HTTP name server, supports establishing a secure communication link. The client-side proxy uses the public key of the server-side proxy to send an encrypted connection request to the server-side proxy. (*Kiuchi* 65; Keromytis Decl. ¶ 56.) Thus, the public key corresponds to the server-side proxy, which is separate from the alleged domain name service system, the C-HTTP name server. (Keromytis Decl. ¶ 56.) The public key of the server-side proxy in *Kiuchi* includes no indication about the capabilities of the *C-HTTP name server itself*, and certainly does not include an indication that the C-HTTP name server supports establishing a secure communication link. (*Id.*)



New Bay Capital, LLC
Ex.1015-Page 2582 of 3151

Moreover, the C-HTTP name server in *Kiuchi* returning a public key of a server-side proxy is consistent with a conventional domain name service system that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*, '504 patent 39:7-42; Keromytis Decl. ¶ 57.) As discussed, the '504 patent indicates that a conventional domain name service system stores *public keys* of different machines so that hosts can request and receive those public keys from the domain name service system. ('504 patent 39:34-42; Keromytis Decl. ¶ 19, 57.) Like the conventional domain name service system described in the '504 patent, the C-HTTP name server of *Kiuchi* returns the public key of a server-side proxy when provided with the host URL for the server-side proxy. (*Kiuchi* 65; Keromytis Decl. ¶ 57.) Thus, the C-HTTP name server returning the public key of the server-side proxy is an aspect of a conventional domain name system that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*, '504 patent 39:7-42; Keromytis Decl. ¶¶ 56-57.)

Accordingly, the C-HTTP name server of *Kiuchi* returning a public key of a server-side proxy does not disclose a domain name service system configured to comprise an indication that the *domain name service system* supports establishing a secure communication link, as claimed.

**(ii) The C-HTTP Name Server Returning the IP Address of the Server-Side Proxy Is Not an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link"**

The Request also contends that the C-HTTP name server returning the IP address of the server-side proxy is an indication that the alleged domain name service system, the C-HTTP name server, supports establishing a secure communication link. (Req. at 24; Req. Ex. F-3 at 11, quoting *Kiuchi* 65.) However, it appears that the Office Action does not adopt the Request's position, as the Office Action only contends that the public key of the server-side proxy in *Kiuchi* is the claimed indication. (*See* OA at 18.) Thus, Patent Owner does not believe the Request's contention—that the C-HTTP name server returning the IP address of the server-side proxy is an indication that the C-HTTP name server supports establishing a secure communication link—is part of the rejection. Nonetheless, out of an abundance of caution, the Request's contention is addressed below.

Like returning the public key, the C-HTTP name server in *Kiuchi* returning the IP address of the server-side proxy is not an indication that the alleged *domain name service system*, the C-HTTP



New Bay Capital, LLC
Ex.1015-Page 2583 of 3151

name server, supports establishing a secure communication link. The IP address is for the server-side proxy, which is separate from the alleged domain name service system, the C-HTTP name server. (Keromytis Decl. ¶ 58.) Moreover, the IP address includes no indication about the capabilities of the C-HTTP name server itself, and certainly does not include an indication that the C-HTTP name server supports establishing a secure communication link. (*Id.*) It is just the address the C-HTTP name server returns when it receives a corresponding URL from the client-side proxy. (*See Kiuchi* 65; Keromytis Decl. ¶ 58.)

Like returning a public key, the C-HTTP name server returning the IP address of the server-side proxy in *Kiuchi* is a feature of a conventional domain name system that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*, '504 patent 39:7-42; Keromytis Decl. ¶ 19, 58.) As discussed, the '504 patent indicates that a conventional domain name service system merely returns an IP address or public key that was requested of it. (Keromytis Decl. ¶ 19.) For instance, the '504 patent explains that "[c]onventional Domain Name Servers (DNSs) provide a look-up function that *returns the IP address* of a requested computer or host. For example, when a computer user types in the web name 'Yahoo.com,' the user's web browser transmits a request to a DNS, which converts the name into a four-part IP address that is returned to the user's browser . . . ." ('504 patent 39:7-13, emphasis added; Keromytis Decl. ¶ 19; *see also* '504 patent 39:14-42.) Similar to the conventional domain name systems described by the '504 patent, the C-HTTP name server in *Kiuchi* returns an IP address corresponding to a provided URL. (*See Kiuchi* 65; Keromytis Decl. ¶ 58.)

Accordingly, the C-HTTP name server of *Kiuchi* returning an IP address of a server-side proxy does not disclose a domain name service system configured to comprise an indication that the *domain name service system* supports establishing a secure communication link, as claimed.

For the reasons provided above, *Kiuchi* has not been shown to disclose a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1.

**(iii) *Pfaffenberger* Does Not Remedy the Deficiencies of *Kiuchi***

In addition to alleging that *Kiuchi* discloses the claimed indication, the Request and the Office Action further contend that *Pfaffenberger* discloses the claimed indication, and that it would have been obvious to combine *Pfaffenberger* with *Kiuchi* to arrive at this subject matter of independent claim 1. (*See* OA at 18; Req. at 24-26; Req. Ex. F-3 at 2-3, 11-12.) The Request and



New Bay Capital, LLC
Ex.1015-Page 2584 of 3151

the Office Action are incorrect, at least because *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* explained above.

*Pfaffenberger* is a guide that describes how to surf the World Wide Web using the Netscape Navigator web browser, including how to use its various features. (*See generally Pfaffenberger.*) While *Pfaffenberger* focuses on the functionality of Netscape Navigator from the standpoint of the end-user, it is not concerned with how domain name service systems operate, and certainly does not disclose a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link, as recited by claim 1. (Keromytis Decl. ¶ 54.)

The Office Action, incorporating the Request by reference, contends that "*Pfaffenberger* also describes indicating support for a secure communication link by providing a visible icon on an http browser (Request, pp. 22-24) and that the addition of an http browser to the C-http system of *Kiuchi* would have been obvious (p. 22)." (OA at 18.) More specifically, the Request contends that "*Pfaffenberger* describes providing a 'Doorkey icon' that provides a visual indication to a user regarding whether a current communication link is secure." (Req. at 22, citing *Pfaffenberger* 13; *see also* Req. Ex. F-3 at 11-12.) According to the Request and the Office Action, *Pfaffenberger*'s doorkey icon discloses an indication that the domain name service system supports establishing a secure communication link. This is incorrect.

The cited portion of *Pfaffenberger* explains that "[t]his [doorkey] icon indicates whether you're accessing a secure server." (*Pfaffenberger* 13.) As illustrated in FIG. 1.2 of *Pfaffenberger*, the doorkey icon is provided in the lower left corner of the web browser window to show that the web browser is accessing a secure server. (*See id.* at 9.) Although *Pfaffenberger*'s doorkey icon shows "whether you're accessing a secure server," it does *not* indicate, specifically, that any purported *domain name service system* in *Pfaffenberger* supports establishing a secure communication link. (Keromytis Decl. ¶ 59.) The icon merely indicates that the end-user's web browser is currently accessing a secure server. (*Id.*) Indeed, *Pfaffenberger* is a user guide for the Netscape Navigator web browser and does not concern the operation of domain name service systems, much less domain name service systems configured to comprise an indication that the domain name service system supports establishing a secure communication link. (*Id.*)

For at least the above reasons, even if *Kiuchi* and *Pfaffenberger* were combined as proposed by the Request and the Office Action, the combination still would not disclose or suggest a domain name service system configured to comprise an indication that the domain name service system



New Bay Capital, LLC
Ex.1015-Page 2585 of 3151

supports establishing a secure communication link, as claimed. Thus, the rejection of claim 1 under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claim should be confirmed.

**(iv) One of Ordinary Skill in the Art Would Not Have Relied on *Pfaffenberger***

One of ordinary skill in the art would not have relied on *Pfaffenberger* to remedy the above noted deficiencies of *Kiuchi* because *Pfaffenberger* is nonanalogous art. One of the problems solved by the embodiments of the '504 patent is that it makes it easy and convenient to enable secure communications using a domain name service system. (Short Decl. in control no. 95/001,788 ¶ 3.) For example, independent claim 1 recites "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link."

*Pfaffenberger* is nonanalogous art because it is not reasonably pertinent to solving this problem. As discussed in the section above, *Pfaffenberger* is a guide for how to use the Netscape Navigator web browser, and is not at all concerned with the inner workings of domain name service systems, much less domain name service systems configured to comprise an indication that the domain name service system supports establishing a secure communication link. (Keromytis Decl. ¶ 60.)

In contrast, one of ordinary skill in the art, faced with the problems solved by the '504 patent, have looked for ways to enable secure communications easily and conveniently using a domain name service system. (*Id.*) But *Pfaffenberger*'s user guide would have been useless in addressing the problems solved by the features recited in claim 1 and disclosed by the '504 patent. Accordingly, *Pfaffenberger* would not "have commended itself to an inventor's attention in considering *his problem*," and is thus non-analogous art to the '504 patent. *In re Clay*, 966 F.2d at 659 (emphasis added).

For at least the reasons set forth above, even if *Kiuchi* and *Pfaffenberger* were combined as proposed by the Request and the Office Action, the combination still would not disclose or suggest a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link, as claimed. Additionally, one of ordinary skill in the art would not have relied upon *Pfaffenberger* to remedy the deficiencies of *Kiuchi* or to address the problems solved by the features recited in claim 1. Thus, the rejection of claim 1 under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claim should be confirmed.



New Bay Capital, LLC
Ex.1015-Page 2586 of 3151

**b. Independent Claims 36 and 60**

Independent claims 36 and 60 include recitations similar to those described above with respect to claim 1. For example, claim 36 recites "instructions executable in a domain name service system, the instructions comprising code for: . . . supporting an indication that the domain name service system supports establishing a secure communication link." And claim 60 recites, for example, "the domain name service system comprising an indication that the domain name service system supports establishing a secure communication link." And the Request incorporates by reference its analysis for independent claim 1 when addressing these features of independent claims 36 and 60. (*See* Req. Ex. F-3 at 41, 48-49.) Thus, for reasons similar to those discussed above in connection with independent claim 1, *Kiuchi* and *Pfaffenberger* would not have rendered obvious the subject matter of independent claims 36 and 60. Accordingly, Patent Owner requests that the rejection of claims 36 and 60 under 35 U.S.C. § 103 be withdrawn, and that the patentability of the claims be confirmed.

**c. Dependent Claims 8, 9, 10, 12, and 13**

Dependent claim 8 depends from independent claim 1 and includes all of its features. Thus, the combination of *Kiuchi* and *Pfaffenberger* does not render obvious claim 8, and the rejection of this claim should be withdrawn, at least for the reasons discussed above in connection with independent claim 1. Claim 8 also distinguishes over the combination for additional reasons. For example, claim 8 recites that "the domain name service system is connectable to a virtual private network through the communication network." Claims 9, 10, 12, and 13 also include these features because they depend from claim 8. Thus, the combination of *Kiuchi* and *Pfaffenberger* also does not disclose or suggest these additional features.

The Request contends that the C-HTTP connection in *Kiuchi* corresponds to the virtual private network recited in claim 8. (*See, e.g.*, Req. Ex. F-3 at 14-16.) One of ordinary skill in the art would not have viewed the C-HTTP connection in *Kiuchi* as a virtual private network in the context of the '504 patent.

One of ordinary skill, having read the '504 patent, would have understood a virtual private network, as recited in claims 8-13, to be a network of computers which privately communicate with each other by encrypting traffic on insecure communication paths between the computers. (Keromytis Decl. ¶ 77.) For instance, suppose two computers A and B reside on a public network and two computers X and Y reside on a private network. (*Id.*) If A establishes a virtual private network with X and Y's network to address data to X, and B separately establishes a virtual private



New Bay Capital, LLC
Ex.1015-Page 2587 of 3151

network with X and Y's network to address data to Y, then A would nevertheless be able to securely address data to B, X, and Y without additional setup. (*Id.*) This is true because A, B, X, and Y would all be part of the same virtual private network. (*Id.*)[3] *Kiuchi* fails to disclose such a virtual private network. (*Id.*)

But *Kiuchi*'s C-HTTP connection is very different from the claimed virtual private network. (*Id.* ¶ 78.) In *Kiuchi*'s C-HTTP system, a specific point-to-point connection is established each time a computer is to communicate with another computer using C-HTTP. (*Id.*) For instance, suppose, according to *Kiuchi*, a computer A (i.e., user agent or client-side proxy) establishes a C-HTTP session with a computer X (i.e., origin server or server-side proxy) in the closed network, and a computer B (i.e., user agent or client-side proxy) separately establishes a C-HTTP session with a computer Y (i.e., origin server or client-side proxy) in the same closed network. (*Id.*) In this case, computer A would be unable to access computer Y via a C-HTTP connection without first ending the existing C-HTTP session with computer X and again engaging in the C-HTTP setup process described above to establish a C-HTTP session with computer Y. (*Kiuchi* 65-66; Keromytis Decl. ¶ 78.) Similarly, computer B would be unable to communicate with computer X using C-HTTP without ending its session with computer Y and again completing the setup process to establish a new C-HTTP session with computer X. (Keromytis Decl. ¶ 78.) This is because *Kiuchi*'s C-HTTP connection is a point-to-point connection in which "[t]he [C-HTTP] session is finished when the client accesses another C-HTTP server." (*Kiuchi* 65; Keromytis Decl. ¶ 78.) In fact, Figure (b) of *Kiuchi* shows that C-HTTP connections to different servers in the closed network have distinct connection IDs:

b. The HTML document rewritten and forwarded to a use agent by the client-side proxy. The string, "6zdDfldfcZLj8V!i", attached to the end of the URLs is a connection ID

```
<TITLE>SAMPLE</TITLE>
<BODY>
<A HREF =
"http://server.in.current.connection/sample.html=@
=6zdDfldfcZLj8V!i">
Please click here.</A>
<A HREF =
"http://another.server.in.closed.network/=@=6zdDfl
dfcZLj8V!i">
Another server.</A>
</BODY>
```

[3] This view of "virtual private network" is supported by the interpretation of the term offered by another expert in the reexamination proceedings for other patents in the Munger patent family. (*See* Decls. of Jason Nieh, Ph.D., in control nos. 95/001,269 and 95/001,270.)



New Bay Capital, LLC
Ex.1015-Page 2588 of 3151

(*Kiuchi* 66; Keromytis Decl. ¶ 78.) And, as discussed, the server-side proxy generates the connection ID for a given C-HTTP connection during the above-described C-HTTP session setup process. (Keromytis Decl. ¶ 78.) This confirms that a C-HTTP connection is of a point-to-point nature and requires a new C-HTTP setup process each time a computer is to communicate with another computer using C-HTTP. (*Id.*)

In light of the additional setup required each time a client communicates with a different server using C-HTTP, one of ordinary skill in the art would have understood *Kiuchi*'s C-HTTP connection as a point-to-point connection rather than the claimed virtual private network. (*Id.* at 79.) Indeed, for at least this reason, one of ordinary skill in the art would not have understood computers connected via C-HTTP to be part of a virtual private network at all. (*Id.*)

Moreover, claim 8 recites that "*the domain name service system* is connectable to a virtual private network through the communication network" (emphasis added). The Request points to the C-HTTP communication between *the client-side proxy and the server-side proxy* as the claimed virtual private network. (Req. Ex. F-3 at 15, citing *Kiuchi* 64.) As explained by *Kiuchi*, "[o]nce the connection [between the proxies] is established, a client-side proxy forwards HTTP/1.0 requests from the user agent in encrypted form using C-HTTP format." (*Kiuchi* 66.) But this C-HTTP communication between the *proxies*, the alleged virtual private network, does not include the alleged *domain name service system*, the C-HTTP name server. (Keromytis Decl. ¶ 80.) Indeed, the alleged virtual private network in *Kiuchi* is not used for communication until after the proxies have already communicated with the C-HTTP name server to obtain each other's public keys and IP addresses. (*See id.* at 65-66; Keromytis Decl. ¶ 80.) Thus, the alleged domain name service system in *Kiuchi*, the C-HTTP name server, is not connectable to the alleged virtual private network, the C-HTTP communication between the proxies. (Keromytis Decl. ¶ 80.)

For the above reasons, *Kiuchi* fails to disclose the features of claims 8-13. Moreover, *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose a domain name service system connectable to a virtual private network, as required by claims 8-10, 12, and 13. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the combination fails to render obvious the features of claims 8-10, 12, and 13, and the rejection of these claims under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claims should be confirmed.



New Bay Capital, LLC
Ex.1015-Page 2589 of 3151

**d. Dependent Claim 10**

Dependent claim 10 depends from independent claim 1 via claim 8, and thus includes all of the features of these claims. Accordingly, the combination of *Kiuchi* and *Pfaffenberger* does not render obvious claim 10, and the rejection of this claim should be withdrawn, at least for the reasons discussed above in connection with claims 1 and 8. Claim 10 also distinguishes over the combination for additional reasons. For example, claim 10 recites that "the virtual private network is based on inserting into each data packet communicated over a secure communication link one or more data values that vary according to a pseudo-random sequence." The combination of *Kiuchi* and *Pfaffenberger* also does not render obvious these additional features.

The Request contends that *Kiuchi*'s disclosure of inserting "[r]andom bytes" "every fourth byte of the request and response before encryption in order to avoid the same encrypted requests or responses being repeated" discloses the above features of claim 10. (Req. Ex. F-3 at 17, citing *Kiuchi* 72.) This is incorrect at least because the random bytes are not inserted into packets of the alleged virtual private network in *Kiuchi, the C-HTTP connection between the proxies*. (Keromytis Decl. ¶ 81.)

The passage of *Kiuchi* cited in the Office Action is from Appendix 2A of *Kiuchi*, entitled "The summary of C-HTTP name server protocol." (*Kiuchi* 72.) As indicated by its title, Appendix 2A relates to name requests/responses *between the proxies and the C-HTTP name server*. (*Id.*; Keromytis Decl. ¶ 82.) In the cited passage, *Kiuchi* explains that random bytes are inserted into *name requests/responses between the proxies and the C-HTTP name server*. (*Kiuchi* 72; Keromytis Decl. ¶ 82.) But, as discussed above in connection with claim 8 from which claim 10 depends, the Request takes the position that the encrypted C-HTTP communication between the client-side proxy and server-side proxy, not the communication of name requests/responses between the proxies and the C-HTTP name server, corresponds to the claimed virtual private network. Thus, these random bytes are not inserted into data packets of the alleged virtual private network but into separate name requests/responses between the proxies and the C-HTTP name server. Thus, *Kiuchi* does not disclose that the alleged *virtual private network* is based on inserting into each data packet communicated over a secure communication link one or more data values that vary according to a pseudo-random sequence, as claimed.

In addition, the Request's analysis of claim 10 is deficient because it improperly relies on different components of *Kiuchi* for the same claimed element. But "[t]he key to supporting any rejection under 35 U.S.C. 103 is the clear articulation of the reason(s) why the claimed invention



New Bay Capital, LLC
Ex.1015-Page 2590 of 3151

would have been obvious." M.P.E.P. § 2142. This involves, among other things, considering all of the elements of the claim. *Id.* § 2143.03. By relying on different components of *Kiuchi* for the same element of claim 10, the Request has failed to provide a clear articulation of the reasons why the claimed invention would have been obvious.

In particular, the Request and Office Action rely on one alleged virtual private network (the encrypted C-HTTP communication between the client-side proxy and server-side proxy) for claim 8 and on another alleged virtual private network (the separate name requests/responses between the proxies and the C-HTTP name server) for claim 10. But claim 10 depends from claim 8 and refers to the same "virtual private network." Because of this inconsistency in the Request's analysis, the Request has not demonstrated, or even properly alleged, that *Kiuchi* discloses or suggests that the alleged *virtual private network* is based on inserting into each data packet communicated over a secure communication link one or more data values that vary according to a pseudo-random sequence, as recited in claim 10.

Moreover, *Kiuchi* discloses that these random bytes are inserted into *C-HTTP name requests and responses*. (*See Kiuchi* 72; Keromytis Decl. ¶ 83.) *Kiuchi* does not disclose that the random bytes are inserted into *each data packet*. Accordingly, even if the random bytes are viewed as the claimed "one or more data values that vary according to a pseudo-random sequence," *Kiuchi* still does not disclose that they are inserted into *each data packet*, as claimed. One of ordinary skill in the art would not have considered *Kiuchi*'s C-HTTP name requests and responses, which are application layer communications, to be data packets. (Keromytis Decl. ¶ 83.)

For the above reasons, *Kiuchi* fails to disclose the features of claim 10. Moreover, *Pfaffenberger* does not remedy the above-noted deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose a virtual private network based on inserting into each data packet communicated over a secure communication link one or more data values that vary according to a pseudo-random sequence, as recited in claim 10. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the combination fails to render obvious the features of claim 10, and the rejection of this claim under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claim should be confirmed.

**e. Dependent Claim 12**

Dependent claim 12 depends from independent claim 1 via claim 8, and thus includes all of the features of these claims. Accordingly, the combination of *Kiuchi* and *Pfaffenberger* does not render obvious claim 12, and the rejection of this claim should be withdrawn, at least for the reasons



New Bay Capital, LLC
Ex.1015-Page 2591 of 3151

discussed above in connection with claims 1 and 8. Claim 12 also distinguishes over the combination for additional reasons. For example, claim 12 recites that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values." The combination of *Kiuchi* and *Pfaffenberger* also does not disclose or render obvious these additional features.

The Request and the Office Action allege that the Nonce values contained in the headers of *Kiuchi*'s C-HTTP requests and responses constitute the claimed "value in each data packet." (OA at 18; Req. Ex. F-3 at 17-19.) Specifically, they allege that, "[i]n the Examples of C-HTTP communication found in Appendix 3, it can be seen that the 'Request-Nonce value' is incremented, moving from '8abd853f' in Example c., to '8abd8540' in Example g., to '8abd8541' in Example i." (*Id.* at 18.) According to *Kiuchi*, Example c. is a "Request for connection to the server-side proxy," Example g. is "Sending C-HTTP requests to the server-side proxy," and Example i. is a "Request for closing the connection." (*Kiuchi* 74.) At best, the Request has just shown that different requests contain different Nonce values.[4] (Keromytis Decl. ¶ 85.) The Request has not shown, and *Kiuchi* does not disclose, however, that these Nonce values are compared to a moving window of valid values, as recited in claim 12. *Kiuchi* just mentions that the "[r]eplay attacks are blocked by *checking* values of the Request-Nonce header field." (*Kiuchi.* at 65, emphasis added.) But *Kiuchi* does not explain *how* the values of the Nonce header field are checked, and certainly does not teach that they are checked by comparing them to a moving window of valid values. Further, there are many ways that the values of the Nonce header field could be checked without comparing them to a moving window of valid values. (Keromytis Decl. ¶ 85.) Thus, this feature is neither disclosed by, nor inherent in, *Kiuchi*.

In addition, as discussed, *Kiuchi* teaches that C-HTTP *requests* and *responses* contain a Nonce value in a Nonce header field. (*See Kiuchi* 65, 71; Keromytis Decl. ¶ 87.) But *Kiuchi* does not teach that Nonce values are inserted into *each data packet*. Accordingly, even if the Nonce values were compared to a moving window of valid values (which they are not), *Kiuchi* still does not disclose that the virtual private network is based on comparing a value in *each data packet* transmitted between the first computer and the second computer to a moving window of valid values, as claimed. One of ordinary skill in the art would not have considered *Kiuchi*'s C-HTTP requests and responses, which are application layer requests, to be data packets. (Keromytis Decl. ¶ 87.)

[4] Indeed, in secure communications, a nonce is a unique, arbitrary number used only once to identify a particular communication.



New Bay Capital, LLC
Ex.1015-Page 2592 of 3151

For the above reasons, *Kiuchi* fails to disclose the features of claim 12. Moreover, *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values," as recited in claim 12. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the combination fails to render obvious the features of claim 12, and the rejection of this claim under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claim should be confirmed.

**f. Dependent Claim 13**

Dependent claim 13 depends from independent claim 1 via claim 8, and thus includes all of the features of these claims. Accordingly, the combination of *Kiuchi* and *Pfaffenberger* does not render obvious claim 13, and the rejection of this claim should be withdrawn, at least for the reasons discussed above in connection with claims 1 and 8. Claim 13 also distinguishes over the combination for additional reasons. For example, claim 13 recites that "the virtual private network is based on a comparison of a discriminator field in a header of each data packet to a table of valid discriminator fields maintained for a first device." The combination of *Kiuchi* and *Pfaffenberger* also does not disclose or render obvious these additional features.

The Request and the Office Action contend that "*Kiuchi* teaches that the virtual private network . . . is based on a comparison of a connection ID field in the header of a request to a table of valid discriminator fields" because the connection ID is compared to a current connection table or list. (OA at 18; Req. Ex. F-3 at 20-21.) As pointed out in the Request, *Kiuchi* explains that the alleged discriminator field, a connection ID, is included in certain *C-HTTP requests*. But, again, *Kiuchi* does not teach that a connection ID is in a header of *each data packet*, as claimed. One of ordinary skill in the art would not have considered *Kiuchi*'s C-HTTP requests, which are application layer requests, to be data packets. (Keromytis Decl. ¶ 87.)

For the above reasons, *Kiuchi* fails to disclose the features of claim 13. Moreover, *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose that "the virtual private network is based on a comparison of a discriminator field in a header of each data packet to a table of valid discriminator fields maintained for a first device," as recited in claim 13. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the combination fails to render obvious the features of claim 13, and the rejection of



New Bay Capital, LLC
Ex.1015-Page 2593 of 3151

this claim under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claim should be confirmed.

**g. Dependent Claims 17 and 41**

Claims 17 and 41 depend from claims 1 and 36, respectively. As discussed above with regard to independent claims 1 and 36, the combination of *Kiuchi* in view of *Pfaffenberger* does not disclose or suggest the recited domain name service system that is configured to comprise an indication that the domain name service system supports establishing a secure communication link. Because the combination does not disclose or suggest such a domain name service system, it also does not disclose or suggest that the domain name service system comprises an indication that the domain name service system supports establishing a secure communication link, as recited in claim 17 and similarly recited in claim 41.

The Request incorporates by reference its analysis of independent claim 1 with respect to these features of claims 17 and 41. (*See id.* at 26, 43.) Accordingly, for reasons similar to those discussed above as to why those alleged features of *Kiuchi* and *Pfaffenberger* do not disclose or suggest a domain name service system that is configured to comprise an indication that the domain name service system supports establishing a secure communication link, those features of *Kiuchi* and *Pfaffenberger* also do not disclose or suggest the recited features of claims 17 and 41. Accordingly, the combination of *Kiuchi* and *Pfaffenberger* fails to render obvious the features of claims 17 and 41, and the rejection of these claims under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claims should be confirmed.

**h. Dependent Claims 24 and 48**

Dependent claims 24 and 48 depend from independent claims 1 and 36, respectively, and thus include all of the features of these claims. Accordingly, the combination of *Kiuchi* and *Pfaffenberger* does not render obvious claims 24 and 48, and the rejection of these claims should be withdrawn, at least for the reasons discussed above in connection with independent claims 1 and 36. Claims 24 and 48 also distinguish over the combination for additional reasons. For example, claim 24 recites that "at least one of the plurality of domain names comprises an indication that the domain name service system supports establishing a secure communication link," and claim 48 similarly recites that "at least one of the plurality of domain names includes an indication that the domain name service system supports the establishment of a secure communication link."

The Request contends that *Kiuchi* discloses these features because "*Kiuchi* provides example domain names that expressly indicate that the server is included within the secure C-HTTP closed



New Bay Capital, LLC
Ex.1015-Page 2594 of 3151

network," and "[t]he example domain name 'another.server.in.closed.network' indicates that the server is in the secure network and therefore supports establishing a secure communication link." (*Id.* at 30, citing *Kiuchi* 66.) But claim 24 recites "at least one of the plurality of domain names comprises an indication that the *domain name service system* supports establishing a secure communication link" (emphasis added). The Request contends that "[t]he example domain name 'another.server.in.closed.network' indicates that the *server is in the secure network* and therefore supports establishing a secure communication link." (*Id.*, emphasis added.) But the status of the server indicates nothing regarding the capabilities of the alleged *domain name service system*, the C-HTTP name server, and certainly would not indicate that *it* supports establishing a secure communication link.

For the above reasons, *Kiuchi* fails to disclose the features of claims 24 and 48. Moreover, *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose that at least one of the plurality of domain names comprises, or includes, an indication that the domain name service system supports establishing, or the establishment of, a secure communication link, as recited in claims 24 and 48. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the combination fails to render obvious the features of claims 24 and 48, and the rejection of these claims under 35 U.S.C. § 103(a) over *Kiuchi* in view of *Pfaffenberger* should be withdrawn and the claims should be confirmed.

**i. Dependent Claim 26**

Dependent claim 26 recites "at least one of the plurality of domain names enables establishment of a secure communication link." *Kiuchi* and *Pfaffenberger*, alone or in combination, have not been shown to disclose or suggest this feature.

With respect to these features of the claim, the Request contends that "*Kiuchi* teaches the use of a secure domain name service to control access to the establishment of secure communication links." (Req. Ex. F-3 at 31-32, citing *Kiuchi* 65, 68.) However, this is not what claim 26 recites. Claim 26 recites that "at least one of the plurality of domain names enables establishment of a secure communication link." Even if the Request is correct that *Kiuchi* teaches the use of a secure domain name service to control access to the establishment of secure communication links, which Patent Owner does not concede, this does not demonstrate anything about what any *domain name* in *Kiuchi*, itself, enables, much less that it enables establishment of a secure communication link.

Accordingly, the Request has not demonstrated that *Kiuchi* discloses that "at least one of the plurality of domain names enables establishment of a secure communication link," as recited in claim



New Bay Capital, LLC
Ex.1015-Page 2595 of 3151

26. Moreover, *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose that at least one of the plurality of domain names enables establishment of a secure communication link, as recited in claim 26. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the rejection of claim 26 in view of *Kiuchi* and *Pfaffenberger* should be withdrawn and the claim should be confirmed.

**j. Dependent Claim 27**

Dependent claim 27 recites that "the domain name service system is configured to enable establishment of a secure communication link between a first location and a second location transparently to a user at the first location." *Kiuchi* and *Pfaffenberger*, alone or in combination, have not been shown to disclose or suggest this feature.

With respect to this feature of the claim, the Request contends that "*Kiuchi* teaches that the secure C-HTTP closed network is designed to include a proxy server at each participating hospital" and that "the client-side proxy that processes a request for a user is located at the same institution as the user." (Req. Ex. F-3 at 32-33, citing *Kiuchi* 65.) However, that *Kiuchi*'s C-HTTP closed network allegedly includes proxies that process a request for a user does not demonstrate anything about what the alleged *domain name service system*, the C-HTTP name server, is configured to do. And it certainly does not disclose that the *C-HTTP name server* is configured to enable establishment of a secure communication link between a first location and a second location transparently to a user at the first location, as recited in claim 26.

Accordingly, the Request has not demonstrated that *Kiuchi* discloses that "the *domain name service system* is configured to enable establishment of a secure communication link between a first location and a second location transparently to a user at the first location" (emphasis added), as recited in claim 26. Moreover, *Pfaffenberger* does not remedy the deficiencies of *Kiuchi* at least because *Pfaffenberger* also fails to disclose that a domain name service system is configured to enable establishment of a secure communication link between a first location and a second location transparently to a user at the first location, as recited in claim 27. Nor does the Office Action or Request rely on *Pfaffenberger* for such disclosure. Accordingly, the rejection of claim 27 in view of *Kiuchi* and *Pfaffenberger* should be withdrawn and the claim should be confirmed.

**k. Dependent Claims 2-4, 6, 7, 11, 14-16, 18, 19, 22, 25, 28-30, 33, 34, 37-40, 42, 43, 46, 49-54, and 57-59**

Remaining claims 2-4, 6, 7, 11, 14-16, 18, 19, 22, 25, 28-30, 33, 34, 37-40, 42, 43, 46, 49-54, and 57-59 depend from one of independent claims 1 and 36 and include all of their features. Thus, the combination of *Kiuchi* and *Pfaffenberger* does not render obvious any of these claims for at least



New Bay Capital, LLC
Ex.1015-Page 2596 of 3151

the reasons discussed above in connection with independent claims 1 and 36. Accordingly, the rejection of claims 2-4, 6, 7, 11, 14-16, 18, 19, 22, 25-30, 33, 34, 37-40, 42, 43, 46, 49-54, and 57-59 based on the combination of *Kiuchi* and *Pfaffenberger* should be withdrawn and the claims should be found patentable.

#### 4. Rejection of Claims 5, 23, and 47 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Rivest* (Issue 17)

The Office Action rejects claims 5, 23, and 47 under 35 U.S.C. § 103(a) based on *Kiuchi* in view of *Pfaffenberger*, and further in view of R.L. Rivest et al., "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems" ("*Rivest*"). (OA at 16.) For the reasons discussed below, the rejection should be withdrawn and the claims should be confirmed.

Claims 5 and 23 depend from independent claim 1, and claim 47 depends from independent claim 36. As explained above, the combination of *Kiuchi* and *Pfaffenberger* does not disclose or suggest the features of claims 1 and 36, and thus does not support the rejection of those claims. The rejection of claims 5, 23, and 47 should also be withdrawn because *Rivest* does not remedy the deficiencies of the combination of *Kiuchi* and *Pfaffenberger* with respect to independent claims 1 and 36. For instance, *Rivest* does not disclose or suggest, and the Request and the Office Action do not rely upon *Rivest* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claims 5, 23, and 47 over *Kiuchi* in view of *Pfaffenberger*, and further in view of *Rivest*, should be withdrawn.

#### 5. Rejection of Claim 7 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Borella* (Issue 18)

The Office Action rejects claim 7 under 35 U.S.C. § 103(a) based on *Kiuchi* in view of *Pfaffenberger*, and further in view of U.S. Patent No. 6,269,099 to Borella et al. ("*Borella*"). (*Id.*) For the reasons discussed below, the rejection should be withdrawn and the claims should be confirmed.

Claim 7 depends from independent claim 1 and includes all of its features. As explained above, the combination of *Kiuchi* and *Pfaffenberger* does not disclose or suggest the features of independent claim 1 and thus does not support the rejection of that claim. The rejection of claim 7 should also be withdrawn because *Borella* does not remedy the deficiencies of the combination of *Kiuchi* and *Pfaffenberger* with respect to independent claim 1. For instance, *Borella* does not disclose or suggest, and the Request and the Office Action do not rely upon *Rivest* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name



New Bay Capital, LLC
Ex.1015-Page 2597 of 3151

service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claim 7 over *Kiuchi* in view of *Pfaffenberger*, and further in view of *Borella*, should be withdrawn.

### 6. Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Martin* (Issue 19)

The Office Action rejects claim 11 under 35 U.S.C. § 103(a) based on *Kiuchi* in view of *Pfaffenberger*, and further in view of *Martin*. (*Id.*) For the reasons discussed below, the rejection should be withdrawn.

Claim 11 depends from independent claim 1. As explained above, the proposed combination of *Kiuchi* and *Pfaffenberger* does not disclose or suggest the features of claim 1 and thus does not support the rejection of that claim. The rejection of claim 11 should also be withdrawn because *Martin* does not remedy the deficiencies of the proposed combination of *Kiuchi* and *Pfaffenberger* with respect to independent claim 1. For instance, *Martin* does not disclose or suggest, and the Request and the Office Action do not rely upon *Martin* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claim 11 over *Kiuchi* in view of *Pfaffenberger*, and further in view of *Martin*, should be withdrawn.

Additionally, dependent claim 11 recites that "the virtual private network is based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." The *Kiuchi-Pfaffenberger* combination does not disclose these features—nor do the Request and the Office Action assert that it does. They also do not even assert that *Martin* discloses the claimed feature. Instead, the Request asserts that "[c]hoosing one of the source addresses at random" shows the elements of claim 11. (Req. Ex. F-3 at 58.) But this is not what the claim recites. Claim 11 recites that "the virtual private network is based on a network address hopping regime that is used to pseudorandomly change network addresses in packets." *Martin* does not disclose pseudorandomly changing network addresses in packets, let alone a network address hopping regime that is used to pseudorandomly change network addresses in packets. (*See Martin* 9.) Because the Request and the Office Action do not even assert that *Martin* discloses this feature or explain how the combination of *Kiuchi* and *Martin* would render this missing feature obvious, the rejection is improper on its face and should be withdrawn. Moreover, because *Kiuchi-Pfaffenberger* and *Martin* do not disclose or suggest the features of claim 11, either alone or in combination, the patentability of claim 11 should be confirmed.



New Bay Capital, LLC
Ex.1015-Page 2598 of 3151

### 7. Rejection of Claims 20, 21, 35, 44, and 45 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Broadhurst* (Issue 20)

The Office Action rejects claims 20, 21, 35, 44, and 45 under 35 U.S.C. § 103(a) based on *Kiuchi* in view of *Pfaffenberger*, and further in view of U.S. Patent No. 6,560,634 to Broadhurst et al. ("*Broadhurst*"). (*Id.* at 17.) For the reasons discussed below, the rejection should be withdrawn.

Claims 20, 21, and 35 depend from independent claim 1, and claims 44 and 45 depend from independent claim 36. As explained above, the combination of *Kiuchi* and *Pfaffenberger* does not disclose or suggest the features of claims 1 and 36, and thus does not support the rejection of those claims. The rejection of claims 20, 21, 35, 44, and 45 should also be withdrawn because *Broadhurst* does not remedy the deficiencies of the combination of *Kiuchi* and *Pfaffenberger* with respect to independent claims 1 and 36. For instance, *Broadhurst* does not disclose or suggest, and the Request and the Office Action do not rely upon *Broadhurst* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claims 20, 21, 35, 44, and 45 over *Kiuchi* in view of *Pfaffenberger*, and further in view of *Broadhurst*, should be withdrawn.

### 8. Rejection of Claims 31, 33, 35, and 56 Under 35 U.S.C. § 103(a) Based on *Kiuchi* in View of *Pfaffenberger* and *Ludwig* (Issue 21)

The Office Action rejects claims 31, 33, 35, and 56 under 35 U.S.C. § 103(a) based on *Kiuchi* in view of *Pfaffenberger*, and further in view of *Ludwig*. (*Id.*) For the reasons discussed below, the rejection should be withdrawn.

Claims 31, 33, and 35 depend from independent claim 1, and claim 56 depends from independent claim 36. As explained above, the combination of *Kiuchi* and *Pfaffenberger* does not disclose or suggest the features of claims 1 and 36, and thus does not support the rejection of those claims. The rejection of claims 31, 33, 35, and 56 should also be withdrawn because *Ludwig* does not remedy the deficiencies of the combination of *Kiuchi* and *Pfaffenberger* with respect to independent claims 1 and 36. For instance, *Ludwig* does not disclose or suggest, and the Request and the Office Action do not rely upon *Ludwig* to show, at least "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." Thus, for at least the reasons set forth above, the rejection of claims 31, 33, 35, and 56 over *Kiuchi* in view of *Pfaffenberger*, and further in view of *Ludwig*, should be withdrawn.



New Bay Capital, LLC
Ex.1015-Page 2599 of 3151

**D. Secondary Considerations Demonstrate Nonobviousness**

Even if the Office had established a prima facie case of obviousness regarding any of claims 1-60 (which it has not), there is substantial evidence to rebut any finding of obviousness. As provided in M.P.E.P. § 2145, "[o]ffice personnel should consider all rebuttal arguments and evidence presented by applicants," including evidence relating to the secondary considerations as set forth in *Graham v. John Deere Co.*, 383 U.S. 1 (1966), which can support the nonobviousness of the claimed inventions. Those secondary considerations include commercial success, acceptance by others in the field, long-felt need, failure of others, and praise by others. Here, evidence related to secondary considerations rebuts any finding of obviousness of the claimed inventions.

Generally, the computer and Internet-security industries have long sought ways to conveniently establish secure communication links, such as VPN communication links. Around the time of the effective filing date of the '504 patent, it was widely recognized that providing secure remote access to a LAN or WAN was extremely difficult for IT support desks. (Short Decl. in control no. 95/001,788 ¶¶ 8, 11.) Specifically, remote access was "a nightmare for support desks. Staffers never kn[e]w what combination of CPU, modem, operating system and software configuration they [were] going to have to support," and adding the commercially available VPN software only made matters worse. (*Id.* ¶ 11.) The computer and Internet-security industries were forced to choose between ease of use and security, but they could not have both. (*Id.* ¶ 9.) The inventions claimed in the '504 patent, which provide a domain name service for establishing a secure communication link, combine both the ease of use *and* the security aspects of secure communication links without sacrificing one or the other. (*Id.*)

Prior to the features claimed in the '504 patent, there was a long-felt need for a system that could establish a secure communication link, such as a VPN communication link, in a simple and straightforward manner because "a solution that was difficult for an end-user to employ would likely have lead [sic] to a lack of use or incorrect use." (*Id.* ¶ 3.) As one example of the manifestation of the long-felt need, the Defense Advanced Research Projects Agency ("DARPA") funded various research programs to further the science and technology of information assurance and survivability. (*Id.* ¶¶ 4-5.) One such program, "Next Generation Internet," received approximately $130 million in funding between 1998 and 2000. (*Id.* ¶ 4.)

Recognizing this long-felt need for these inventions, both In-Q-Tel, a venture capital firm that invests in companies developing cutting-edge technology, and SAIC (the original owner of the '504 patent) also spent significant resources on their development. (*Id.* ¶¶ 6-7.) In fact, in the year



New Bay Capital, LLC
Ex.1015-Page 2600 of 3151

the inventions claimed in the '504 patent were developed, SAIC spent approximately 85% of its entire research and development budget for that year on developing these and other similar inventions. (*Id.* ¶ 7.)

Other attempts to provide an easy-to-use solution were unsuccessful. For example, the DARPA-funded research programs discussed above fell far short of the claimed inventions of the '504 patent. (*Id.* ¶¶ 4-5, 10.) One such program, "Dynamic Coalitions," was specifically created to address the ability of the Department of Defense to quickly and easily set up secure communications over the Internet. (*Id.* ¶¶ 4-5.) More than fifteen prestigious organizations took part in the "Dynamic Coalitions" research program, but none of them came up with a solution, in the relevant time frame, that was even close to the solutions provided in the claimed inventions of the '504 patent. (*Id.*) That is, they did not develop a solution that provided a domain name service for establishing a secure communication link. (*Id.*) By providing a domain name service for establishing a secure communication link, the inventions of the '504 patent succeeded where others failed. (*Id.* ¶ 11.)

The claimed inventions have also experienced commercial success. In particular, SafeNet, a leading provider of Internet-security technology that is the de facto standard in the VPN industry, entered into a portfolio license in July 2002 with the original owner of the application from which the '504 patent issued. (*Id.* ¶ 12.) SafeNet licensed the patents because of features disclosed and claimed in the patents, including those in the '504 patent. (*Id.*) In addition, Microsoft has entered into a similar license that includes the '504 patent. (*Id.*) Indeed, as noted, Microsoft was found to willfully infringe two of the patents in the Munger patent family, leading to a damages award of over one hundred million dollars. (*Id.*) And on May 3, 2012, Aastra USA, Inc. entered into a license with VirnetX that includes the '504 patent. (*See* Ex. A-5 at 1.)

The claimed inventions of the '504 patent were also contrary to the accepted wisdom at the time of the inventions. (Short Decl. in control no. 95/001,788 ¶ 13.) For example, there was a general understanding that reliable security could only be achieved through difficult-to-provision VPNs and that easy-to-set-up connections could not be secure. (*Id.*)

The technology of the '504 patent was also met with skepticism by those skilled in the art who learned of the patented inventions. (*Id.* ¶ 15.) For example, a DARPA program manager informed one of the coinventors of the '504 patent that the technology disclosed in the '504 patent would never be adopted. (*Id.*) Moreover, the IT offices of many large companies and institutions expressed skepticism that secure connections could ever be enabled easily by regular computer users. (*Id.*)



New Bay Capital, LLC
Ex.1015-Page 2601 of 3151

Several events also demonstrate praise for the inventions in the '504 patent by those in the field. As discussed above, SAIC invested a disproportionately large percentage of its internal resources in the technology. (*Id.* ¶ 17.) SafeNet, Microsoft, and Aastra have all licensed the technology. (*Id.*; Ex. A-5 at 1.) A study done by CSMG praised the inventions. (Short Decl. in control no. 95/001,788 ¶ 17.) Jim Rutt at Network Solutions, which was eventually acquired by Verisign, praised and expressed significant interest in the technology and would have invested but for a change in circumstances at his company. (*Id.*) This evidence showing that the claimed inventions met a long-felt need, succeeded where others have failed, have been commercially successful, were contrary to the accepted wisdom at the time of the invention, were met by skepticism by those skilled in the art, and were praised by others in the field, rebuts any finding that the claimed inventions would have been obvious.

## IV. Conclusion

For at least these reasons, VirnetX requests reconsideration and withdrawal of the rejections in the Office Action and confirmation of the patentability of all of the claims of the '504 patent.

VirnetX notes that the Request, Order, and Office Action contain a number of assertions and allegations concerning the disclosure, claims, and cited references. VirnetX does not subscribe to any assertion or allegation in the Request, Order, and Office Action regardless of whether it is addressed specifically herein.

Please grant any extension of time and charge any required fees to Deposit Account No. 06-0916.

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: June 1, 2012

By: /Joseph E. Palys/
Joseph E. Palys
Reg. No. 46,508



New Bay Capital, LLC
Ex.1015-Page 2602 of 3151

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | |
|---|---|
| In re *Inter Partes* Reexamination of: | |
| Victor Larson et al. | Control No.: 95/001,851 |
| U.S. Patent No. 7,418,504 | Group Art Unit: 3992 |
| Issued: August 26, 2008 | Examiner: Roland Foster |
| For: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | Confirmation No.: 1688 |

Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

**Declaration of Angelos D. Keromytis, Ph.D.**

I declare that the following statements are true to the best of my knowledge, information, and belief, formed after reasonable inquiry under the circumstances.

I, ANGELOS D. KEROMYTIS, declare as follows:

1. I have been retained by VirnetX Inc. ("VirnetX") for the above-referenced reexamination proceeding. I understand that this reexamination involves U.S. Patent No. 7,418,504 ("the '504 patent"). I further understand that the '504 patent is assigned to VirnetX and that it is part of a family of patents ("Munger patent family") that stems from U.S. provisional application nos. 60/106,261 ("the '261 application"), filed on October 30, 1998, and 60/137,704 ("the '704 application"), filed on June 7, 1999. I understand that the '504 patent is a continuation of U.S. application no. 09/558,210 ("the '210 application"), filed on April 26, 2000 (now abandoned), which is a continuation-in-part of U.S. application no. 09/504,783 (now U.S. Patent No. 6,502,135, "the '135 patent"). I also understand that the '135 patent is a continuation-in-part of U.S. application no. 09/429,643 (now U.S. Patent No. 7,010,604), which claims priority to the '261 and '704 applications.



New Bay Capital, LLC
Ex.1015-Page 2603 of 3151

## I. RESOURCES I HAVE CONSULTED

2. I have reviewed the '504 patent, including claims 1-60. I have also reviewed a Request for *Inter Partes* Reexamination of the '504 patent filed by Cisco Systems, Inc. with the U.S. Patent and Trademark Office on December 13, 2011 ("Request" or "Req."), as well as its accompanying exhibits.[1] Additionally, I have reviewed an Order Granting Request for *Inter Partes* Reexamination of the '504 patent ("the Order") and an Office Action ("the Office Action"), both mailed on March 1, 2012.[2]

3. I have also studied the following documents cited in and included with the Request and/or Office Action: Rolf Lendenmann, "Understanding OSF DCE 1.1 for AIX and OS/2" ("*Lendenmann*"); U.S. Patent No. 6,199,234 to Aziz, Jr. et al., entitled "Method and Apparatus for Client-Host Communication over a Computer Network," issued September 12, 2000 ("*Aziz*"); Takahiro Kiuchi and Shigekoto Kaihara, "C-HTTP — The Development of a Secure, Closed HTTP-based Network on the Internet" ("*Kiuchi*"); Bryan Pfaffenberger, "Netscape Navigator 3.0: Surfing the Web and Exploring the Internet" ("*Pfaffenberger*"); Information Sciences Institute, "Transmission Control Protocol," RFC 793 ("RFC 793"); D. Eastlake and C. Kaufman, "Domain Name System Security Extensions," RFC 2065 ("RFC 2065"); U.S. Patent No. 5,898,830 to Wesinger et al., entitled "Firewall Providing Enhanced Network Security and User Transparency," issued April 27, 1999 ("*Wesinger*"); U.S. Patent No. 5,689,641 to Ludwig et al., entitled "Multimedia Collaboration System Arrangement for Routing Compressed AV Signal Through a Participant Site Without Decompressing the AV Signal," issued November 18, 1997 ("*Ludwig*"); David M. Martin, "A Framework for Local Anonymity in the Internet" ("*Martin*"); Bruce Schneier, "Applied Cryptography" ("*Schneier*"); George Lawton, "New Top-Level Domains Promise Descriptive Names" ("*Lawton*"); Jean-Paul Gaspoz, "VPN on DCE: From Reference Configuration to Implementation" ("*Gaspoz*"); U.S. Patent No. 6,269,099 to Borella et al., entitled "Protocol and Method for Peer Network Device Discovery," issued July 31, 2001; U.S. Patent No. 6,560,634 to Broadhurst et al., entitled "Method of Determining Unavailability of an Internet Domain Name," issued May 6, 2003 ("*Broadhurst*"); Mark Pallen, "The World Wide Web" ("*Pallen*"); R.L. Rivest

---

[1] I refer to the Request for *Inter Partes* Reexamination as "the Request" and, correspondingly, I will refer to Cisco Systems, Inc. as "the Requester."

[2] The Office Action incorporates nearly all of the Request by reference. For that reason, when I sometimes refer to "the Request," I am also referring to the Office Action.



New Bay Capital, LLC
Ex.1015-Page 2604 of 3151

et al., "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems" ("*Rivest*"); U.S. Patent No. 4,952,930 to Franaszek et al., entitled "Multipath Hierarchical Network," issued August 28, 1990 ("*Franaszek*"); and Frederic Gittler et al., "The DCE Security Service" ("*Gittler*").

4. I am familiar with the level of ordinary skill in the art with respect to the inventions of the '504 patent as of February 15, 2000, when the application for the parent '135 patent was filed. Specifically, based on my review of the technology, the educational level of active workers in the field, and drawing on my own experience, I believe a person of ordinary skill in the art at that time would have had a master's degree in computer science or computer engineering, as well as two years of experience in computer networking with some accompanying exposure to network security.

5. I have been asked to consider how one of ordinary skill in the art would have understood the references mentioned above. My findings are set forth below.

## II. QUALIFICATIONS

6. I have a great deal of experience and familiarity with computer and network security, and have been working in this field since 1993.

7. I am currently an Associate Professor of Computer Science at Columbia University, as well as Director of the University's Network Security Laboratory. I joined Columbia in 2001 as an Assistant Professor, after receiving my M.Sc. and Ph.D. degrees in Computer Science, both from the University of Pennsylvania. My Ph.D. dissertation work was on the topic of secure access control for distributed systems and, in particular, on the management of trust in distributed computer networks.

8. I received my B.Sc. in Computer Science from the University of Crete, in Greece, in 1996. During my undergraduate studies, I worked as system administrator in the Computing Center at the University of Crete. Following that, I worked as network engineer at the first commercial Internet Service Provider ("ISP") in Greece, FORTHnet SA, where I was exposed to many network security issues.

9. I have actively participated in the Internet Engineering Task Force ("IETF"), a standards-setting body for the Internet, since 1995. In the late 1990s and early 2000s, my work with the IETF was primarily within the Internet Protocol Security ("IPsec") Working Group. In addition to contributing to the specification of the IPsec standards, I wrote the first implementation of the Photuris key management protocol (now RFC 2522). I also contributed to the first open-source implementation of the IKSAMP/IKE key management protocol for the open-source BSD operating system (now RFC 2409), and developed the first such implementation for the Linux operating



New Bay Capital, LLC
Ex.1015-Page 2605 of 3151

system. My Linux implementation, named Pluto, was adopted by the National Institute of Standards and Technology ("NIST") in 1999. In addition, my implementation of IPsec for the open-source BSD operating system is currently used by many companies and governments around the world, and serves as the basis for several commercial products that employ cryptographic communications. In 1999, I architected and implemented the first open-source framework for supporting hardware cryptographic accelerators. This framework is used in the open-source OpenBSD, NetBSD, FreeBSD, and Linux operating systems. My work in implementing firewalls and other cryptographic and network protocols has resulted in commercial systems and publications in refereed technical conferences and academic journals. I served as Working Group Secretary for the IETF IPsec Working Group (2003-2005) and as Security Area Advisor to the IETF at large (2003-2008).

10. In my current position at Columbia University, I work with a large group of graduate and postgraduate students in the area of cybersecurity. My past students now work in this field as university professors, as technical researchers for research laboratories, or as engineers for telecommunications companies. I have received federal, state, and corporate sponsorship to conduct cybersecurity research from the Department of Defense, the National Security Agency, the Defense Advanced Research Projects Agency ("DARPA"), the National Science Foundation, the Department of Homeland Security, the Air Force, the Office for Naval Research, the Army Research Office, the Department of the Interior, the National Reconnaissance Office, New York State, Google, Intel, Cisco, and others. In my ten years as a professor, I have received over thirty-six million dollars to support my research in cybersecurity. I also regularly teach courses on cybersecurity, in addition to more general courses in computer science.

11. I have published over 200 technical papers in refereed journals, conferences, and workshops, all of which are directed to various areas of cybersecurity. I have also authored a book, coauthored another book, and contributed chapters for many other books that relate to cybersecurity. Between 1999 and 2010, I have drafted or codrafted eight standards documents that were published as Request for Comments ("RFCs"). Several of these RFCs are directly related to IP security. For example, RFC 6042 relates to transport layer security; RFC 5708, RFC 2792, and RFC 2704 relate to key signature and encoding for trust management; and RFC 3586 relates to IP security policy requirements. Additionally, I am a coinventor on twelve issued U.S. patents, and have several other applications pending. Most of these patents and pending applications are related to network and systems security.

12. I have chaired several international technical conferences and workshops in cybersecurity, including, for example, the International Conference on Financial Cryptography and Data Security (FC), ACM Computer and Communication Security (CCS), and the New Security Paradigms Workshop (NSPW). I have also served in over eighty technical program committees for such events. From 2004-2010, I served as Associate Editor for the premier technical journal on cybersecurity—the ACM Transactions on Information and Systems Security (TISSEC). Additionally, I have served on several advisory workshops to the United States Government on cybersecurity, including, among others, the Office of the Director of National Intelligence (ODNI)/National Security Agency (NSA) Invitational Workshop on Computational Cybersecurity in Compromised Environments (C3E) (2011), the Office of Naval Research (ONR) Workshop on Host Computer Security (2010), the Intelligence Community Technical Exchange on Moving Target (2010), the Lockheed Martin Future Security Threats Workshop (2009), and the ARO/FSTC Workshop on Insider Attack and Cyber Security.

13. In addition to this work, I have cofounded two companies in cybersecurity. One company, StackSafe Inc. (formerly Revive Systems Inc.), was a provider of a virtualized preproduction staging environment that includes automated testing, analysis, and reporting for IT operations teams. I was with this company from its founding in 2005 until 2009. The second company, Allure Security Technologies (founded in 2010), develops deception-based solutions for detecting and mitigating the malicious cyber-insider threat, commercializing technology developed at Columbia through DHS and DARPA grants and a DARPA SBIR contract.

14. My curriculum vitae, which is appended to this declaration, details my background and technical qualifications. Although I am being compensated at my standard rate of $500/hour for my work on this declaration, the compensation in no way affects the statements in this declaration.

## III. BACKGROUND OF THE '504 PATENT

15. Before turning to a discussion of the references relied on in the Request and the Office Action, I summarize my understanding of certain embodiments disclosed in the '504 patent. Generally speaking, the '504 patent discloses, among other things, systems and methods for providing a domain name service ("DNS") for establishing a secure communication link.

16. The '504 patent discloses several embodiments of a domain name service system for establishing a secure communication link, such as a virtual private network ("VPN") communication link. In one such embodiment, a novel, specialized DNS server receives a traditional DNS request, and the DNS server automatically facilitates the establishment of a secure communication link, such

as a VPN link, between a target node and a user. ('504 patent 39:46-51.) This specialized DNS server is different from a conventional DNS server known at the time of invention for at least the reason that the specialized DNS server supports the establishment of a secure communication link beyond merely returning a requested IP address or public key.

17. For example, in the exemplars of FIGS. 26 and 27 of the '504 patent, reproduced below, a DNS server 2602 including a DNS proxy 2610 supports establishing a VPN link between a computer 2601 and a secure target site 2604. (*Id.* at 39:67-41:59.)

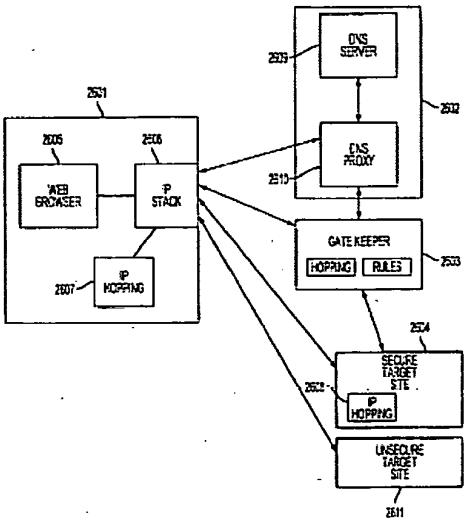


FIG. 26

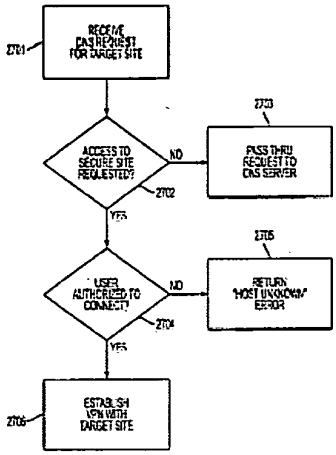


FIG. 27

18. In one embodiment, the DNS server 2602 receives a DNS request for a target site from computer 2601. (*Id.* at 40:49-52.) The DNS proxy 2610 determines whether the target site is a secure site. (*Id.* at 40:6-8, 40:49-56.) If access to a secure site has been requested, the DNS proxy 2610 determines whether the computer 2601 is authorized to access the site. (*Id.* at 40:57-59.) If so, the DNS proxy 2610 transmits a message to gatekeeper 2603 to facilitate the creation of a VPN link between computer 2601 and secure target site 2604. (*Id.* at 40:18-24.) DNS proxy 2610 then responds to the computer's 2601 DNS request with an address received from the gatekeeper 2603. (*Id.* at 40:19-22.) A secure VPN link is then established between the computer 2601 and the secure target site 2604. (*Id.* at 41:5-8.) As shown in this example, the specialized DNS server supports creating a secure communication link, or, in other words, does more than a conventional DNS server at the time of invention.



New Bay Capital, LLC
Ex.1015-Page 2608 of 3151

19. In fact, the '504 patent highlights this distinction between the specialized DNS server disclosed in its specification and a conventional DNS scheme, which merely returns a requested IP address or public key:

> Conventional Domain Name Servers (DNSs) provide a look-up function that returns the IP address of a requested computer or host. For example, when a computer user types in the web name "Yahoo.com," the user's web browser transmits a request to a DNS, which converts the name into a four-part IP address that is returned to the user's browser . . . .
>
> . . .
>
> One conventional scheme that provides secure virtual private networks over the Internet provides the DNS server with the public keys of the machines that the DNS server has the addresses for. This allows hosts to retrieve automatically the public keys of a host that the host is to communicate with so that the host can set up a VPN without having the user enter the public key of the destination host. One implementation of this standard is presently being developed as part of the FreeS/WAN project (RFC 2535).
>
> The conventional scheme suffers from certain drawbacks. For example, any user can perform a DNS request. Moreover, DNS requests resolve to the same value for all users.
>
> According to certain aspects of the invention, a specialized DNS server traps DNS requests and, if the request is from a special type of user (e.g., one for which secure communication services are defined), the server does not return the true IP address of the target node, but instead automatically sets up a virtual private network between the target node and the user.

(*Id.* at 39:7-51.) Compared with a conventional DNS known at the time of the filing of the '504 patent, the specialized DNS disclosed in the '504 patent supports establishing a secure communication link. The claims of the '504 patent are also directed to a domain name service for establishing a secure communication link. (*See, e.g.*, *id.* at 55:49-56, 57:48-58, 60:3-14.)

## IV. REFERENCES CITED AGAINST CLAIMS 1, 36, AND 60

### A. *Lendenmann*

20. Generally, *Lendenmann* discloses a distributed computing environment ("DCE"), which "is a layer of services that allows distributed applications to communicate with a collection of computers, operating systems, and networks." (*Lendenmann* 7.) As illustrated in Figure 3, *Lendenmann*'s DCE may include several different components, including security services, time services, and directory services. (*Id.* at 8.)



New Bay Capital, LLC
Ex.1015-Page 2609 of 3151

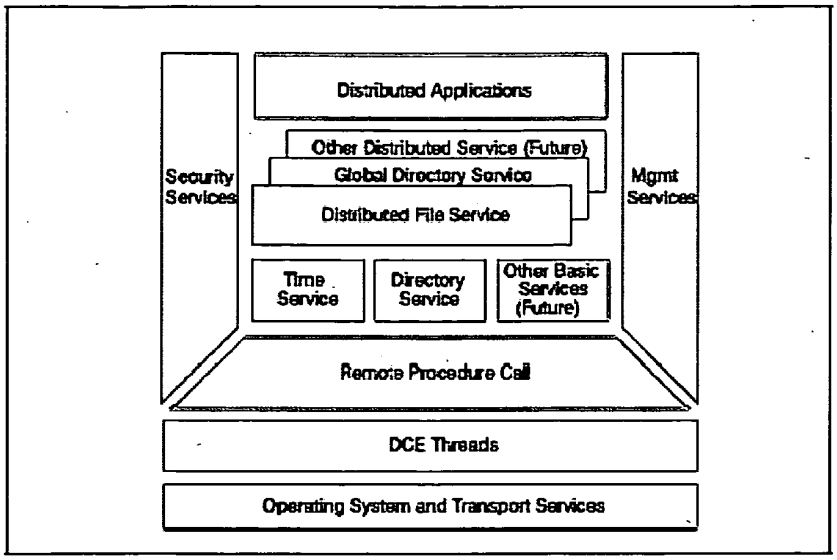


*Figure 3. DCE Architecture*

(*Id.*)

21. It further discloses that a collection of machines, operating systems, and networks managed by a single set of DCE services constitutes a "DCE cell." (*Id.* at 7.) At a minimum, a cell must contain a Security Server, a Cell Directory Service ("CDS"), and Distributed Time Servers. (*Id.* at 9.) These components provide different services for establishing remote procedure calls ("RPCs") between clients and servers.

22. For example, *Lendenmann* explains that the CDS may receive a request and then "return[] the network address of the named resource." (*Id.* at 21, 29.) Additionally, the CDS may provide other identification-related information to the client in a partially complete "binding handle," for example, the appropriate protocol for communications with the server (e.g., TCP/IP, UDP/IP, etc.) or the object UUID. (*Id.* at 182-85.)

23. *Lendenmann* also explains that the Security Service provides services to implement security measures for RPCs between clients and servers, although no security is actually necessary for RPCs. (*Id.* at 191-94, 207, explaining that a client may "optionally set up authenticated RPC.") *Lendenmann* first explains that in order to pursue security-enhanced communications between clients and servers, a client needs to "add[] . . . security information to the server binding handle" received from the CDS. (*Id.* at 191.) Afterwards, a client must obtain a "session key" from the Security Service for implementing any authentication or encryption measures. (*Id.* at 192, 194.) The client may then send the session key to the server, which provides the client with a challenge. (*Id.* at 194.)

If the client matches the challenge criteria, "everything is set for authenticated RPC," and security-enhanced communications may proceed. (*Id.*)

**1. *Lendenmann* Does Not Disclose a Domain Name Service System Configured to Comprise an Indication That the Domain Name Service System Supports Establishing a Secure Communication Link**

24. I understand that independent claim 1 recites a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link. I understand that independent claims 36 and 60 recite similar features. Thus, while I refer to independent claim 1 in the discussion below, the discussion applies equally to independent claims 36 and 60.

25. As I understand the Request, it contends that six different features of *Lendenmann* disclose the indication recited in claim 1. (Req. Ex. F-1 at 13-18.) First, the Request contends that "by returning the network address corresponding to a secure domain name," the CDS provides an indication that the domain name service system supports establishing a secure communication link. (*Id.* at 13, citing *Lendenmann* 21.) The Request also contends that by only performing operations for users authorized by access control lists ("ACLs"), the CDS provides the recited indication. (*Id.* at 13-14, citing *Lendenmann* 34.) The Request also contends that the CDS, by providing clients with binding handles containing identification information for servers, provides the recited indication. (*Id.* at 14-16.) The Request further contends that the authentication challenge disclosed by *Lendenmann* discloses the recited indication. (*Id.* at 16, citing *Lendenmann* 194.) The Request also contends that the server status counters described in *Lendenmann* disclose the recited indication. (*Id.* at 16-17, citing *Lendenmann* 37.) The Request finally contends that *Lendenmann*'s "online documentation" provides the indication recited in claim 1. (*Id.* at 17-18, citing *Lendenmann* 13.)

26. However, it is my opinion these six different features do not disclose or suggest a domain name service system configured to comprise an indication that the domain name service supports establishing a secure communication link. Instead, these features disclose nothing more than a conventional DNS system that is both recognized and distinguished by the '504 patent.

27. First, it is my opinion that the CDS "returning the network address corresponding to a secure domain name" does not disclose the indication recited in claim 1. Instead, *Lendenmann* discloses a conventional name service function for the CDS: "when given a name, CDS returns the network address of the named resource." (*Lendenmann* 21.) As *Lendenmann* explains with Fig. 15 reproduced below, the CDS works as follows:

(1) the client sends a lookup request to the CDS clerk;

(2) the CDS clerk checks its cache and, not finding the name there, contacts the CDS server;

(3) the CDS server checks to see if the name is in the clearinghouse;

(4) the CDS server obtains the requested information if the name exists in the clearinghouse;

(5) the CDS server then returns the information to the CDS clerk; and

(6) the CDS clerk caches the information and passes the requested information to the client.

(*Id.* at 29-30.)

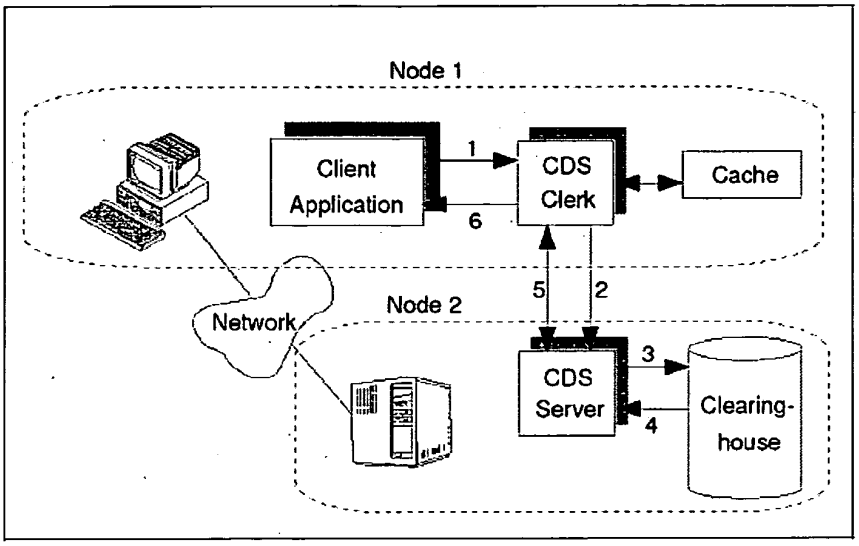


*Figure 15. CDS Components Performing a CDS Look-up*

28. One of ordinary skill would not have understood that *Lendenmann*'s CDS, by returning server identification information to the client, discloses or suggests the indication recited in claim 1 of the '504 patent. Instead, the CDS disclosed in *Lendenmann* is consistent with the conventional domain name service systems that the '504 patent distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link." ('504 patent 39:7-41:59.) As a result, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS feature to disclose or suggest an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1.

29. Second, it is my opinion that *Lendenmann*'s CDS, by only performing operations for users authorized by ACLs, also does not disclose the indication recited in claim 1. *Lendenmann* discloses that the Security Service plays a gatekeeper role to control access to the information in the CDS. (*Lendenmann* 34.) For example, in response to a name request, "ACL management software examines the ACL entry associated with that name or principal name and grants or denies the [CDS]

operation." (*Id.*) The Security Service's gatekeeping function, however, has no bearing on the operations of the alleged domain name service system, the CDS. Even if a client has authorization to receive the information in the CDS, the CDS still merely "returns the network address of the named resource" when given a name in a request, as discussed above. (*Id.* at 21.)

30. By contrast, *Lendenmann* discloses a Security Service function, separate from the CDS, to handle any security-related measures for communications after the CDS has returned server identification information to the client. (*Id.* at 191-94.) For example, the Security Service provides for *Lendenmann*'s authentication procedures, employing a "Security Server" to run the authentication process without any action by or reference to a CDS. (*Id.* at 53-55, Fig. 21.) Although the Security Service may permit or deny access to the CDS, the CDS itself only returns server identification information as illustrated in Figure 15. (*Id.* at 29.) As a result, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS feature to disclose or suggest an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1.

31. Third, it is my opinion that the CDS, by providing clients with binding handles containing identification information for servers, also does not disclose the recited indication. *Lendenmann* discloses that the CDS may provide binding handles to a client to find a target server for a RPC. (*Id.* at 182-85.) These binding handles may provide a network address and other identification-related information, for example, the appropriate protocol for communications with the server (e.g., TCP/IP, UDP/IP, etc.) or the object UUID. (*Id.*)

32. But I disagree that with the Request's implication that the security annotations in the binding handles originated from the CDS. The CDS, in fact, only provides partial binding information. (*Id.* at 184.) During the RPC process, the client first sends a request to the CDS, and the CDS returns a "partly bound" binding handle. (*Id.* at 190.) Then, as *Lendenmann* explains, a client desiring to establish a RPC with certain security measures must first "specify . . . the authentication service, protection level and authorization service that it wants to use in its communications with a server." (*Id.* at 191.) Specifically, "[t]he client [places] a call to rpc#binding#set#auth#info(), which adds this security information to the server binding handle." (*Id.*) And afterwards, "[t]he client then uses this extended binding handle in its further RPC calls." (*Id.*) Therefore, the CDS simply returns identifying information for the target server, and the client must supplement this incomplete binding information in order to implement any desired security measures for communications. As a result, one of ordinary skill in the art would not have understood



New Bay Capital, LLC
Ex.1015-Page 2613 of 3151

*Lendenmann*'s CDS and its incomplete binding handles to disclose or suggest the indication recited in claim 1.

33. Fourth, it is my opinion that the authentication challenge does not disclose the indication recited in claim 1. As I understand the Request, it relies on the *Lendenmann* section concerning key management and mutual authentication, which does not involve the CDS in any manner. (Req. Ex. F-1 at 16; *Lendenmann* 193-94.) During the mutual authentication process, the client's RPC runtime component first requests a service ticket from the Security Service, which contains the session key for the upcoming client/server communication. (*Lendenmann* 194.) After the client sends the session key to the server, the server challenges the client. (*Id.*) Then, if the client successfully matches the challenge criteria, "everything is set for the authenticated RPC." (*Id.*) While the Security Service assists the client and the server during this process, the CDS is conspicuously absent. Thus, contrary to the assertions of the Request, the authentication challenge does not occur "in response" to any name service query. Instead, it occurs in response to the client separately sending the session key to the target server, for which the network address is already known. (*See id.*) One of ordinary skill in the art therefore would not have understood *Lendenmann*'s authentication challenges to disclose or suggest the recited indication.

34. Fifth, it is my opinion that the server status counters do not disclose the indication recited in claim 1. *Lendenmann* discloses that the CDS clerk, CDS server, and CDS clearinghouse all maintain a set of counters "to keep track of the operations performed since it was last started up." (*Id.* at 37.) An example display of such counters is shown below:

```
cdscp> show server
                          SHOW
                         SERVER
                             AT   1995-06-01-10:06:33
                  Creation Time = 1995-05-31-15:24:20.077
              Future Skew Time = 0
               Read Operations = 6409
              Write Operations = 33
              Skulks Initiated = 7
              Skulks Completed = 7
    Times Lookup Paths Broken = 0
              Crucial Replicas = 0
        Child Update Failures = 0
            Security Failures = 0
         Known Clearinghouses = /.../itsc1.austin.ibm.com/ev1_ch
```

(*Id.*)

35. Although the Request asserts that the "Security Failures" counter provides an indication that the CDS supports establishing a secure communication link, nowhere does *Lendenmann* describe the meaning of the "Security Failures" counter or explain that this counter is related in any manner to a secure communication link. This server counter does not provide any indication that the CDS itself supports establishing a secure communication link. Rather, *Lendenmann* provides for a separate Security Service to handle any steps for establishing security measures in communications. (*Id.* at 191-94.) These functions are entirely separate from the conventional name-returning features of the CDS. As one example of this, the authentication process between a client and server is managed by the Security Server, while the CDS plays no role:

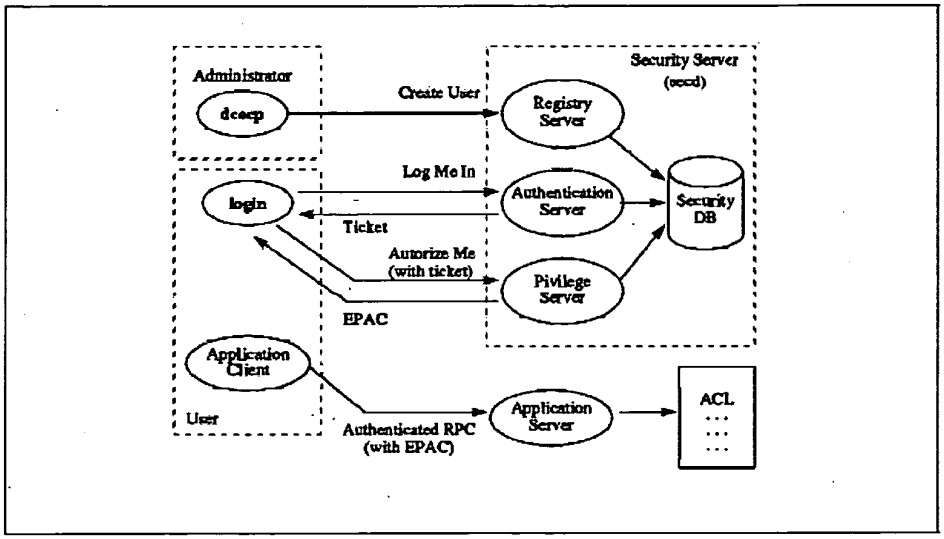


*Figure 21. Authentication Process.*

(*Id.* at 53-55.) The Security Service similarly handles any steps for any authorization or encryption measures, without CDS involvement. (*Id.* at 191-94.) Given the CDS's conventional characteristics, one of ordinary skill in the art would not have understood *Lendenmann*'s CDS and server counters to disclose or suggest the recited indication.

36. Sixth and finally, it is my opinion that the "online documentation" described in *Lendenmann* does not provide the indication recited in claim 1. *Lendenmann* explains that this "online documentation" is nothing more than a collection of DCE manuals provided in softcopy form. (*Id.* at 14.) Thus, as I understand the Request, it is asserting that the various undisclosed DCE manuals briefly and vaguely mentioned in *Lendenmann* somehow provide the indication recited in claim 1. But the mere presence of "online documentation" manuals having indeterminate content does not indicate anything about the CDS, much less that the CDS supports establishing a secure



New Bay Capital, LLC
Ex.1015-Page 2615 of 3151

communication link. Given the dearth of information about this "online documentation" in *Lendenmann*, a person of ordinary skill in the art could not possibly understand how it reflects on the CDS at all.

**B. *Aziz***

37. Generally, *Aziz* discloses a system "for dynamically configuring authorized clients with the address of a protected host and the key and address of an intermediate device (e.g., encrypting firewall, encrypting router, secure gateway) which is protecting a number of hosts on a private network" behind the intermediate device. (*Aziz* 4:3-9.) Fig. 1 of *Aziz*, reproduced below, discloses such a system:

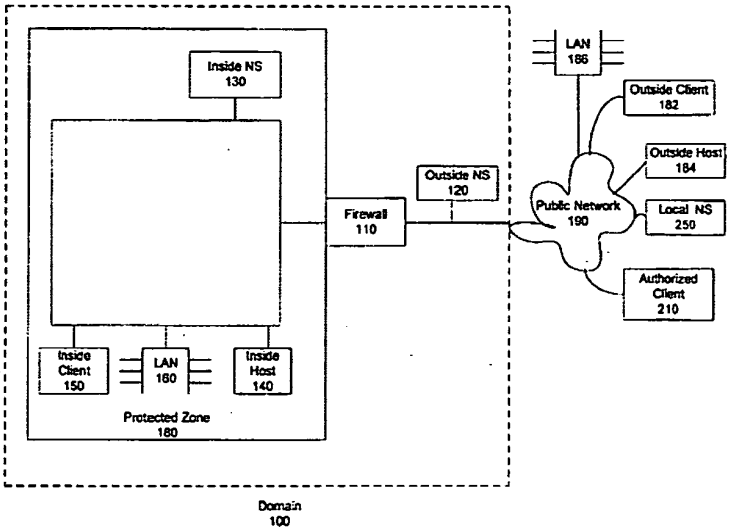


38. *Aziz* explains that "outside NS" 120 may receive a query for a host address located within domain 100 and may determine whether an SX record exists for that host name. (*Id.* at 9:49-53.) An SX record is a DNS resource record that "contains the identifier (e.g., name or address) of a 'secure exchanger,'" such as firewall 110. (*Id.* at 6:23-40.) If an SX record exists, then outside NS 120 may include the SX record in the response to the requester, which may also include the requested host address, if available. (*Id.* at 9:54-10:5.) *Aziz* also discloses that SIG (signature) and KEY resource records may be included in the response. (*Id.* at 9:35-41.)

39. *Aziz* also discloses a resolver 225, which is included in the "authorized client" 210, (*id.* at 8:5-50, Figs. 2A-2C), and receives a response to the query for a host address (*id.* at 10:39-41). If the response includes an SX record and the requested host address, then resolver 225 creates a tunnel map entry that provides the information "authorized client" 210 needs to encrypt messages to "inside host" 140. (*Id.* at 11:13-60.) Resolver 225 then returns the requested host address to an



New Bay Capital, LLC
Ex.1015-Page 2616 of 3151

application 215, also located in "inside host" 210. (*Id.* at 11:55-60.) According to *Aziz*, "[t]his completes the execution" of the configuration process. (*Id.* at 11:60-62.)

**1. *Aziz* Does Not Disclose a Domain Name Service System Configured to Comprise an Indication That the Domain Name Service System Supports Establishing a Secure Communication Link**

40. As I understand the Request, it asserts that four different features of *Aziz* disclose the recited indication. (Req. Ex. F-2 at 7-11.) I also note that the Requester argues that practically all of the elements shown in Fig. 1 of *Aziz* are included in the alleged domain name service system. (*Id.* at 5.) But it is my opinion that one of ordinary skill in the art would not have understood all of these elements in *Aziz* to be included in the recited domain name service system. Such an expansive reading is not consistent with the plain and ordinary meaning of a "domain name service system," which would not include all of the components pointed to by the Requester. (*Id.*) Additionally, this assertion is belied by other parts of the Request, which appear to assume that NS 120 is the recited domain name service system. (*See, e.g.*, *id.* at 7-11, relying on resource records such as SX, KEY, and SIG resource records that are stored in *NS 120* as being the recited "indication.")

41. Regarding the four different features alleged to disclose the recited indication, the Request first asserts that providing SX records in the response from the NS 120 (the alleged domain name service system) provides the indication recited in claim 1. I disagree for the following reasons.

42. The SX record only contains the identifier (e.g., name or address) of a "secure exchanger" (e.g., firewall 110) associated with the owner of the record. (*Aziz* 6:27-38.) Thus, the SX record includes the name or address of firewall 110, which is separate from the alleged domain name service system, NS 120. Although the SX record may include the name or address of firewall 110, the SX record includes no indication about the capabilities of the alleged domain name service system itself or about the capabilities of firewall 110, and certainly does not include an indication that the domain name service system supports establishing a secure communication link.

43. Instead, a person of ordinary skill would have understood that returning an SX record with the name or address of firewall 110 is a feature consistent with the conventional domain name service systems recognized and distinguished by the '504 patent. (*See, e.g.*, '504 patent 39:7-45.) As one example, the '504 patent explains that "[c]onventional Domain Name Servers (DNSs) provide a look-up function that returns the IP address of a requested computer or host[:] . . . when a computer user types in the web name 'Yahoo.com,' the user's web browser transmits a request to a DNS, which converts the name into a four-part IP address that is returned to the user's browser . . . ." (*Id.*



New Bay Capital, LLC
Ex.1015-Page 2617 of 3151

at 39:7-13.) As a result, one of ordinary skill in the art would not have understood *Aziz*'s SX record to disclose or suggest the recited indication.

44. Second, it is my opinion that the KEY and SIG records of *Aziz* do not disclose the indication recited in claim 1. The Request asserts that with these KEY and SIG records, *Aziz* discloses the recited indication "by providing secure DNS service." (Req. Ex. F-2 at 9-10.) I disagree for the following reasons.

45. The KEY and SIG records do not provide an indication that the alleged domain name service system (NS 120) supports establishing a secure communication link. Instead, the KEY and SIG records provided in *Aziz* correspond to resource records, as whenever a resource record is added to a response, the "appropriate SIG and KEY records are also added (i.e., one SIG record for each record type and record owner combination and the KEY record used to generate the SIG record)." (*Aziz* 9:35-41.) These resource records (the A and SX records), however, correspond to the inside host 140 and the firewall 110, which are both separate from the NS 120. A person of ordinary skill would therefore have understood that the KEY and SIG records do not provide any indication about the capabilities of the NS 120 itself, and therefore cannot disclose or suggest the recited indication that the "domain name service system supports establishing a secure communication link." Moreover, returning KEY and SIG resource records is consistent with a conventional domain name system that the '504 patent recognizes and distinguishes. (*See, e.g.*, '504 patent 39:7-42, explaining that conventional domain name systems merely store public keys of different machines so that hosts can request and receive those public keys from the domain name service system.)

46. Third, the Request asserts that "by providing the information needed for secure communications between a client and a protected host," *Aziz* discloses the recited indication. (Req. Ex. F-2 at 8-9.) I disagree for the following reasons.

47. The Request cites to two portions of *Aziz* for this argument. First, the Request argues that the data in the SX record is used by a client for secure communications with protected hosts, and therefore this constitutes the recited indication. (*Id.*, citing *Aziz* 6:57-60.) But as discussed above, the data in the SX record is simply the name or address of the secure exchanger, and providing the SX record with this name or address in *Aziz* does not disclose or suggest the indication recited in claim 1.

48. The second portion of *Aziz* cited by the Request states that the resolver "update[s] a data structure on a client containing information used for secure communications with protected hosts .... Such a data structure comprises data sets whose fields typically contain 'tunnel



New Bay Capital, LLC
Ex.1015-Page 2618 of 3151

information' (e.g., destination and secure exchanger addresses) and related cryptographic data (e.g., secure exchanger's key or algorithm)." (*Id.*, citing *Aziz* 7:28-35.) It is my opinion, for two reasons, that the "data structure on a client containing information used for secure communications" does not disclose a *domain name service system* configured to comprise an indication that the domain name service system supports establishing a secure communication link. First, *Aziz* clearly discloses that "the data structure" is "on a client," which is separate from the alleged domain name service system, NS 120. (*Id.* at 7:30, Fig. 1.) Thus, "the data structure" does not disclose a domain name service system configured to comprise any indication. Second, *Aziz* discloses that the data structure includes "destination and secure exchanger addresses" (i.e., A records and SX records) and secure exchanger keys or algorithms (i.e., SIG and KEY records for the secure exchanger). As discussed above, however, providing these records in *Aziz* is a conventional domain name service system feature. As a result, a person of ordinary skill would have understood that these features of *Aziz* do not disclose or suggest the indication recited in claim 1.

49. Additionally, it is my opinion that the reference to RFC 2065 relied upon by the Request also does not disclose the recited indication. According to the Request, this reference to RFC 2065 discloses "indicating that the connection to the domain name server itself can be encrypted," which, in turn, discloses the recited indication. I disagree.

50. Other than the arguments directed to the KEY and SIG records, which as discussed above cannot disclose the recited indication, the Request relies on RFC 2065 to describe "Bit 8" within the KEY resource record in closer detail. (Req. Ex. F-2 at 10, citing RFC 2065 at 11.) But because the KEY resource record cannot be the recited indication, as discussed above, a component of the KEY resource record also cannot be the recited indication. As a result, a person of ordinary skill would have understood that these features of *Aziz* do not disclose or suggest the indication recited in claim 1.

**C. *Kiuchi* and *Pfaffenberger***

51. Generally, *Kiuchi* proposes a technique called "closed HTTP" (C-HTTP) for providing secure HTTP communications "within a closed group of institutions on the Internet, where each member is protected by its own firewall." (*Kiuchi* 64.) According to *Kiuchi*, C-HTTP is useful in the medical community, where "there is a strong need for closed networks among hospitals and related institutions" to handle patient data and other sensitive medical information. (*Id.*)

52. C-HTTP requires three main components: "1) a client-side proxy on the firewall of one institution, 2) a server-side proxy on the firewall of another institution, and 3) a C-HTTP name

server, which manages a given C-HTTP-based network and the information for [all of its] proxies." (*Id.*) When an institution wants to participate in a C-HTTP network, it must, among other things, install a client-side and/or server-side proxy on its firewall, register an IP address and a hostname for its proxy, and give the proxy's public key to the C-HTTP name server. (*Id.* at 65.) During C-HTTP communications, "[a] client-side proxy and server-side proxy communicate with each other using a secure, encrypted protocol (C-HTTP)." (*Id.* at 64.)

53. When a user agent computer behind a client-side proxy wants to establish a C-HTTP session with a server behind a server-side proxy, the following C-HTTP setup process occurs:

> (1) The client-side proxy asks the C-HTTP name server whether it can communicate with the server.
>
> (2) The C-HTTP name server determines whether the server-side proxy is in the closed network and whether the connection is permitted.
>
> (3) If so, the C-HTTP name server sends the IP address and public key of the server-side proxy, as well as request and response Nonce values, to the client-side proxy.
>
> (4) The client-side proxy sends a connection request to the server-side proxy, encrypted with the server-side proxy's public key.
>
> (5) The server-side proxy asks the C-HTTP name server whether the client-side proxy is also in the closed network and whether the connection is permitted.
>
> (6) If so, the C-HTTP name server sends to the server-side proxy the IP address and public key of the client-side proxy, as well as the same request and response Nonce values previously sent to the client-side proxy.
>
> (7) The server-side proxy then authenticates the client-side proxy, generates a connection ID, generates a second symmetric key for C-HTTP response encryption, and sends this information to the client-side proxy. When the client-side proxy accepts and checks this information, the connection is established.
>
> (8) Once the connection is established, a client-side proxy forwards requests from the user agent in encrypted form using C-HTTP format.

(*Id.* at 65-66.) *Kiuchi* explains that "[t]he [C-HTTP] session is finished when the client accesses another C-HTTP server." (*Id.* at 65.)

54. *Pfaffenberger* is generally a guide that describes how to surf the World Wide Web using the Netscape Navigator web browser, including how to use its various features. (*See generally*

*Pfaffenberger.*) One of the features of Netscape Navigator that *Pfaffenberger* describes is a "doorkey icon," which is displayed in the lower left corner of the web browser to show when the web browser is accessing a secure server. (*See id.* at 9, 13, FIG. 1.2.) While *Pfaffenberger* focuses on the functionality of Netscape Navigator from the standpoint of the end-user, it is not concerned with the inner workings of domain name service systems.

**1. *Kiuchi* and *Pfaffenberger* Do Not Disclose a Domain Name Service System Configured to Comprise an Indication That the Domain Name Service System Supports Establishing a Secure Communication Link**

55. As I understand the Request, it asserts that returning a "public key" of the server-side proxy in *Kiuchi* is an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1. (Req. Ex. F-3 at 11.) It also asserts that returning the IP address of the server-side proxy in *Kiuchi* discloses the recited indication. (*Id.*) The Request also contends that the "doorkey" icon on a web browser disclosed in *Pfaffenberger* provides the recited indication. (*Id.* at 11-12.) Furthermore, the Request asserts that it would have been obvious to combine *Pfaffenberger* with *Kiuchi* to provide the indication recited in claim 1. I disagree for the following reasons.

56. First, it is my opinion that returning the public key of the server-side proxy, as disclosed in *Kiuchi*, does not disclose or suggest the recited indication. The client-side proxy uses the public key of the server-side proxy to send an encrypted connection request to the server-side proxy. (*Kiuchi* 65.) Accordingly, the public key corresponds to the server-side proxy, which is separate from the C-HTTP name server—the alleged domain name service system. The public key of the server-side proxy in *Kiuchi* includes no indication about the capabilities of the *C-HTTP name server itself*, and certainly does not include an indication that the C-HTTP name server supports establishing a secure communication link.

57. Rather, returning a public key of a server-side proxy is consistent with a conventional domain name service system that the '504 patent recognizes and distinguishes from a "domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1. (*See, e.g.*,'504 patent 39:7-45.) In particular, the '504 patent indicates that a conventional domain name service system stores public keys of different machines so that hosts can request and receive those public keys from the domain name service system. (*Id.* at 39:34-42.) As a result, a person of ordinary skill would have understood that these features of *Kiuchi* do not provide the indication recited in claim 1.



New Bay Capital, LLC
Ex.1015-Page 2621 of 3151

58. Second, it is my opinion that *Kiuchi*'s C-HTTP name server returning the IP address of the server-side proxy does not disclose or suggest the indication recited in claim 1. The IP address is for the server-side proxy, which is separate from the alleged domain name service system, the C-HTTP name server. (*Kiuchi* 65.) And this IP address includes no indication about the capabilities of the C-HTTP name server itself, and therefore does not include an indication that the C-HTTP name server supports establishing a secure communication link, as recited in claim 1. Further, like returning a public key, as discussed above, returning the IP address of the server-side proxy in *Kiuchi* is another feature of a conventional domain name system that the '504 patent recognizes and distinguishes from the claimed domain name service system. (*See, e.g.*, '504 patent 39:7-42, explaining that conventional domain name systems merely return a requested IP address corresponding to a provided domain name.) As a result, a person of ordinary skill would have understood that returning this IP address in *Kiuchi* does not provide the indication recited in claim 1.

59. Third, it is my opinion that *Pfaffenberger*'s "doorkey icon" does not provide the indication recited in claim 1. *Pfaffenberger* explains that "[t]his icon indicates whether you're accessing a secure server," but this does not indicate that a *domain name service system* supports establishing a secure communication link in any manner. (*Pfaffenberger* 13.) All the icon indicates is that the web browser is currently accessing a secure server. *Pfaffenberger* is simply a user guide for the Netscape Navigator web browser and does not concern the operation of domain name service systems. Therefore, even if *Kiuchi* and *Pfaffenberger* were combined as proposed by the Request, a person of ordinary skill in the art would not have understood that this combination discloses or suggests a domain name service system configured to comprise an indication that the domain name service system supports establishing a secure communication link, as recited in claim 1.

60. Moreover, it is my opinion that a person of ordinary skill would not have relied on *Pfaffenberger* because *Pfaffenberger* is nonanalogous art to the claimed inventions. One of the problems solved by the embodiments of the '504 patent is that it makes it easy and convenient to enable secure communications using a domain name service system. But *Pfaffenberger* is not reasonably pertinent to solving this problem. As discussed above, *Pfaffenberger* is a guide for how to use the Netscape Navigator web browser, and is not at all concerned with the inner workings of domain name service systems. Thus, a person of ordinary skill would not have relied on *Pfaffenberger* in any manner in attempting to solve the problems recognized and addressed by the '504 patent.



New Bay Capital, LLC
Ex.1015-Page 2622 of 3151

## V. REFERENCES CITED AGAINST CLAIMS 5, 23, AND 47

### A. *Lendenmann*

61. I understand that the Request asserts that by only performing operations for users authorized by ACLs, *Lendenmann*'s CDS discloses a "domain name service system . . . configured to *authenticate the query* [for a network address]" (emphasis added), as recited in claims 5 and 23, and similarly recited in claim 47. (Req. Ex. F-1 at 19-20, 36-37, 55.) For the following reasons, I disagree.

62. The passage of *Lendenmann* cited in the Request discloses that the Security Service—not the CDS—performs the alleged authentication process by acting as a gatekeeper to the information in the CDS. (Req. Ex. F-1 at 19-20, citing *Lendenmann* 34.) Specifically, *Lendenmann* states that "[t]he CDS server only completes an operation over the clearinghouse if the user is authenticated and authorized *by the Security Service*." (*Lendenmann* 34, emphasis added.) As a result, a person of ordinary skill would not have understood the cited feature of *Lendenmann* to disclose or suggest that the purported domain name service system (the CDS) is configured to authenticate the query for a network address, as recited in claims 5, 23, and 47.

63. In addition, I disagree with the Request's contentions that the two cited portions of *Lendenmann* disclose the CDS authenticating a query using a cryptographic technique. (Req. Ex. F-1 at 20-22, citing *Lendenmann* 54-55, 193, 194.) The first portion does not discuss the purported domain name service system, the CDS, at all, but rather the secret-key authentication procedures between the separate Security Service component and principals. (*Lendenmann* 53-55, Fig. 21.) Similarly, the second portion illustrates the security measures performed by the Security Service, without reference to the CDS. (*Id.* at 193, 194.) Accordingly, a person of ordinary skill would not have understood the cited passages of *Lendenmann* to disclose or suggest the purported domain name service system configured to authentication a query using a cryptographic technique.

### B. *Aziz*

64. I understand that the Request also relies on *Aziz* as disclosing that the SIG resource record can be used to authenticate data in other resource records, asserting that this feature "authenticate[s] the query [for a network address]," as recited in claim 5 and similarly recited in claims 23 and 47. I disagree. One of ordinary skill in the art would understand that a domain name service resource record is generally not a query for a network address. Furthermore, simply because a SIG resource record *can* be used to authenticate data does not disclose that the domain name service system itself is configured to authenticate the query for a network address.



New Bay Capital, LLC
Ex.1015-Page 2623 of 3151

65. Additionally, the Request relies on *Aziz* as disclosing that "[a]uthentication means that a host is assured that the message is from the client that the message claims," while listing several standard cryptographic methods. (*Aziz* 3:22-29; *see* Req. Ex. F-2 at 11-13, 28-29, 41.) But a person of ordinary skill in the art would not have understood such general statements about the meaning of authentication to disclose or suggest that "the domain name service system is configured to authenticate the query *using a cryptographic technique*" (emphasis added), as recited in claim 5.

## VI. REFERENCES CITED AGAINST CLAIMS 8-13

### A. *Lendenmann*

66. I understand that the Request contends that *Lendenmann* discloses a domain name service system "connectable to a virtual private network," as recited in claim 8, because a client in one DCE cell can engage in shared-secret key authentication with a client in a foreign DCE cell. I disagree.

67. *Lendenmann* discloses that clients in different cells can communicate with each other while employing the secret key authentication feature. (*Lendenmann* 68.) With this feature, *Lendenmann* illustrates that the Security Server—not the CDS—manages any security measures for subsequent communications, including any authentication, authorization, or encryption. (*Id.* at 53-55; *see also id.* at 191-94.) As one example, the Security Server authenticates users without any action by or reference to a CDS, as illustrated in Figure 21 below:

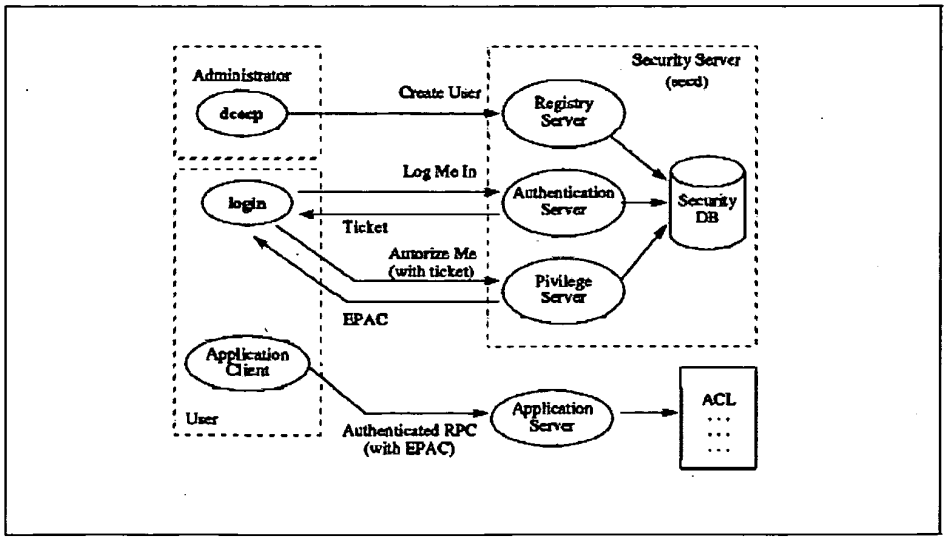


*Figure 21. Authentication Process.*

(*Id.* at 53.) Indeed, throughout this authentication process, *Lendenmann*'s CDS, the alleged domain name server, is conspicuously absent.

68. The Request asserts that because a CDS is "a required part of a DCE 'cell'," and because clients within one cell can communicate with another cell, this satisfies the claim element of a "domain name service system . . . connectable to a virtual private network through the communication network." (Req. Ex. F-1 at 125-26.) A person of ordinary skill in the art, however, would not have interpreted the entire DCE cell of *Lendenmann* to be the recited "domain name service system." *Lendenmann* distinguishes between the subcomponents of a cell and their capabilities, explaining that, "[a]t a minimum, a cell must contain a Security Server, a Cell Directory Server and Distributed Time Servers." (*Lendenmann* 9.) As discussed above, *Lendenmann* provides that the Security Server facilitates and manages security measures for security systems, while the CDS performs no functions beyond those that the '504 patent recognizes as conventional.

69. Additionally, *Gaspoz* would not aid in disclosing or suggesting the elements of claim 8. The Request asserts that because *Gaspoz* discusses implementing a VPN in conjunction with an OSF DCE platform, "it would have been obvious to connect the virtual private network to the Cell Directory Service (CDS)." (Req. Ex. F-1 at 127.) But *Gaspoz* actually discloses that implementing a VPN on a DCE was extraordinarily difficult because of a computer language discrepancy. (*Gaspoz* 257-58.) Although the Request claims that connecting a CDS to a VPN would be obvious because *Gaspoz* discloses that its objects connected to a VPN "were thoroughly represented as DCE servers," *Gaspoz* in fact admits that the object-oriented computer language (C++) made it difficult to utilize objects in establishing a VPN. (*Id.*) *Gaspoz* had to resort to manual file creation to create objects, since it "found no other means for creating persistent objects, because DCE lacks supporting tools in this area." (*Id.*) Accordingly, a person of ordinary skill in the art would not have understood that *Gaspoz* and *Lendenmann* disclose or suggest the elements of claim 8, further recognizing that they should not be combined in the first place.

70. *Gaspoz* also recognized that OSF DCE utilizes a CDS, and instead employed *UserAddr* and *DialList* objects to designate and permit look-ups of the addresses for various other objects within a VPN. (*Id.* at 256-57.) Thus, *Gaspoz* recognized the limitations of employing a VPN together with a DCE, and disclosed using an alternative to a CDS for assisting in establishing VPNs. (*Id.* at 256-59.) And because the CDS disclosed in *Lendenmann* performs no functions other than returning server identification information, which the '504 patent distinguishes as conventional, there would have been no reason to combine *Lendenmann* and *Gaspoz* because *Gaspoz* had already considered the CDS component of the OSF DCE and developed an alternative with the *UserAddr* and *DialList* objects. (*Id.* at 256-57; *see also Lendenmann* 21; '504 patent 39:7-13, 39:34-42.)



New Bay Capital, LLC
Ex.1015-Page 2625 of 3151

**B. *Aziz***

71. I also understand that the Request relies on *Aziz* and *Franaszek* as disclosing or suggesting the elements of claim 9, contending that the connection to inside NS 130 in *Aziz* is "the virtual private network" recited in claim 9. I disagree. The Request asserts that *Aziz* "contemplates imposing some kind of organization on the network structure." (Req. Ex. F-2 at 91.) Meanwhile, the Request asserts that *Franaszek* discloses "[a] hierarchy of multipath networks." (*Id.* at 91-92, quoting *Franaszek* Abstract.) But to one of ordinary skill in the art, simply disclosing "some kind of organization on the network structure" or a "hierarchy of multipath networks" falls far short of disclosing or suggesting that a "virtual private network is one of a plurality of secure communication links in a hierarchy of secure communication links," as recited in claim 9.

72. Moreover, a person of ordinary skill in the art would have recognized that combining *Franaszek* with *Aziz* would have fundamentally changed the operating principles of *Aziz*, and thus would not have made the combination. *Franaszek* is directed to an interconnection system that can be used in a shared-memory computer system, such as a multiple instruction multiple data ("MIMD") computer system, that includes multiple processors interconnected to multiple memory systems. (*See, e.g.*, *Franaszek* 1:20-28, 3:7-31.) *Franaszek* shows an exemplary system in Fig. 1, reproduced below:

FIG. 1

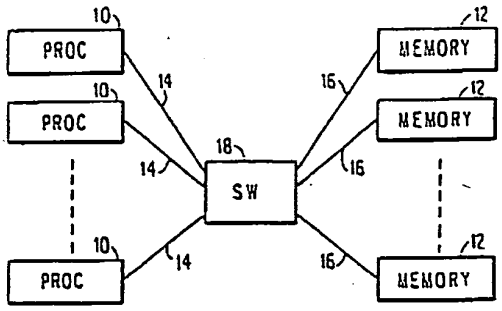


73. The particular interconnection system disclosed by *Franaszek* includes "a hierarchical network comprising multiple levels of multistage networks." (*Id.* at 2:3-5.) In particular, *Franaszek* explains that the "hierarchy of multipath networks" relied on by the Request is a "network hierarchy . . . configured with multiple levels of multistage networks, each one with different message transfer latencies." (*Id.* at 2:13-19.) *Franaszek* admits that these "multistage networks" are known and discloses an exemplary architecture of one multistage network, the Delta network, in Fig. 2 reproduced below. (*Id.* at 3:32-34.) *Franaszek* explains that the Delta network is



New Bay Capital, LLC
Ex.1015-Page 2626 of 3151

"one component of" the hierarchical network included in the disclosed interconnection system. (*Id.* at 2:59-60.)

FIG. 2 PRIOR ART

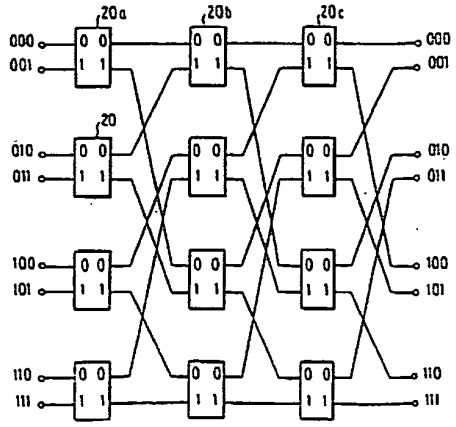


74. The Request proposes combining *Aziz* and *Franaszek* to "organize the secure communication links into a hierarchy as taught by *Franaszek*." (Req. Ex. F-2 at 91.) But the hierarchy taught by *Franaszek* includes a hierarchy of multistage networks, such as the one shown above, that are used for communicating between multiple processors and memories within shared-memory computer systems. Connecting the hierarchy of multistage networks of *Franaszek* would fundamentally change the principle of operation of *Aziz*, which discloses that the devices are connected to each other via traditional computer networking architectures such as LANs, WANs, the Internet, etc. (*Aziz* 4:50-55.) Attempting to incorporate the multistage network architecture described in *Franaszek* with the system of *Aziz* would change the principles of operation of *Aziz* because it would require the computer networks of *Aziz* to be modified to accommodate hardware components traditionally used for connecting processors and memories within a shared-memory computer. Thus, a person of ordinary skill would have recognized that these references are incompatible and should not be combined.

75. Furthermore, a person of ordinary skill in the art would not have relied on *Franaszek* because it is nonanalogous art. One of the problems solved by embodiments of the '504 patent is that it makes it easy and convenient to enable secure communications. But *Franaszek* does not provide any information whatsoever about making it easy and convenient to enable secure communications. Instead, *Franaszek* is concerned with developing a hierarchy technology that allows for an increased number of nodes without increased transfer delays. (*Franaszek* 2:6-12.) Thus, a person of ordinary



New Bay Capital, LLC
Ex.1015-Page 2627 of 3151

skill in the art would not have looked to *Franaszek* to solve the various problems addressed by the '504 patent with its various disclosed embodiments, for example, "a domain name service system configured to . . . comprise an indication that the domain name service system supports establishing a secure communication link," as recited in claim 1.

C. ***Kiuchi* and *Pfaffenberger***

76. I further understand that the Request relies on *Kiuchi* as disclosing claims 8 and 9 because the C-HTTP connection between the proxies corresponds to the virtual private network of claim 8. (*See, e.g.*, Req. Ex. F-3 at 14-16.) I disagree, because one of ordinary skill in the art would not have viewed the C-HTTP connection in *Kiuchi* as a virtual private network in the context of the '504 patent.

77. One of ordinary skill, having read the '504 patent, would have understood a virtual private network, as recited in claims 8-13, to be a network of computers that privately communicate with each other by encrypting traffic on insecure communication paths between the computers. For instance, suppose two computers A and B reside on a public network and two computers X and Y reside on a private network. If A establishes a virtual private network with X and Y's network to address data to X, and B separately establishes a virtual private network with X and Y's network to address data to Y, then A would nevertheless be able to securely address data to B, X, and Y without additional setup. This is true because A, B, X, and Y would all be part of the same virtual private network. *Kiuchi* fails to disclose or suggest such a virtual private network.

78. *Kiuchi*'s C-HTTP connection is different from the claimed virtual private network. In *Kiuchi*'s C-HTTP system, a specific point-to-point connection is established each time a computer is to communicate with another computer using C-HTTP. For instance, suppose, according to *Kiuchi*, a computer A (i.e., user agent or client-side proxy) establishes a C-HTTP session with a computer X (i.e., origin server or server-side proxy) in the closed network, and a computer B (i.e., user agent or client-side proxy) separately establishes a C-HTTP session with a computer Y (i.e., origin server or client-side proxy) in the same closed network. In this case, computer A would be unable to access computer Y via a C-HTTP connection without first ending the existing C-HTTP session with computer X and again engaging in the multistep C-HTTP setup process described above to establish a C-HTTP session with computer Y. (*Kiuchi* 65-66.) Similarly, computer B would be unable to communicate with computer X using C-HTTP without ending its session with computer Y and again completing the setup process to establish a new C-HTTP session with computer X. This is because *Kiuchi*'s C-HTTP connection is a point-to-point connection in which "[t]he [C-HTTP]



New Bay Capital, LLC
Ex.1015-Page 2628 of 3151

session is finished when the client accesses another C-HTTP server." (*Id.* at 65.) In fact, Figure (b) of *Kiuchi* shows that C-HTTP connections to different servers in the closed network have distinct connection IDs:

**b. The HTML document rewritten and forwarded to a use agent by the client-side proxy. The string, "6zdDfldfcZLj8V!i", attached to the end of the URLs is a connection ID**

```
<TITLE>SAMPLE</TITLE>
<BODY>
<A HREF =
"http://server.in.current.connection/sample.html=@
=6zdDfldfcZLj8V!i">
Please click here.</A>
<A HREF =
"http://another.server.in.closed.network/=@=6zdDfl
dfcZLj8V!i">
Another server.</A>
</BODY>
```

(*Id.* at 66.) And, as discussed, the server-side proxy generates the connection ID for a given C-HTTP connection during the above-described C-HTTP session setup process. This confirms that a C-HTTP connection is of a point-to-point nature and requires a new C-HTTP setup process each time a computer is to communicate with another computer using C-HTTP.

79. In light of the additional setup required each time a client communicates with a different server using C-HTTP, one of ordinary skill in the art would have understood *Kiuchi*'s C-HTTP connection as a point-to-point connection rather than the claimed virtual private network. Indeed, for at least this reason, one of ordinary skill in the art would not have understood computers connected via C-HTTP to be part of a virtual private network at all.

80. Additionally, claim 8 specifies that "*the domain name service system* is connectable to a virtual private network through the communication network" (emphasis added). The Request points to the C-HTTP communication between the client-side proxy and the server-side proxy as the claimed virtual private network. (Req. Ex. F-3 at 15.) In particular, as explained by *Kiuchi*, "[o]nce the connection [between the proxies] is established, a client-side proxy forwards HTTP/1.0 requests from the user agent in encrypted form using C-HTTP format." (*Kiuchi* 66.) But this C-HTTP communication between the *proxies*, the alleged virtual private network, does not include the alleged *domain name service system*, the C-HTTP name server. Indeed, the alleged virtual private network in *Kiuchi* is not used for communication until after the proxies have already communicated with the C-HTTP name server to obtain each other's public keys and IP addresses. (*See id.* at 65-66.) Thus,



New Bay Capital, LLC
Ex.1015-Page 2629 of 3151

the alleged domain name service system in *Kiuchi*, the C-HTTP name server, is not connectable to the alleged virtual private network, the C-HTTP communication between the proxies.

81. I also understand that the Request contends that *Kiuchi*'s disclosure of inserting "[r]andom bytes" "every fourth byte of the request and response before encryption in order to avoid the same encrypted requests or responses being repeated" discloses the elements of claim 10, which provides that "the virtual private network is based on inserting into each data packet communicated over a secure communication link one or more data values that vary according to a pseudo-random sequence." (Req. Ex. F-3 at 17, citing *Kiuchi* 72.) I disagree.

82. The passage of *Kiuchi* cited in the Office Action is from Appendix 2A of *Kiuchi*, entitled "The summary of C-HTTP name server protocol." (*Id.* at 72.) As indicated by its title, Appendix 2A relates to name requests/responses between the proxies and the C-HTTP name server. (*Id.*) In the cited passage, *Kiuchi* explains that random bytes are inserted into name requests/responses between the proxies and the C-HTTP name server. (*Id.*) But, as discussed above, the Request takes the position that the encrypted C-HTTP communication between the client-side proxy and server-side proxy corresponds to the claimed virtual private network. Thus, these random bytes are not inserted into data packets of the alleged virtual private network, but into separate name requests/responses between the proxies and the C-HTTP name server. Thus, *Kiuchi* does not disclose that the alleged *virtual private network* is based on inserting into each data packet communicated over a secure communication link one or more data values that vary according to a pseudo-random sequence, as claimed.

83. Furthermore, *Kiuchi* teaches these random bytes are inserted into *C-HTTP name requests and responses*. (*See id.*) But *Kiuchi* does not teach that the random bytes are inserted into *each data packet*. Accordingly, even if the random bytes are viewed as the claimed "one or more data values that vary according to a pseudo-random sequence," *Kiuchi* still does not disclose that they are inserted into *each data packet*, as claimed. Thus, a person of ordinary skill would have understood that these features of *Kiuchi* fail to disclose or suggest the elements of claim 10.

84. The Request also contends that the Nonce values contained in the headers of *Kiuchi*'s C-HTTP requests and responses constitute the "value in each data packet," as recited in claim 12. (Req. Ex. F-3 at 17-19.) I disagree.

85. To support its position, the Request contends that "[i]n the Examples of C-HTTP communication found in Appendix 3, it can be seen that the 'Request-Nonce value' is incremented, moving from '8abd853f' in Example c., to '8abd8540' in Example g., to '8abd8541' in Example i."



New Bay Capital, LLC
Ex.1015-Page 2630 of 3151

(*Id.* at 18.) According to *Kiuchi*, Example c. is a "Request for connection to the server-side proxy," Example g. is "Sending C-HTTP requests to the server-side proxy," and Example i. is a "Request for closing the connection." (*Kiuchi* 74.) But the Request has just shown that different requests contain different Nonce values. Indeed, in secure communications, a nonce is a unique, arbitrary number used only once to identify a particular communication.

86. But *Kiuchi* does not disclose that these Nonce values are compared to a moving window of valid values, as recited in claim 12. *Kiuchi* mentions that the "[r]eplay attacks are blocked by checking values of the Request-Nonce header field." (*Kiuchi* 65.) But *Kiuchi* does not explain *how* the values of the Nonce header field are checked, and certainly does not teach that they are checked by comparing them to a moving window of valid values. There are many ways that the Nonce values could be checked without comparing them to a moving window of valid values.

87. In addition, *Kiuchi* teaches that Nonce values are inserted into C-HTTP requests and responses, not into *each data packet*. (*See Kiuchi* 65, 71.) Accordingly, even if the Nonce values were compared to a moving window of valid values (which they are not), *Kiuchi* still would not disclose that the virtual private network is based on comparing a value in *each data packet* transmitted between the first computer and the second computer to a moving window of valid values, as claimed. One of ordinary skill in the art would not have considered *Kiuchi*'s C-HTTP requests and responses, which are application layer requests, to be data packets. Thus, a person of ordinary skill would have understood that these features of *Kiuchi* fail to disclose or suggest the elements of claim 12.

## VII. REFERENCES CITED AGAINST CLAIMS 18 AND 42

88. I understand that the Request also relies on *Aziz* as disclosing that "one of the plurality of domain names is reserved for secure communication links," as recited in claim 18 and similarly recited in claim 42. The Request asserts that having a domain name "zone" configured to ensure that communications are secure demonstrates these claim features. I disagree, because *Aziz* does not disclose reserving any domain names for secure communication links, and the portions of *Aziz* relied on by the Request do not support the Request's assertions.

89. I understand that the Request cites to three portions in the Background section of *Aziz* as allegedly disclosing the claimed feature. (Req. Ex. F-2 at 24-25, citing *Aziz* 1:58-68, 2:9-22, 3:14-24.) The first portion discloses that organizations may divide a domain into zones and that each zone may have a database that contains names, addresses, and other information for that zone. (*Aziz* 1:58-68.) The second portion discloses that the organization can configure the registered name servers for



New Bay Capital, LLC
Ex.1015-Page 2631 of 3151

the zone to either contain information regarding a machine (i.e., in a publicly visible zone) or not contain information regarding the machine (i.e., in a visibility-limited or "protected" zone). (*Id.* at 2:9-22.) The Request asserts that this portion teaches that "machines in one zone can be reserved for secure communication links by being put in a 'protected zone.'" (Req. Ex. F-2 at 24-25.) But this portion of *Aziz* is silent regarding reserving machines for secure communication links, let alone that a domain name is reserved for secure communication links. The third portion of *Aziz* relied on by the Request simply discloses that when a "client and host want to ensure that their communications are 'secure,'" the client may require information in addition to the address of the host. (*Aziz* 3:14-24.) This portion also does not disclose that a domain name is reserved for secure communication links. Thus, a person of ordinary skill would not have understood these portions of *Aziz* to disclose or suggest the elements of claims 18 and 42, as these portions are silent regarding reserving domain names.

## VIII. REFERENCES CITED AGAINST CLAIMS 24 AND 48

90. I understand that the Request also relies on *Aziz* as disclosing these claim elements because "[h]aving a machine be in the visibility-limited 'protected zone' is an indication that the domain name service system supports establishing a secure communication link." (Req. Ex. F-2 at 30, citing *Aziz* 2:9-22.) I disagree.

91. The portion of *Aziz* relied on by the Request only discloses that if a machine is in a "visibility-limited" or "protected zone," the registered name servers may not have any information about the zone name servers for that zone. (*Aziz* 2:9-22.) It is silent regarding what the domain names of machines within that zone include, and fails to disclose that any one of the plurality of domain names comprises an indication that the domain name service supports establishing a secure communication link. The Request's assertion that "[h]aving a machine be in the visibility-limited 'protected zone' is an indication . . ." says nothing about what a domain name comprises. (Req. Ex. F-2 at 30, citing *Aziz* 2:9-22.) Rather, just because a machine is in a protected zone does not mean that the domain name of the machine comprises an indication that a domain name service system supports establishing a secure communication link. As a result, a person of ordinary skill would have understood that *Aziz* does not disclose or suggest these elements of claims 24 and 48.

**Truth and Accuracy of Statements**

I further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that willful false statements or the like may jeopardize the validity of the '504 patent.

Signed at New York, New York, this 31th day of May, 2012.

/Angelos D. Keromytis/
Angelos D. Keromytis



New Bay Capital, LLC
Ex.1015-Page 2633 of 3151

**Angelos D. Keromytis - *Curriculum Vitae***

**Positions Held**

- **January 2006 - Present**
  Associate Professor, Department of Computer Science, Columbia University, New York.
- **January 2009 - January 2010**
  Senior Research Engineer, Symantec Research Labs Europe, Sophia Antipolis, France.
- **July 2001 - December 2005**
  Assistant Professor, Department of Computer Science, Columbia University, New York.
- **September 1996 - July 2001**
  Research Assistant, Computer and Information Science Department, University of Pennsylvania, Philadelphia.
- **January 1993 - October 1995**
  Member of the Technical Staff, FORTHnet S.A., Heraclion, Greece.
- **September 1991 - January 1993**
  Member of the Technical Staff, Education Team, Computer Center of the University of Crete, Heraclion, Greece.

**Education**

- **November 2001**
  Ph.D. (Computer Science), University of Pennsylvania, USA.
- **August 1997**
  M.Sc. (Computer Science), University of Pennsylvania, USA.
- **June 1996**
  B.Sc. (Computer Science), University of Crete, Greece.

**Service and Teaching**

**Editorial Boards and Steering Committees**

- Associate Editor, Encyclopedia of Cryptography and Security (2nd Edition), Springer, 2010 - 2011.
- Associate Editor, IET (formerly IEE) Proceedings Information Security, 2005 - 2010.
- Steering Committee, *ISOC Symposium on Network and Distributed System Security (SNDSS)*, 2006 - 2009.
- Steering Committee, *New Security Paradigms Workshop (NSPW)*, 2007 onward.
- Associate Editor, ACM Transactions on Information and System Security (TISSEC), 2004 - 2010.
- Steering Committee, *USENIX Workshop on Hot Topics in Security (HotSec)*, 2006 - 2009.
- Steering Committee, *Computer Security Architecture Workshop (CSAW)*, 2007 - 2009.

**Program Chair**

- Program Chair, 16th International Conference on Financial Cryptography and Data Security (FC), 2012.
- Program co-Chair, 17th ACM Computer and Communication Security (CCS), 2010.
- Program co-Chair, 16th ACM Computer and Communication Security (CCS), 2009.



New Bay Capital, LLC
Ex.1015-Page 2634 of 3151

- Program co-Chair, New Security Paradigms Workshop (NSPW), 2008.
- Program co-Chair, New Security Paradigms Workshop (NSPW), 2007.
- Chair, 27th International Conference on Distributed Computing Systems (ICDCS), *Security Track,* 2007.
- Chair, 16th World Wide Web (WWW) Conference, *Security, Privacy, Reliability and Ethics Track,* 2007.
- Chair, 15th USENIX Security Symposium, 2006.
- Deputy Chair, 15th World Wide Web (WWW) Conference, *Security, Privacy and Ethics Track,* 2006.
- Chair, 3rd Workshop on Rapid Malcode (WORM), 2005.
- Program co-Chair, 3rd Applied Cryptography and Network Security (ACNS) Conference, 2005.
- Program co-Chair, OpenSig Workshop, 2003.

**Program Organization**

- General Chair, New Security Paradigms Workshop (NSPW), 2010.
- General Vice Chair, New Security Paradigms Workshop (NSPW), 2009.
- Co-chair, Invited Talks, 17th USENIX Security Symposium, 2008.
- General co-chair, Applied Cryptography and Network Security (ACNS) Conference, 2008.
- Co-chair, Invited Talks, 16th USENIX Security Symposium, 2007.
- Organizing Committee, Columbia/IBM/Stevens Security & Privacy Day (bi-annual event).
    - Organizer, Columbia/IBM/Stevens Security & Privacy Day, December 2010.
    - Organizer, Columbia/IBM/Stevens Security & Privacy Day, June 2007.
- Co-organizer, ARO/FSTC Workshop on Insider Attack and Cyber Security, 2007.
- Publicity co-Chair, ACM Conference on Computer and Communications Security, 2006.
- General co-Chair, OpenSig Workshop, 2003.

**Program Committees**

- Program Committee, ISOC Symposium on Network and Distributed Systems Security (SNDSS), 2003, 2004, 2006, 2007, 2008, 2012.
- Program Committee, International Workshop on Security (IWSEC), 2006, 2007, 2008, 2009, 2010, 2011.
- Program Committee, ACM Conference on Computer and Communications Security (CCS), 2005, 2007, 2008, 2009, 2010.
- Program Committee, Applied Cryptography and Network Security (ACNS) Conference, 2005, 2006, 2010, 2011, 2012.
- Program Committee, USENIX Security Symposium, 2004, 2005, 2006, 2008.
- Program Committee, International Conference on Distributed Computing Systems (ICDCS), *Security Track*, 2005, 2006, 2007, 2008.
- Program Committee, Workshop on Rapid Malcode (WORM), 2004, 2005, 2006, 2007.
- Program Committee, Information Security Conference (ISC), 2005, 2007, 2009, 2011.
- Program Committee, World Wide Web Conference (WWW), 2005, 2006, 2007.
- Program Committee, USENIX Workshop on Hot Topics in Security (HotSec), 2006, 2007, 2010.
- Program Committee, Financial Cryptography (FC) Conference, 2002, 2010, 2011, 2012.
- Program Committee, European Workshop on Systems Security (EuroSec), 2009, 2010, 2011.



New Bay Capital, LLC
Ex.1015-Page 2635 of 3151

- Program Committee, Annual Computer Security Applications Conference (ACSAC), 2006, 2007, 2011.
- Program Committee, USENIX Technical Conference, *Freely Distributable Software (Freenix) Track,* 1998, 1999, 2003.
- Program Committee, IEEE Security & Privacy Symposium, 2006, 2008.
- Program Committee, ACM SIGCOMM Workshop on Large Scale Attack Defense (LSAD), 2006, 2007.
- Program Committee, New Security Paradigms Workshop (NSPW), 2007, 2008.
- Program Committee, IEEE WETICE Workshop on Enterprise Security, 2002, 2003.
- Program Committee, International Conference on Mathematical Methods, Models and Architectures for Computer Network Security (MMM-ACNS), 2007, 2010.
- Program Committee, USENIX Annual Technical Conference (ATC), 2008, 2011.
- Program Committee, European Sumposium on Research in Computer Security (ESORICS), 2011.
- Program Committee, International Workshop on Mobile Security (WMS), 2010.
- Program Committee, 40th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Dependable Computing and Communication Symposium (DCCS), 2010.
- Program Committee, Computer Forensics in Software Engineering Workshop, 2009.
- Program Committee, USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET), 2008.
- Program Committee, 23rd International Information Security Conference (IFIP SEC), 2008.
- Program Committee, Joint iTrust and PST Conferences on Privacy, Trust Management and Security (IFIPTM), 2008.
- Program Committee, 1st Computer Security Architecture Workshop (CSAW), 2007.
- Program Committee, 8th IEEE Information Assurance Workshop (IAW), 2007.
- Program Committee, Anti-Phishing Working Group (APWG) eCrime Researchers Summit, 2007.
- Program Committee, 4th GI International Conference on Detection of Intrusions & Malware, and Vulnerability Assessment (DIMVA), 2007.
- Program Committee, 2nd ACM Symposium on InformAtion, Computer and Communications Security (AsiaCCS), 2007.
- Program Committee, 6th International Conference on Cryptology and Network Security (CANS), 2007.
- Program Committee, 2nd Workshop on Advances in Trusted Computing (WATC), 2006.
- Program Committee, International Conference on Information and Communications Security (ICICS), 2006.
- Program Committee, 2nd Workshop on Secure Network Protocols (NPSec), 2006.
- Program Committee, 1st Workshop on Hot Topics in System Dependability (HotDep), 2005.
- Program Committee, 20th ACM Symposium on Applied Computing (SAC), Trust, Recommendations, Evidence and other Collaboration Know-how (TRECK) Track, 2005.
- Program Committee, 1st Workshop on Operating System and Architecture Support for the on demand IT Infrastructure (OASIS), 2004.
- Program Committee, Workshop on Information Security Applications (WISA), 2004.
- Program Committee, Workshop on Logical Foundations of an Adaptive Security Infrastructure (WOLFASI), 2004.
- Program Committee, 29th IEEE Conference on Local Computer Networks (LCN), 2004.
- Program Committee, 2nd International Conference on Trust Management, 2004.



New Bay Capital, LLC
Ex.1015-Page 2636 of 3151

- Program Committee, Asia BSD Conference, 2004.
- Program Committee, 2nd Annual New York Metro Area Networking Workshop (NYMAN), 2002.
- Program Committee, Cloud Computing Security Workshop (CCSW), 2009.
- Program Committee, Workshop on Grid and Cloud Security (WGC-Sec), 2011.
- Program Committee, Workshop on Cyber Security Experimentation and Test (CSET), 2011.

**Advisory Workshops**

- ODNI/NSA Invitational Workshop on Computational Cybersecurity in Compromised Environments (C3E), Keystone, CO, September 2011.
- ONR Workshop on Host Computer Security, Chicago, IL, October 2010.
- Intel Workshop on Trust Evidence and End-to-end Trust in Heterogeneous Environments, Santa Clara, CA, May 2010.
- Intelligence Community Technical Exchange on Moving Target, Washington, DC, April 2010.
- Lockheed Martin Future Security Threats Workshop, New York, NY, November 2009.
- Air Force Office for Scientific Research (AFOSR) Invitational Workshop on Homogeneous Enclave Software *vs* Heterogeneous Enclave Software, Arlington, VA, October 2007.
- NSF Future Internet Network Design Working Meeting, Arlington, VA, June 2007.
- ARO/FSTC Workshop on Insider Attack and Cyber Security, Arlington, VA, June 2007.
- NSF Invitational Workshop on Future Directions for the CyberTrust Program, Pittsburgh, PA, October 2006.
- ARO/HSARPA Invitational Workshop on Malware Detection, Arlington, VA, August 2005.
- Department of Defense Invitational Workshop on the Complex Behavior of Adaptive, Network-Centric Systems, College Park, MD, July 2005.
- ARDA Next Generation Malware Invitational Workshop, Annapolis Junction, MD, March 2005.
- Co-leader of session on "Securing software environments", joint NSF and Department of Treasury Invitational Workshop on Resilient Financial Information Systems, Washington, DC, March 2005.
- DARPA Application Communities Invitational Workshop, Arlington, VA, October 2004.
- DARPA APNets Invitational Workshop, Philadelphia, PA, December 2003.
- NSF/NIST Invitational Workshop on Cybersecurity Workforce Needs Assessment and Educational Innovation, Arlington, VA, August 2003.
- NSF Invitational Workshop on Large Scale Cyber-Security, Lansdowne, VA, March 2003.
- IP Security Working Group Secretary, Internet Engineering Task Force (IETF), 2003 - 2008.
- Session moderator, Workshop on Intelligence and Research, Florham Park, NJ, October 2001.
- DARPA Composable High Assurance Trusted Systems #2 (CHATS2) Invitational Workshop, Napa, CA, November 2000.

**Other Professional Activities**

- Co-chair, ACM Computing Classification System Update Committee ("Security and Privacy" top-level node), 2011.
- Member, ACM Computing Classification System Update Committee (top two levels), 2010.
- External Advisory Board member, *"i-code: Real-time Malicious Code Identification"*, EU



New Bay Capital, LLC
Ex.1015-Page 2637 of 3151

project, 2010 - 2012.

- Reviewer (grant applications), Greek Ministry of Education, 2010.
- Reviewer (grant applications), Danish National Research Foundation, 2010.
- Member of the Scientific Advisory Board, Centre for Research and Technology, Hellas (CERTH), 2008 - 2011.
- Senior Member of the ACM, 2008 onward.
- Senior Member of the IEEE, 2009 onward.
- Visiting Scientist, Institute for Infocomm Research ($I^2R$), Singapore, February - May 2007.
- Columbia Representative to the Institute for Information Infrastructure Protection (I3P), 2006 - 2008.
- Technical Advisory Board, *StackSafe Inc. (formerly Revive Systems Inc.),* 2006 - 2009.
- Technical Advisory Board, *Radiuz Inc.,* 2006.
- Reviewer (grant applications), Institute for Security Technology Studies (ISTS), Dartmouth College, 2006.
- Reviewer, Singapore National Science and Technology Awards (NSTA), 2006.
- Board of Directors, *StackSafe Inc.(formerly Revive Systems Inc.),* 2005 - 2009.
- Founder, *StackSafe Inc. (formerly Revive Systems Inc.),* 2005 - 2009.
- Expert witness in criminal and intellectual property litigation cases, 2005, 2006, 2007, 2009, 2010, 2011.
- Science Fair Judge, Middle School for Democracy and Leadership, Brooklyn, NY, 2005, 2006.
- Reviewer (grant applications), Swiss National Science Foundation, 2007.
- Reviewer (grant applications), Netherlands Organisation for Scientific Research, 2005, 2006.
- Reviewer (grant applications), US/Israel Binational Science Foundation, 2003, 2005.
- NSF reviewer & panelist, 2002, 2003, 2006, 2008, 2009, 2011.
- Internet Engineering Task Force (IETF) Security Area Advisor, 2001 - 2008.

**Ph.D. Thesis Committee Service**

- Michalis Polychronakis, *"Generic Code Injection Attack Detection using Code Emulation",* Computer Science Department, University of Crete, October 2009.
- Spyros Antonatos, *"Defending against Known and Unknown Attacks using a Network of Affined Honeypots",* Computer Science Department, University of Crete, October 2009.
- Van-Hau Pham, *"Honeypot Traces Forensics by Means of Attack Event Identification",* Computer Science Group, Communications and Electronics Department, Ecole Nationale Superieure des Telecommunications, September 2009.
- Gabriela F. Ciocarlie, *"Towards Self-Adaptive Anomaly Detection Sensors",* Department of Computer Science, Columbia University, September 2009.
- Vanessa Frias-Martinez, *"Behavior-Based Admission and Access Control for Network Security",* Department of Computer Science, Columbia University, September 2008.
- Wei-Jen Li, *"SPARSE: A Hybrid System for Malcode-Bearing Document Detection",* Department of Computer Science, Columbia University, June 2008.
- Raj Kumar Rajendran, *"The Method for Strong Detection for Distributed Routing",* Electrical Engineering Department, Columbia University, March 2008.
- Constantin Serban, *"Advances in Decentralized and Stateful Access Control",* Computer Science Department, Rutgers University, December 2007.
- Ricardo A. Baratto, *"THINC: A Virtual and Remote Display Architecture for Desktop Computing",* Computer Science Department, Columbia University, October 2007.



New Bay Capital, LLC
Ex.1015-Page 2638 of 3151

- Zhenkai Liang, *"Techniques in Automated Cyber-Attack Response and Recovery"*, Computer Science Department, Stony Brook University, November 2006.
- Ke Wang, *"Network Payload-based Anomaly Detection and Content-based Alert Correlation"*, Computer Science Department, Columbia University, August 2006.
- Seoung-Bum Lee, *"Adaptive Quality of Service for Wireless Ad hoc Networks"*, Electrical Engineering Department, Columbia University, June 2006.
- Shlomo Hershkop, *"Behavior-based Email Analysis with Application to Spam Detection"*, Computer Science Department, Columbia University, August 2005.
- Gaurav S. Kc, *"Defending Software Against Process-subversion Attacks"*, Computer Science Department, Columbia University, April 2005.
- Gong Su, *"MOVE: A New Virtualization Approach to Mobile Communication"*, Computer Science Department, Columbia University, May 2004.
- Jonathan M. Lennox, *"Services for Internet Telephony"*, Computer Science Department, Columbia University, December 2003.
- Michael E. Kounavis, *"Programming Network Architectures"*, Electrical Engineering Department, Columbia University, June 2003.
- Wenyu Jiang, *"QoS Measurement and Management for Internet Real-time Multimedia Services"*, Computer Science Department, Columbia University, April 2003.

**Post-doctoral Students**

- Hyung Chan Kim (October 2007 - October 2008)
- Stelios Sidiroglou (October 2008 - December 2008)
- Georgios Portokalidis (March 2010 - present)
- Michalis Polychronakis (May 2010 - present)
- Dimitris Geneiatakis (June 2010 - present)

**Current Ph.D. Students**

- Georgios Kontaxis (September 2011)
- Vasilis Pappas (September 2009 - present)
- Vasileios Kemerlis (September 2008 - present)
- Kangkook Jee (January 2008 - present)
- Sambuddho Chakravarty (January 2007 - present)
- Angelika Zavou (September 2006 - present)

**Graduated Ph.D. Students**

- Debra Cook (January 2002 - June 2006)
  - Thesis title: *"Elastic Block Ciphers"*
  - Post-graduation: Member of the Technical Staff, Bell Labs
  - Currently: Research Staff Member, Telcordia Research
- Angelos Stavrou (January 2003 - August 2007)
  - Thesis title: *"An Overlay Architecture for End-to-End Service Availability"* (awarded with distinction)
  - Post-graduation: Assistant Professor, Computer Science Department, George Mason University (GMU)



New Bay Capital, LLC
Ex.1015-Page 2639 of 3151

    - Currently: Assistant Professor, Computer Science Department, George Mason University (GMU)
- Michael E. Locasto (September 2002 - December 2007)
    - Thesis title: *"Integrity Postures for Software Self-Defense"* (awarded with distinction)
    - Post-graduation: ISTS Research Fellow, Dartmouth College
    - Currently: Assistant Professor, Department of Computer Science, University of Calgary
- Stelios Sidiroglou (June 2003 - May 2008)
    - Thesis title: *"Software Self-healing Using Error Virtualization"*
    - Post-graduation: Research Scientist, Columbia University
    - Currently: Research Scientist, MIT CSAIL
- Mansoor Alicherry (September 2006 - October 2010)
    - Thesis title: *"A Distributed Policy Enforcement Architecture for Mobile Ad Hoc Networks"*
    - Post-graduation: Member of the Technical Staff, Alcatel-Lucent Bell Labs
    - Currently: Member of the Technical Staff, Alcatel-Lucent Bell Labs
- Brian Bowen (September 2007 - December 2010; co-advised with Salvatore J. Stolfo)
    - Thesis title: *"Design and Analysis of Decoy Systems for Computer Security"*
    - Post-graduation: Member of the Technical Staff, Sandia National Laboratories
    - Currently: Member of the Technical Staff, Sandia National Laboratories

**Service at Columbia**

- Computer Science Department Ph.D. Committee, 2010 - 2011
- Computer Science Department Facilities committee, 2001 - 2008, 2010 - current
    - Chair, Facilities committee, 2003 - 2005, 2011 - current
- M.Sc. Admissions committee, 2007 - current.
- M.Sc. Committee, 2008 - current.
- Computer Science Department Faculty Recruiting committee, 2002, 2008
- Columbia committee on Research Conflict of Interest Policy, 2007 - 2008
- Co-organizer, Computer Science Faculty Retreat, Fall 2007
- Advisor for the School of Engineering Computer Science Majors, Freshmen & Sophomores, 2004 - 2005
- Computer Science Department Undergraduate Admissions Representative, 2003 - 2008
- Advisor for the School of Engineering Computer Science Majors, Seniors, 2003 - 2004, 2006 - 2007
- Computer Science Department Space Allocation Policy committee, 2002 - 2010
- Computer Science Department Events Representative, 2002 - 2008
- Advisor for the School of Engineering Computer Science Majors, Juniors, 2002 - 2003, 2005 - 2006
- Computer Science Department CRF Director Hiring committee, 2003
- Advisor for the School of Engineering Computer Science Majors, Sophomores, 2001 - 2002
- Computer Science Department Faculty Recruiting committee, 2001 - 2002
- Executive Vice Provost committee on Columbia's response to the 9/11 events, Fall 2001

**Teaching**

*(Scores indicate mean course quality rating from student survey; survey not conducted for summer*



New Bay Capital, LLC
Ex.1015-Page 2640 of 3151

*sessions)*

- Instructor, COMS E6183-1 - Advanced Topics in Network Security, Columbia University
  - Fall 2006: 17 on-campus students *(4.58/5)*
- Instructor, COMS W6998.1 - Advanced Topics in Network Security, Columbia University
  - Fall 2004: 17 on-campus students *(4.62/5)*
  - Spring 2003: 18 on-campus students *(N/A)*
- Instructor, COMS W4180 - Network Security, Columbia University
  - Spring 2011: 4 CVN students *(N/A)*
  - Fall 2010: 2 CVN students *(N/A)*
  - Spring 2010: 25 on-campus and 5 CVN students *(4.48/5)*
  - Summer 2006: 7 CVN students *(N/A)*
  - Spring 2006: 63 on-campus and 9 CVN students *(4.14/5)*
  - Summer 2005: 4 CVN students *(N/A)*
  - Spring 2005: 41 on-campus and 5 CVN students *(4.25/5)*
  - Summer 2004: 6 CVN students *(N/A)*
  - Fall 2003: 45 on-campus and 12 CVN students *(3.74/5)*
  - Summer 2003: 5 CVN students *(N/A)*
  - Fall 2002: 43 on-campus and 9 CVN students *(3.21/5)*
  - Fall 2001: 23 on-campus students *(3.6/5)*
- Instructor, COMS W4118 - Operating Systems, Columbia University
  - Summer 2007: 8 CVN students *(N/A)*
  - Fall 2006: 59 on-campus and 7 CVN students *(3.73/5)*
  - Summer 2006: 15 CVN students *(N/A)*
  - Fall 2005: 52 on-campus and 9 CVN students *(3.86/5)*
  - Spring 2004: 32 on-campus and 4 CVN students *(3.39/5)*
  - Spring 2002: 37 on-campus students *(3.13/5)*
- Instructor, COMS W3157 - Advanced Programming, Columbia University
  - Fall 2010: 37 on-campus students *(3.25/5)*
  - Fall 2007: 30 on-campus students *(4.16/5)*
- Instructor, CIS700/002 - Building Secure Systems, University of Pennsylvania, Spring 1998

**Support for Research and Teaching (Gifts and Grants)**

1. PI (co-PIs: Roxana Geambasu, Junfeng Yang, Simha Sethumadhavan, Sal Stolfo), *"MEERKATS: Maintaining EnterprisE Resiliency via Kaleidoscopic Adaptation & Transformation of Software Services"*, DARPA MRC, **$6,619,270** (09/2011 - 09/2015; leading team that includes George Mason University and Symantec Corp.)
2. PI, *"NSF Support for the 2011 New Security Paradigms Workshop Financial Aid (Supplement)"*, NSF Trustworthy Computing, **$10,000** (06/2011 - 07/2012)
3. PI, *"Leveraging the Cloud to Audit Use of Sensitive Infomation"*, Google (research gift), **$60,200** (05/2011)
4. co-PI (with Sal Stolfo), *"ADAMS Advanced Behavioral Sensors (ABS)"*, DARPA ADAMS, **$780,996** (05/2011 - 04/2013)
5. PI, *"Tracking Sensitive Information Flows in Modern Enterprises"*, Intel, **$84,951** (12/2010 - 12/2011)
6. co-PI (with Simha Sethumadhavan, Sal Stolfo, Junfeng Yang, and David August @ Princeton), *"SPARCHS: Symbiotic, Polymorphic, Autotomic, Resilient, Clean-slate, Host*



New Bay Capital, LLC
Ex.1015-Page 2641 of 3151

*Security"*, DARPA CRASH, **$6,424,180** (10/2010 - 09/2014)

7. PI, *"NSF Support for the 2010 New Security Paradigms Workshop Financial Aid"*, NSF Trustworthy Computing, **$10,000** (09/2010 - 08/2011)
8. PI (co-PIs: Junfeng Yang, Sal Stolfo), *"MINESTRONE"*, IARPA, **$7,530,113** (08/2010 - 07/2014; leading team that includes Stanford University, George Mason University, and Symantec Corp.)
9. co-PI (with Junfeng Yang and Dawson Engler @ Stanford), *"Seed: CSR: Large: Collaborative Research: SemGrep: Improving Software Reliability Through Semantic Similarity Bug Search"*, NSF CSR, CNS-10-12107, **$325,000** (07/2010 - 06/2011)
10. PI, *"Tracking Sensitive Information Flows in Modern Enterprises"*, Intel, **$82,286** (08/2009 - 07/2010)
11. PI, *"Supplement for International Research Collaborations"*, NSF Trustworthy Computing, $41,769 (09/2009 - 08/2011)
12. PI, *"NSF Support for the 2009 New Security Paradigms Workshop Financial Aid"*, NSF Trustworthy Computing, **$10,000** (09/2009 - 08/2010)
13. PI, *"Measuring the Health of Internet Routing: A Longitudinal Study"*, Google (research gift), **$60,000** (07/2009)
14. PI, *"CSR: Small: An Information Accountability Architecture for Distributed Enterprise Systems"*, NSF Trustworthy Computing, CNS-09-14312, **$450,000** (07/2009 - 06/2012)
15. co-PI (with Jason Nieh), *"TC: Small: Exploiting Software Elasticity for Automatic Software Self-Healing"*, NSF Trustworthy Computing, CNS-09-14845, **$450,000** (07/2009 - 06/2012)
16. co-PI (with Steve Bellovin and Sal Stolfo), *"Pro-actively Removing the Botnet Threat"*, Office of Naval Research (ONR), **$294,625** (04/2009 - 09/2010)
17. co-PI (with Simha Sethumadhavan and Sal Stolfo), *"SCOPS: Secure Cyber Operations and Parallelization Studies Cluster"*, Air Force Office for Scientific Research (AFOSR), **$650,000** (04/15/2009 - 04/14/2010)
18. PI (co-PIs: Sal Stolfo), *"Program Whitelisting, Vulnerability Analytics and Risk Assessment"*, Symantec (research gift), **$65,000** (12/2008)
19. co-PI (with Sal Stolfo), *"Automated Creation of Network and Content Traffic For the National Cyber Range"*, DARPA/STO, **$85,000** (01/01/2009 - 06/30/2011; part of a larger project)
20. co-PI (with Steve Bellovin, Tal Malkin, and Sal Stolfo), *"Secure Encrypted Search"*, IARPA, $648,787 (09/2008 - 02/2010)
21. PI, *"Tracking Sensitive Information Flows in Modern Enterprises"*, Intel (research gift), $64,000 (05/2008)
22. PI, *"Privacy and Search: Having it Both Ways in Web Services"*, Google (research gift), $50,000 (03/2008)
23. PI (co-PI: Sal Stolfo), *"Continuation: Safe Browsing Through Web-based Application Communities"*, Google (research gift), **$50,000** (03/2008)
24. co-PI (with Steve Bellovin, Vishal Misra, Henning Schulzrinne, Dan Rubenstein, Nick Maxemchuck), *"Zero Outage Dynamic Intrinsically Assurable Communities (ZODIAC)"*, DARPA/STO, **$835,357** (11/2007 - 05/2009; part of a larger project with Telcordia, Sparta, GMU, and the University of Pennsylvania)
25. PI, *"Travel Supplement under the US/Japan Critical Infrastructure Protection Cooperation Program"*, NSF CyberTrust, **$38,640** (09/2007 - 08/2009)
26. PI, *"PacketSpread: Practical Network Capabilities"*, NSF CyberTrust, CNS-07-14277, **$280,000** (09/2007 - 08/2010)
27. PI, *"Integrated Enterprise Security Management"*, NSF CyberTrust, CNS-07-14647,



New Bay Capital, LLC
Ex.1015-Page 2642 of 3151

**$286,486** (08/2007 - 07/2009)

28. PI, *"Safe Browsing Through Web-based Application Communities"*, NY State/Polytechnic CAT, **$25,000** (06/2007 - 06/2009)
29. PI, *"MURI: Foundational and Systems Support for Quantitative Trust Management"*, Office of Naval Research (ONR), **$750,000** (05/2007 - 04/2012; part of a larger project with the University of Pennsylvania and Georgia Institute of Technology)
30. PI (co-PIs: Jason Nieh, Sal Stolfo), *"MURI: Autonomic Recovery of Enterprise-Wide Systems After Attack or Failure with Forward Correction"*, Air Force Office of Scientific Research (AFOSR), **$1,368,000** (05/2007 - 04/2012; part of a larger project with GMU and Penn State University)
31. co-PI (with Sal Stolfo), *"Human Behavior, Insider Threat, and Awareness"*, DHS/I3P, **$616,442** (04/2007 - 03/2009)
32. PI (co-PI: Sal Stolfo), *"Safe Browsing Through Web-based Application Communities"*, Google (research gift), **$50,000** (01/2007)
33. PI (co-PI: Sal Stolfo), *"Supplement to Behavior-based Access Control and Communication in MANETs grant"*, DARPA/IPTO and NRO, **$96,627** (09/2006 - 07/2007)
34. PI, *"Secure Overlay Services"*, NY State/Polytechnic CAT, **$10,000** (09/2006 - 06/2007)
35. PI (co-PIs: Gail Kaiser, Sal Stolfo), *"Enabling Collaborative Self-healing Software Systems"*, NSF CyberTrust, CNS-06-27473, $800,000 (09/2006 - 08/2010)
36. PI (co-PI: Sal Stolfo), *"Behavior-based Access Control and Communication in MANETs"*, DARPA/IPTO, **$100,000** (07/2006 - 06/2007)
37. co-PI (with Steve Bellovin and Sal Stolfo), *"Large-Scale System Defense"*, DTO, **$535,555** (07/2006 - 12/2007)
38. PI, *"Active Decoys for Spyware"*, NY State/Polytechnic CAT, **$25,000** (06/2006 - 12/2007)
39. PI, *"Retrofitting A Flow-oriented Paradigm in Commodity Operating Systems for High-Performance Computing"*, NSF CPA, CCF-05-41093, **$378,091** (01/2006 - 12/2008)
40. co-PI (with Jason Nieh, Gail Kaiser), *"Broadening Participation in Research"*, NSF BPC, **$133,565** (09/2005 - 08/2006)
41. PI, *"Secure Overlay Services"*, NY State/Polytechnic CAT, **$12,500** (09/2005 - 06/2006)
42. co-PI (with Dan Rubenstein, Vishal Misra), *"Secure Overlay Services"*, Intel Corp. (research gift), **$75,000** (08/2005)
43. PI, *"Snakeyes"*, New York State Center for Advanced Technology, **$14,999** (07/2005 - 06/2006)
44. PI, *"Self-protecting Software"*, Columbia Science and Technology Ventures (research gift), $65,000 (06/2005 - 09/2005)
45. co-PI (with Gail Kaiser), *"Trustworthy Computing Curriculum Development"*, Microsoft Research (research gift), **$50,000** (12/2004 - 12/2005)
46. co-PI (with Jason Nieh, Gail Kaiser), *"Secure Remote Computing Services"*, NSF ITR, CNS-04-26623, $1,200,000 (09/2004 - 08/2009)
47. PI, *"Secure Overlay Services"*, NY State/Polytechnic CAT, **$12,500** (09/2004 - 06/2005)
48. co-PI (with Dan Rubenstein, Vishal Misra), *"Secure Overlay Services"*, Intel Corp. (research gift), **$90,000** (06/2004)
49. co-PI (with Dan Rubenstein, Vishal Misra), *"Secure Overlay Services"*, Intel Corp. (research gift), **$120,000** (08/2003)
50. PI (co-PIs: Dan Rubenstein, Vishal Misra), *"Secure Overlay Services"*, Cisco Corp. (research gift), **$76,000** (07/2003)
51. co-PI (with Sal Stolfo, Tal Malkin, Vishal Misra), *"Distributed Intrusion Detection Feasibility Study"*, Department of Defense, **$300,000** (03/2003 - 03/2004)



New Bay Capital, LLC
Ex.1015-Page 2643 of 3151

52. PI, *"STRONGMAN"*, DARPA/ATO, **$23,782** (09/2002 - 08/2003; part of a larger project with the University of Pennsylvania)
53. PI, *"POSSE"*, DARPA/ATO, $16,341 (09/2002 - 08/2003; part of a larger project with the University of Pennsylvania)
54. PI, *"GRIDLOCK"*, NSF Trusted Computing, CCR-TC-02-08972, **$207,000** (07/2002 - 06/2005; part of a larger project with the University of Pennsylvania and Yale University)
55. PI (co-PIs: Dan Rubenstein, Vishal Misra), *"Secure Overlay Services"*, Cisco Corp. (research gift), **$70,000** (07/2002)
56. PI (co-PIs: Dan Rubenstein, Vishal Misra), *"Secure Overlay Services"*, DARPA/ATO, **$695,000** (06/2002 - 05/2004)
57. PI, *"Code Security Analysis Kit (CoSAK)"*, DARPA/ATO, **$37,000** (07/2001 - 06/2003; part of a larger project with Drexel University)

- **Total:** $34,240,062
- **Total as PI:** $20,625,555

**Select Invited Talks**

- *"Collaborative, Adaptive Software Defense"*, invited talk, ONR Workshop on Host Computer Security, Chicago, IL, October 2010.
- *"Using Decoys to Identify Malicious Insiders"*, invited talk, Computer Science Department, National University of Singapore, Singapore, August 2010.
- *"Behavior-based Access Control in Wired and Wireless Networks"*, invited talk, 5th Ph.D. School on Security in Wireless Networking (SWING), Bertinoro, Italy, June/July 2010.
- *"MANET Security: Background and Distributed Defense"*, invited talk, 5th Ph.D. School on Security in Wireless Networking (SWING), Bertinoro, Italy, June/July 2010.
- *"Detecting Insider Attackers"*, invited talk, 5th Ph.D. School on Security in Wireless Networking (SWING), Bertinoro, Italy, June/July 2010.
- *"Self-healing and Collaborative Software Defenses"*, invited talk, 5th Ph.D. School on Security in Wireless Networking (SWING), Bertinoro, Italy, June/July 2010.
- *"Voice over IP: Risks, Threats, and Vulnerabilities"*, invited talk, 5th Ph.D. School on Security in Wireless Networking (SWING), Bertinoro, Italy, June/July 2010.
- *"Determining Device Trustworthiness in Heterogeneous Environments"*, invited talk, Intel Workshop on Trust Evidence and End-to-end Trust in Heterogeneous Environments, Santa Clara, CA, May 2010.
- *"Moving Code: Instruction Set Randomization"*, invited talk, IC Technical Exchange on Moving Target, Washington, DC, April 2010.
- *"Voice over IP: Risks, Threats and Vulnerabilities"*, invited talk, AT&T Labs Research, Florham Park, NJ, April 2010.
- *"Voice over IP: Risks, Threats and Vulnerabilities"*, keynote talk, 5th International Conference on Information Systems Security (ICISS), Kolkata, India, December 2009.
- *"Voice over IP: Risks, Threats and Vulnerabilities"*, Cyber Infrastructure Protection (CIP) Conference, New York, June 2009.
- *"Voice over IP: Risks, Threats and Vulnerabilities"*, keynote talk, Applied Cryptography and Network Security (ACNS) Conference, Paris, France, June 2009.
- *"Automatic Software Self-Healing: Present and Future"*, keynote talk, European Workshop on Systems Security (EuroSec), Nuremberg, Germany, March 2009.
- *"VAMPIRE Project Overview"*, Symantec Research Labs, Culver City, CA, March 2009.



New Bay Capital, LLC
Ex.1015-Page 2644 of 3151

- *"Survey of IMS/VoIP Security Work"*, Agence Nationale de Reserche (ANR), Paris, France, February 2009.
- *"Simulating a Global Passive Adversary for Attacking Tor-like Anonymity Systems"*, National Institute for Advanced Industrial Science and Technology (AIST), Japan, November 2008.
- *"Denial of Service Attacks and Resilient Overlay Networks"*, ENISA-FORTH Summer School on Network & Information Security, Heraklion, Greece, September 2008.
- *"von Neumann and the Current Computer Security Landscape"*, Onassis Foundation Lectures in Science, Heraklion, Greece, July 2008.
- *"Simulating a Global Passive Adversary for Attacking Tor-like Anonymity Systems"*, Institute of Computer Science/FORTH, Heraklion, Greece, July 2008.
- *"Race to the bottom: Malicious Hardware"*, 1st FORWARD Invitational Workshop for Identifying Emerging Threats in Information and Communication Technology Infrastructures, Goteborg, Sweden, April 2008.

**Publications**

(Student co-authors are underlined.)

**Patents**

1. *"Microbilling using a trust management system"*
   Matthew A. Blaze, John Ioannidis, and Angelos D. Keromytis. U.S. Patent Number 7,996,325. Issued on August 9th 2011.
2. *"Methods, systems and media for software self-healing"*
   Michael E. Locasto, Angelos D. Keromytis, Salvatore J. Stolfo, Angelos Stavrou, Gabriela Cretu, Stylianos Sidiroglou, Jason Nieh, and Oren Laadan. U.S. Patent Number 7,962,798. Issued on June 14th, 2011.
3. *"Systems and methods for detecting and inhibiting attacks using honeypots"*
   Stylianos Sidiroglou, Angelos D. Keromytis, and Kostas G. Anagnostakis. U.S. Patent Number 7,904,959. Issued on March 8th, 2011.
4. *"Systems and methods for correlating and distributing intrusion alert information among collaborating computer systems"*
   Salvatore J. Stolfo, Angelos D. Keromytis, Vishal Misra, Michael Locasto, and Janak Parekh. U.S. Patent Number 7,784,097. Issued on August 24th, 2010.
5. *"Systems and methods for correlating and distributing intrusion alert information among collaborating computer systems"*
   Salvatore J. Stolfo, Tal Malkin, Angelos D. Keromytis, Vishal Misra, Michael Locasto, and Janak Parekh. U.S. Patent Number 7,779,463. Issued on August 17th, 2010.
6. *"Systems and methods for computing data transmission characteristics of a network path based on single-ended measurements"*
   Angelos D. Keromytis, Sambuddho Chakravarty, and Angelos Stavrou. U.S. Patent Number 7,660,261. Issued on February 9th, 2010.
7. *"Microbilling using a trust management system"*
   Matthew A. Blaze, John Ioannidis, and Angelos D. Keromytis. U.S. Patent Number 7,650,313. Issued on January 19th 2010.
8. *"Methods and systems for repairing applications"*
   Angelos D. Keromytis, Michael E. Locasto, and Stylianos Sidiroglou. U.S. Patent Number 7,490,268. Issued on February 10th 2009.
9. *"System and method for microbilling using a trust management system"*



New Bay Capital, LLC
Ex.1015-Page 2645 of 3151

Matthew A. Blaze, John Ioannidis, and Angelos D. Keromytis. U.S. Patent Number 6,789,068. Issued on September 7th 2004.

10. *"Secure and reliable bootstrap architecture"*
William A. Arbaugh, David J. Farber, Angelos D. Keromytis, and Jonathan M. Smith. U.S. Patent Number 6,185,678. Issued on February 6th 2001.

**Journal Publications**

1. *"A Comprehensive Survey of Voice over IP Security Research"*
Angelos D. Keromytis. To appear in the *IEEE Communications Surveys and Tutorials.*
2. *"A System for Generating and Injecting Indistinguishable Network Decoys"*
Brian M. Bowen, Vasileios P. Kemerlis, Pratap Prabhu, Angelos D. Keromytis, and Salvatore J. Stolfo. To appear in the *Journal of Computer Security (JCS).*
3. *"The Efficient Dual Receiver Cryptosystem and Its Applications"*
Ted Diament, Homin K. Lee, Angelos D. Keromytis, and Moti Yung. In *International Journal of Network Security (IJNS),* vol 13, no. 3, pp. 135 - 151, November 2011.
4. *"On the Infeasibility of Modeling Polymorphic Shellcode: Re-thinking the Role of Learning in Intrusion Detection Systems"*
Yingbo Song, Michael E. Locasto, Angelos Stavrou, Angelos D. Keromytis, and Salvatore J. Stolfo. In *Machine Learning Journal (MLJ),* vol. 81, no. 2, pp. 179 - 205, November 2010.
5. *"On The General Applicability of Instruction-Set Randomization"*
Stephen W. Boyd, Gaurav S. Kc, Michael E. Locasto, Angelos D. Keromytis, and Vassilis Prevelakis. In *IEEE Transactions on Dependable and Secure Computing (TDSC),* vol. 7, no. 3, pp. 255 - 270, July - September 2010.
6. *"Shadow Honeypots"*
Michalis Polychronakis, Periklis Akritidis, Stelios Sidiroglou, Kostas G. Anagnostakis, Angelos D. Keromytis, and Evangelos Markatos. In *International Journal of Computer and Network Security (IJCNS),* vol. 2, no. 9, pp. 1 - 15, September 2010.
7. *"Ethics in Security Vulnerability Research"*
Andrea M. Matwyshyn, Ang Cui, Salvatore J. Stolfo, and Angelos D. Keromytis. In *IEEE Security & Privacy Magazine,* vol. 8, no. 2, pp. 67 - 72, March/April 2010.
8. *"Voice over IP Security: Research and Practice"*
Angelos D. Keromytis. In *IEEE Security & Privacy Magazine,* vol. 8, no. 2, pp. 76 - 78, March/April 2010.
9. *"A Market-based Bandwidth Charging Framework"*
David Michael Turner, Vassilis Prevelakis, and Angelos D. Keromytis. In *ACM Transactions on Internet Technology (ToIT),* vol. 10, no. 1, pp. 1 - 30, February 2010.
10. *"A Look at VoIP Vulnerabilities"*
Angelos D. Keromytis. In *USENIX ;login: Magazine,* vol. 35, no. 1, pp. 41 - 50, February 2010.
11. *"Designing Host and Network Sensors to Mitigate the Insider Threat"*
Brian M. Bowen, Malek Ben Salem, Shlomo Hershkop, Angelos D. Keromytis, and Salvatore J. Stolfo. In *IEEE Security & Privacy Magazine,* vol. 7, no. 6, pp. 22 - 29, November/December 2009.
12. *"Elastic Block Ciphers: Method, Security and Instantiations"*
Debra L. Cook, Moti Yung, and Angelos D. Keromytis. In *Springer International Journal of Information Security (IJIS),* vol 8, no. 3, pp. 211 - 231, June 2009.
13. *"On the Deployment of Dynamic Taint Analysis for Application Communities"*
Hyung Chan Kim and Angelos D. Keromytis. In *IEICE Transactions,* vol. E92-D, no. 3, pp.



New Bay Capital, LLC
Ex.1015-Page 2646 of 3151

548 - 551, March 2009.

14. *"Dynamic Trust Management"*
    Matt Blaze, Sampath Kannan, Insup Lee, Oleg Sokolsky, Jonathan M. Smith, Angelos D. Keromytis, and Wenke Lee. In *IEEE Computer Magazine,* vol. 42, no. 2, pp. 44 - 52, February 2009.
15. *"Randomized Instruction Sets and Runtime Environments: Past Research and Future Directions"*
    Angelos D. Keromytis. In *IEEE Security & Privacy Magazine,* vol. 7, no. 1, pp. 18 - 25, January/February 2009.
16. *"Anonymity in Wireless Broadcast Networks"*
    Matt Blaze, John Ioannidis, Angelos D. Keromytis, Tal Malkin, and Avi Rubin. In *International Journal of Network Security (IJNS),* vol. 8, no. 1, pp. 37 - 51, January 2009.
17. *"Decentralized Access Control in Networked File Systems"*
    Stefan Miltchev, Jonathan M. Smith, Vassilis Prevelakis, Angelos D. Keromytis, and Sotiris Ioannidis. In *ACM Computing Surveys,* vol. 40, no. 3, pp. 10:1 - 10:30, August 2008.
18. *"Robust Reactions to Potential Day-Zero Worms through Cooperation and Validation"*
    Kostas G. Anagnostakis, Michael Greenwald, Sotiris Ioannidis, and Angelos D. Keromytis. In *Springer International Journal of Information Security (IJIS), ISC 2006 Special Issue,* vol.6, no. 6, pp. 361 - 378, October 2007. (Extended version of the ISC 2006 paper.)
19. *"Requirements for Scalable Access Control and Security Management Architectures"*
    Angelos D. Keromytis and Jonathan M. Smith. In *ACM Transactions on Internet Technology (ToIT),* vol. 7, no. 2, pp. 1 - 22, May 2007.
20. *"Virtual Private Services: Coordinated Policy Enforcement for Distributed Applications"*
    Sotiris Ioannidis, Steven M. Bellovin, John Ioannidis, Angelos D. Keromytis, Kostas G. Anagnostakis, and Jonathan M. Smith. In *International Journal of Network Security (IJNS),* vol. 4, no. 1, pp. 69 - 80, January 2007.
21. *"Countering DDoS Attacks with Multi-path Overlay Networks"*
    Angelos Stavrou and Angelos D. Keromytis. In *Information Assurance Technology Analysis Center (IATAC) Information Assurance Newsletter (IAnewsletter),* vol. 9, no. 3, pp. 26 - 30, Winter 2006. *(Invited paper, based on the CCS 2005 paper.)*
22. *"Conversion Functions for Symmetric Key Ciphers"*
    Debra L. Cook and Angelos D. Keromytis. In *Journal of Information Assurance and Security (JIAS),* vol. 1, no. 2, pp. 119 - 128, June 2006. *(Extended version of the IAS 2005 paper.)*
23. *"Execution Transactions for Defending Against Software Failures: Use and Evaluation"*
    Stelios Sidiroglou and Angelos D. Keromytis. In *Springer International Journal of Information Security (IJIS),* vol. 5, no. 2, pp. 77 - 91, April 2006. *(Extended version of the ISC 2005 paper.)*
24. *"Worm Propagation Strategies in an IPv6 Internet"*
    Steven M. Bellovin, Bill Cheswick, and Angelos D. Keromytis. In *USENIX ;login,* vol. 31, no. 1, pp. 70 - 76, February 2006.
25. *"Cryptography As An Operating System Service: A Case Study"*
    Angelos D. Keromytis, Theo de Raadt, Jason Wright, and Matthew Burnside. In *ACM Transactions on Computer Systems (ToCS),* vol. 24, no. 1, pp. 1 - 38, February 2006. *(Extended version of USENIX Technical 2003 paper.)*
26. *"Countering Network Worms Through Automatic Patch Generation"*
    Stelios Sidiroglou and Angelos D. Keromytis. In *IEEE Security & Privacy,* vol. 3, no. 6, pp. 41 - 49, November/December 2005.
27. *"WebSOS: An Overlay-based System For Protecting Web Servers From Denial of Service*



New Bay Capital, LLC
Ex.1015-Page 2647 of 3151

*Attacks"*
Angelos Stavrou, Debra L. Cook, William G. Morein, Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In *Elsevier Journal of Computer Networks, special issue on Web and Network Security,* vol. 48, no. 5, pp. 781 - 807, August 2005. *(Extended version of the CCS 2003 paper.)*

28. *"Hardware Support For Self-Healing Software Services"*
Stelios Sidiroglou, Michael E. Locasto, and Angelos D. Keromytis. In *ACM SIGARCH Computer Architecture News, Special Issue on Workshop on Architectural Support for Security and Anti-Virus (WASSA),* vol. 33, no. 1, pp. 42 - 47, March 2005. Also appeared in the Proceedings of the *Workshop on Architectural Support for Security and Anti-Virus (WASSA),* held in conjunction with the *11th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XI),* pp. 37 - 43. October 2004, Boston, MA.
29. *"The Case For Crypto Protocol Awareness Inside The OS Kernel"*
Matthew Burnside and Angelos D. Keromytis. In *ACM SIGARCH Computer Architecture News, Special Issue on Workshop on Architectural Support for Security and Anti-Virus (WASSA),* vol. 33, no. 1, pp. 58 - 64, March 2005. Also appeared in the Proceedings of the *Workshop on Architectural Support for Security and Anti-Virus (WASSA),* held in conjunction with the *11th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XI),* pp. 54 - 60. October 2004, Boston, MA.
30. *"Patch-on-Demand Saves Even More Time?"*
Angelos D. Keromytis. In *IEEE Computer,* vol. 37, no. 8, pp. 94 - 96, August 2004.
31. *"Just Fast Keying: Key Agreement In A Hostile Internet"*
William Aiello, Steven M. Bellovin, Matt Blaze, Ran Canetti, John Ioannidis, Angelos D. Keromytis, and Omer Reingold. In *ACM Transactions on Information and System Security (TISSEC),* vol. 7, no. 2, pp. 1 - 32, May 2004. *(Extended version of the CCS 2002 paper.)*
32. *"SOS: An Architecture for Mitigating DDoS Attacks"*
Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In *IEEE Journal on Selected Areas in Communications (JSAC), special issue on Recent Advances in Service Overlay Networks,* vol. 22, no. 1, pp. 176 - 188, January 2004. *(Extended version of the SIGCOMM 2002 paper.)*
33. *"A Secure PLAN"*
Michael Hicks, Angelos D. Keromytis, and Jonathan M. Smith. In *IEEE Transactions on Systems, Man, and Cybernetics (T-SMC) Part C: Applications and Reviews, Special issue on technologies promoting computational intelligence, openness and programmability in networks and Internet services: Part I,* vol. 33, no. 3, pp. 413 - 426, August 2003. *(Extended version of the DANCE 2002 paper.)*
34. *"Drop-in Security for Distributed and Portable Computing Elements"*
Vassilis Prevelakis and Angelos D. Keromytis. In MCB Press *Emerald Journal of Internet Research: Electronic Networking, Applications and Policy,* vol. 13, no. 2, pp. 107 - 115, 2003. *(Extended version of the INC 2002 paper.)*
35. *"Trust Management for IPsec"*
Matt Blaze, John Ioannidis, and Angelos D. Keromytis. In *ACM Transactions on Information and System Security (TISSEC),* vol. 5, no. 2, pp. 1 - 24, May 2002. *(Extended version of the NDSS 2001 paper.)*
36. *"The Price of Safety in an Active Network"*
D. Scott Alexander, Paul B. Menage, Angelos D. Keromytis, William A. Arbaugh, Kostas G.



New Bay Capital, LLC
Ex.1015-Page 2648 of 3151

Anagnostakis, and Jonathan M. Smith. In *Journal of Communications and Networks (JCN), special issue on programmable switches and routers,* vol. 3, no. 1, pp. 4 - 18, March 2001. Older versions are available as *University of Pennsylvania Technical Report MS-CIS-99-04* and *University of Pennsylvania Technical Report MS-CIS-98-02.*

37. *"Secure Quality of Service Handling (SQoSH)"*
D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, Steve Muir, and Jonathan M. Smith. In *IEEE Communications Magazine,* vol. 38, no. 4, pp. 106 - 112, April 2000. An older version is available as *University of Pennsylvania Technical Report MS-CIS-99-05.*
38. *"Safety and Security of Programmable Network Infrastructures"*
D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith. In *IEEE Communications Magazine, issue on Programmable Networks,* vol. 36, no. 10, pp. 84 - 92, October 1998.
39. *"A Secure Active Network Environment Architecture"*
D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith. In *IEEE Network Magazine, special issue on Active and Controllable Networks,* vol. 12, no. 3, pp. 37 - 45, May/June 1998.
40. *"The SwitchWare Active Network Architecture"*
D. Scott Alexander, William A. Arbaugh, Michael Hicks, Pankaj Kakkar, Angelos D. Keromytis, Jonathan T. Moore, Carl A. Gunter, Scott M. Nettles, and Jonathan M. Smith. In *IEEE Network Magazine, special issue on Active and Programmable Networks,* vol. 12, no. 3, pp. 29 - 36, May/June 1998.

**Peer-Reviewed Conference Proceedings**

1. *"A Multilayer Overlay Network Architecture for Enhancing IP Services Availability Against DoS"*
Dimitris Geneiatakis, Georgios Portokalidis, and Angelos D. Keromytis. To appear in the Proceedings of the *7th International Conference on Information Systems Security (ICISS).* December 2011, Kolkata, India. *(Acceptance rate: 22.8%)*
2. *"ROP Payload Detection Using Speculative Code Execution"*
Michalis Polychronakis and Angelos D. Keromytis. To appear in the Proceedings of the *6th International Conference on Malicious and Unwanted Software (MALWARE).* October 2011, Fajardo, PR.
3. *"Detecting Traffic Snooping in Tor Using Decoys"*
Sambuddho Chakravarty, Georgios Portokalidis, Michalis Polychronakis, and Angelos D. Keromytis. To appear in Proceedings of the *14th International Symposium on Recent Advances in Intrusion Detection (RAID).* September 2011, Menlo Park, CA. *(Acceptance rate: 23%)*
4. *"Measuring the Deployment Hiccups of DNSSEC"*
Vasilis Pappas and Angelos D. Keromytis. In Proceedings of the *International Conference on Advances in Computing and Communications (ACC), Part III,* pp. 44 - 54. July 2011, Kochi, India. *(Acceptance rate: 39%)*
5. *"Misuse Detection in Consent-based Networks"*
Mansoor Alicherry and Angelos D. Keromytis. In Proceedings of the *9th International Conference on Applied Cryptography and Network Security (ACNS),* pp. 38 - 56. June 2011, Malaga, Spain. *(Acceptance rate: 18%)*
6. *"Retrofitting Security in COTS Software with Binary Rewriting"*
Padraig O'Sullivan, Kapil Anand, Aparna Kothan, Matthew Smithson, Rajeev Barua, and Angelos D. Keromytis. In Proceedings of the *26th IFIP International Information Security*



New Bay Capital, LLC
Ex.1015-Page 2649 of 3151

*Conference (SEC),* pp. 154 - 172. June 2011, Lucerne, Switzerland. *(Acceptance rate: 24%)*

7. *"Fast and Practical Instruction-Set Randomization for Commodity Systems"*
   Georgios Portokalidis and Angelos D. Keromytis. In Proceedings of the *26th Annual Computer Security Applications Conference (ACSAC),* pp. 41 - 48. December 2010, Austin, TX. *(Acceptance rate: 17%)*
8. *"An Adversarial Evaluation of Network Signaling and Control Mechanisms"*
   Kangkook Jee, Stelios Sidiroglou-Douskos, Angelos Stavrou, and Angelos D. Keromytis. In Proceedings of the *13th International Conference on Information Security and Cryptology (ICISC).* December 2010, Seoul, Korea.
9. *"Evaluation of a Spyware Detection System using Thin Client Computing"*
   Vasilis Pappas, Brian M. Bowen, and Angelos D. Keromytis. In Proceedings of the *13th International Conference on Information Security and Cryptology (ICISC),* pp. 222 - 232. December 2010, Seoul, Korea.
10. *"Crimeware Swindling without Virtual Machines"*
    Vasilis Pappas, Brian M. Bowen, and Angelos D. Keromytis. In Proceedings of the *13th Information Security Conference (ISC),* pp. 196 - 202. October 2010, Boca Raton, FL. *(Acceptance rate: 27.6%)*
11. *"iLeak: A Lightweight System for Detecting Inadvertent Information Leaks"*
    Vasileios P. Kemerlis, Vasilis Pappas, Georgios Portokalidis, and Angelos D. Keromytis. In Proceedings of the *6th European Conference on Computer Network Defense (EC2ND),* pp. 21 - 28. October 2010, Berlin, Germany.
12. *"Traffic Analysis Against Low-Latency Anonymity Networks Using Available Bandwidth Estimation"*
    Sambuddho Chakravarty, Angelos Stavrou, and Angelos D. Keromytis. In Proceedings of the *15th European Symposium on Research in Computer Security (ESORICS),* pp. 249 - 267. September 2010, Athens, Greece. *(Acceptance rate: 20%)*
13. *"BotSwindler: Tamper Resistant Injection of Believable Decoys in VM-Based Hosts for Crimeware Detection"*
    Brian M. Bowen, Pratap Prabhu, Vasileios P. Kemerlis, Stelios Sidiroglou, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *13th International Symposium on Recent Advances in Intrusion Detection (RAID),* pp. 118 - 137. September 2010, Ottawa, Canada. *(Acceptance rate: 23.5%)*
14. *"An Analysis of Rogue AV Campaigns"*
    Marco Cova, Corrado Leita, Olivier Thonnard, Angelos D. Keromytis, and Marc Dacier. In Proceedings of the *13th International Symposium on Recent Advances in Intrusion Detection (RAID),* pp. 442 - 463. September 2010, Ottawa, Canada. *(Acceptance rate: 23.5%)*
15. *"DIPLOMA: Distributed Policy Enforcement Architecture for MANETs"*
    Mansoor Alicherry and Angelos D. Keromytis. In Proceedings of the *4th International Conference on Network and System Security (NSS),* pp. 89 - 98. September 2010, Melbourne, Australia. *(Acceptance rate: 26%)*
16. *"Automating the Injection of Believable Decoys to Detect Snooping" (Short Paper)*
    Brian M. Bowen, Vasileios Kemerlis, Pratap Prabhu, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *3rd ACM Conference on Wireless Network Security (WiSec),* pp. 81 - 86. March 2010, Hoboken, NJ. *(Acceptance rate: 21%)*
17. *"BARTER: Behavior Profile Exchange for Behavior-Based Admission and Access Control in MANETs"*
    Vanessa Frias-Martinez, Salvatore J. Stolfo, and Angelos D. Keromytis. In Proceedings of the *5th International Conference on Information Systems Security (ICISS),* pp. 193 - 207.



New Bay Capital, LLC
Ex.1015-Page 2650 of 3151

December 2009, Kolkata, India. *(Acceptance rate: 19.8%)*

18. *"A Survey of Voice Over IP Security Research"*
Angelos D. Keromytis. In Proceedings of the *5th International Conference on Information Systems Security (ICISS)*, pp. 1 - 17. December 2009, Kolkata, India. *(Invited paper)*
19. *"A Network Access Control Mechanism Based on Behavior Profiles"*
Vanessa Frias-Martinez, Joseph Sherrick, Salvatore J. Stolfo, and Angelos D. Keromytis. In Proceedings of the *25th Annual Computer Security Applications Conference (ACSAC)*, pp. 3 - 12. December 2009, Honolulu, HI. *(Acceptance rate: 20%)*
20. *"Gone Rogue: An Analysis of Rogue Security Software Campaigns"*
Marco Cova, Corrado Leita, Olivier Thonnard, Angelos D. Keromytis, and Marc Dacier. In Proceedings of the *5th European Conference on Computer Network Defense (EC2ND)*, pp. 1 - 3. November 2009, Milan, Italy. *(Invited paper)*
21. *"Baiting Inside Attackers Using Decoy Documents"*
Brian M. Bowen, Shlomo Hershkop, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *5th International ICST Conference on Security and Privacy in Communication Networks (SecureComm)*, pp. 51 - 70. September 2009, Athens, Greece. *(Acceptance rate: 25.3%)*
22. *"Deny-by-Default Distributed Security Policy Enforcement in Mobile Ad Hoc Networks (Short Paper)"*
Mansoor Alicherry, Angelos D. Keromytis, and Angelos Stavrou. In Proceedings of the *5th International ICST Conference on Security and Privacy in Communication Networks (SecureComm)*, pp. 41 - 50. September 2009, Athens, Greece. *(Acceptance rate: 34.7%)*
23. *"Adding Trust to P2P Distribution of Paid Content"*
Alex Sherman, Angelos Stavrou, Jason Nieh, Angelos D. Keromytis, and Clifford Stein. In Proceedings of the *12th Information Security Conference (ISC)*, pp. 459 - 474. September 2009, Pisa, Italy. *(Acceptance rate: 27.6%)*
24. *"A2M: Access-Assured Mobile Desktop Computing"*
Angelos Stavrou, Ricardo A. Baratto, Angelos D. Keromytis, and Jason Nieh. In Proceedings of the *12th Information Security Conference (ISC)*, pp. 186 - 201. September 2009, Pisa, Italy. *(Acceptance rate: 27.6%)*
25. *"F3ildCrypt: End-to-End Protection of Sensitive Information in Web Services"*
Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *12th Information Security Conference (ISC)*, pp. 491 - 506. September 2009, Pisa, Italy. *(Acceptance rate: 27.6%)*
26. *"DoubleCheck: Multi-path Verification Against Man-in-the-Middle Attacks"*
Mansoor Alicherry and Angelos D. Keromytis. In Proceedings of the *IEEE Symposium on Computers and Communications (ISCC)*, pp. 557 - 563. July 2009, Sousse, Tunisia. *(Acceptance rate: 36%)*
27. *"Voice over IP: Risks, Threats and Vulnerabilities"*
Angelos D. Keromytis. In Proceedings (electronic) of the *Cyber Infrastructure Protection (CIP) Conference*. June 2009, New York, NY. *(Invited paper)*
28. *"Capturing Information Flow with Concatenated Dynamic Taint Analysis"*
Hyung Chan Kim, Angelos D. Keromytis, Michael Covington, and Ravi Sahita. In Proceedings of the *4th International Conference on Availability, Reliability and Security (ARES)*, pp. 355 - 362. March 2009, Fukuoka, Japan. *(Acceptance rate: 25%)*
29. *"ASSURE: Automatic Software Self-healing Using REscue points"*
Stelios Sidiroglou, Oren Laadan, Nico Viennot, Carlos-René Pérez, Angelos D. Keromytis, and Jason Nieh. In Proceedings of the *14th International Conference on Architectural Support*



New Bay Capital, LLC
Ex.1015-Page 2651 of 3151

*for Programming Languages and Operating Systems (ASPLOS),* pp. 37 - 48. March 2009, Washington, DC. *(Acceptance rate: 25.6%)*

30. *"Spectrogram: A Mixture-of-Markov-Chains Model for Anomaly Detection in Web Traffic"* Yingbo Song, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *16th Internet Society (ISOC) Symposium on Network and Distributed Systems Security (SNDSS),* pp. 121 - 135. February 2009, San Diego, CA. *(Acceptance rate: 11.7%)*
31. *"Constructing Variable-Length PRPs and SPRPs from Fixed-Length PRPs"* Debra L. Cook, Moti Yung, and Angelos D. Keromytis. In Proceedings of the *4th International Conference on Information Security and Cryptology (Inscrypt),* pp. 157 - 180. December 2008, Beijing, China. *(Acceptance rate: 17.5%)*
32. *"Behavior-Profile Clustering for False Alert Reduction in Anomaly Detection Sensors"* Vanessa Frias-Martinez, Salvatore J. Stolfo, and Angelos D. Keromytis. In Proceedings of the *24th Annual Computer Security Applications Conference (ACSAC),* pp. 367 - 376. December 2008, Anaheim, CA. *(Acceptance rate: 24.2%)*
33. *"Authentication on Untrusted Remote Hosts with Public-key Sudo"* Matthew Burnside, Mack Lu, and Angelos D. Keromytis. In Proceedings of the *22nd USENIX Large Installation Systems Administration (LISA) Conference,* pp. 103 - 107. November 2008, San Diego, CA.
34. *"Behavior-Based Network Access Control: A Proof-of-Concept"* Vanessa Frias-Martinez, Salvatore J. Stolfo, and Angelos D. Keromytis. In Proceedings of the *11th Information Security Conference (ISC),* pp. 175 - 190. Taipei, Taiwan, September 2008. *(Acceptance rate: 23.9%)*
35. *"Path-based Access Control for Enterprise Networks"* Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *11th Information Security Conference (ISC),* pp. 191 - 203. Taipei, Taiwan, September 2008. *(Acceptance rate: 23.9%)*
36. *"Methods for Linear and Differential Cryptanalysis of Elastic Block Ciphers"* Debra L. Cook, Moti Yung, and Angelos D. Keromytis. In Proceedings of the *13th Australasian Conference on Information Security and Privacy (ACISP),* pp. 187 - 202. July 2008, Wollongong, Australia.*(Acceptance rate: 29.7%)*
37. *"Pushback for Overlay Networks: Protecting against Malicious Insiders"* Angelos Stavrou, Michael E. Locasto, and Angelos D. Keromytis. In Proceedings of the *6th International Conference on Applied Cryptography and Network Security (ACNS),* pp 39 - 54. June 2008, New York, NY. *(Acceptance rate: 22.9%)*
38. *"Casting out Demons: Sanitizing Training Data for Anomaly Sensors"* Gabriela F. Cretu, Angelos Stavrou, Michael E. Locasto, Salvatore J. Stolfo, and Angelos D. Keromytis. In Proceedings of the *IEEE Symposium on Security & Privacy,* pp. 81 - 95. May 2008, Oakland, CA. *(Acceptance rate: 11.2%)*
39. *"Taming the Devil: Techniques for Evaluating Anonymized Network Data"* Scott E. Coull, Charles V. Wright, Angelos D. Keromytis, Fabian Monrose, and Michael K. Reiter. In Proceedings of the *15th Internet Society (ISOC) Symposium on Network and Distributed Systems Security (SNDSS),* pp. 125 - 135. February 2008, San Diego, CA. *(Acceptance rate: 17.8%)*
40. *"SSARES: Secure Searchable Automated Remote Email Storage"* Adam J. Aviv, Michael E. Locasto, Shaya Potter, and Angelos D. Keromytis. In Proceedings of the *23rd Annual Computer Security Applications Conference (ACSAC),* pp. 129 - 138. December 2007, Miami Beach, FL. *(Acceptance rate: 22%)*
41. *"On the Infeasibility of Modeling Polymorphic Shellcode"*



New Bay Capital, LLC
Ex.1015-Page 2652 of 3151

Yingbo Song, Michael E. Locasto, Angelos Stavrou, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *13th ACM Conference on Computer and Communications Security (CCS),* pp. 541 - 551. October/November 2007, Alexandria, VA. *(Acceptance rate: 18.1%)*

42. *"Defending Against Next Generation Attacks Through Network/Endpoint Collaboration and Interaction"*
Spiros Antonatos, Michael E. Locasto, Stelios Sidiroglou, Angelos D. Keromytis, and Evangelos Markatos. In Proceedings of the *3rd European Conference on Computer Network Defense (EC2ND).* October 2007, Heraclion, Greece. *(Invited paper)*
43. *"Elastic Block Ciphers in Practice: Constructions and Modes of Encryption"*
Debra L. Cook, Moti Yung, and Angelos D. Keromytis. In Proceedings of the *3rd European Conference on Computer Network Defense (EC2ND).* October 2007, Heraclion, Greece.
44. *"The Security of Elastic Block Ciphers Against Key-Recovery Attacks"*
Debra L. Cook, Moti Yung, and Angelos D. Keromytis. In Proceedings of the *10th Information Security Conference (ISC),* pp. 89 - 103. Valparaiso, Chile, October 2007. *(Acceptance rate: 25%)*
45. *"Characterizing Self-healing Software Systems"*
Angelos D. Keromytis. In Proceedings of the *4th International Conference on Mathematical Methods, Models and Architectures for Computer Networks Security (MMM-ACNS),* pp. 22 - 33. September 2007, St. Petersburg, Russia. *(Invited paper)*
46. *"A Study of Malcode-Bearing Documents"*
Wei-Jen Li, Salvatore J. Stolfo, Angelos Stavrou, Elli Androulaki, and Angelos D. Keromytis. In Proceedings of the *4th GI International Conference on Detection of Intrusions & Malware, and Vulnerability Assessment (DIMVA),* pp. 231 - 250. July 2007, Lucerne, Switzerland. *(Acceptance rate: 21%)*
47. *"From STEM to SEAD: Speculative Execution for Automated Defense"*
Michael E. Locasto, Angelos Stavrou, Gabriela F. Cretu, and Angelos D. Keromytis. In Proceedings of the *USENIX Annual Technical Conference,* pp. 219 - 232. June 2007, Santa Clara, CA. *(Acceptance rate: 18.75%)*
48. *"Using Rescue Points to Navigate Software Recovery (Short Paper)"*
Stelios Sidiroglou, Oren Laadan, Angelos D. Keromytis, and Jason Nieh. In Proceedings of the *IEEE Symposium on Security & Privacy,* pp. 273 - 278. May 2007, Oakland, CA. *(Acceptance rate: 8.3%)*
49. *"Mediated Overlay Services (MOSES): Network Security as a Composable Service"*
Stelios Sidiroglou, Angelos Stavrou, and Angelos D. Keromytis. In Proceedings of the *IEEE Sarnoff Symposium.* May 2007, Princeton, NJ. *(Invited paper)*
50. *"Elastic Block Ciphers: The Basic Design"*
Debra L. Cook, Moti Yung, and Angelos D. Keromytis. In Proceedings of the *2nd ACM Symposium on InformAtion, Computer and Communications Security (ASIACCS),* pp. 350 - 355. March 2007, Singapore.
51. *"Robust Reactions to Potential Day-Zero Worms through Cooperation and Validation"*
Kostas G. Anagnostakis, Michael B. Greenwald, Sotiris Ioannidis, and Angelos D. Keromytis. In Proceedings of the *9th Information Security Conference (ISC),* pp. 427 - 442. August/September 2006, Samos, Greece. *(Acceptance rate: 20.2%)*
52. *"Low Latency Anonymity with Mix Rings"*
Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *9th Information Security Conference (ISC),* pp. 32 - 45. August/September 2006, Samos, Greece. *(Acceptance rate: 20.2%)*



New Bay Capital, LLC
Ex.1015-Page 2653 of 3151

53. *"W3Bcrypt: Encryption as a Stylesheet"*
Angelos Stavrou, Michael E. Locasto, and Angelos D. Keromytis. In Proceedings of the *4th International Conference on Applied Cryptography and Network Security (ACNS),* pp. 349 - 364. June 2006, Singapore.
54. *"Software Self-Healing Using Collaborative Application Communities"*
Michael E. Locasto, Stelios Sidiroglou, and Angelos D. Keromytis. In Proceedings of the *13th Internet Society (ISOC) Symposium on Network and Distributed Systems Security (SNDSS),* pp. 95 - 106. February 2006, San Diego, CA. *(Acceptance rate: 13.6%)*
55. *"Remotely Keyed Cryptographics: Secure Remote Display Access Using (Mostly) Untrusted Hardware"*
Debra L. Cook, Ricardo A. Baratto, and Angelos D. Keromytis. In Proceedings of the *7th International Conference on Information and Communications Security (ICICS),* pp. 363 - 375. December 2005, Beijing, China. *(Acceptance rate: 17.4%)*
56. *"e-NeXSh: Achieving an Effectively Non-Executable Stack and Heap via System-Call Policing"*
Gaurav S. Kc and Angelos D. Keromytis. In Proceedings of the *21st Annual Computer Security Applications Conference (ACSAC),* pp. 259 - 273. December 2005, Tucson, AZ. *(Acceptance rate: 19.6%)*
57. *"Action Amplification: A New Approach To Scalable Administration"*
Kostas G. Anagnostakis and Angelos D. Keromytis. In Proceedings of the *13th IEEE International Conference on Networks (ICON),* vol. 2, pp. 862 - 867. November 2005, Kuala Lumpur, Malaysia.
58. *"A Repeater Encryption Unit for IPv4 and IPv6"*
Norimitsu Nagashima and Angelos D. Keromytis. In Proceedings of the *13th IEEE International Conference on Networks (ICON),* vol. 1, pp. 335 - 340. November 2005, Kuala Lumpur, Malaysia.
59. *"Countering DoS Attacks With Stateless Multipath Overlays"*
Angelos Stavrou and Angelos D. Keromytis. In Proceedings of the *12th ACM Conference on Computer and Communications Security (CCS),* pp. 249 - 259. November 2005, Alexandria, VA. *(Acceptance rate: 15.2%)*
60. *"A Dynamic Mechanism for Recovering from Buffer Overflow Attacks"*
Stelios Sidiroglou, Giannis Giovanidis, and Angelos D. Keromytis. In Proceedings of the *8th Information Security Conference (ISC),* pp. 1 - 15. September 2005, Singapore. *(Acceptance rate: 14%)*
61. *"gore: Routing-Assisted Defense Against DDoS Attacks"*
Stephen T. Chou, Angelos Stavrou, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *8th Information Security Conference (ISC),* pp. 179 - 193. September 2005, Singapore. *(Acceptance rate: 14%)*
62. *"FLIPS: Hybrid Adaptive Intrusion Prevention"*
Michael E. Locasto, Ke Wang, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *8th International Symposium on Recent Advances in Intrusion Detection (RAID),* pp. 82 - 101. September 2005, Seattle, WA. *(Acceptance rate: 20.4%)*
63. *"Detecting Targeted Attacks Using Shadow Honeypots"*
Kostas G. Anagnostakis, Stelios Sidiroglou, Periklis Akritidis, Konstantinos Xinidis, Evangelos Markatos, and Angelos D. Keromytis. In Proceedings of the *14th USENIX Security Symposium,* pp. 129 - 144. August 2005, Baltimore, MD. *(Acceptance rate: 14%)*
64. *"The Bandwidth Exchange Architecture"*
David Michael Turner, Vassilis Prevelakis, and Angelos D. Keromytis. In Proceedings of the



New Bay Capital, LLC
Ex.1015-Page 2654 of 3151

*10th IEEE Symposium on Computers and Communications (ISCC),* pp. 939 - 944. June 2005, Cartagena, Spain.

65. *"An Email Worm Vaccine Architecture"*
Stelios Sidiroglou, John Ioannidis, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *1st Information Security Practice and Experience Conference (ISPEC),* pp. 97 - 108. April 2005, Singapore.
66. *"Building a Reactive Immune System for Software Services"*
Stelios Sidiroglou, Michael E. Locasto, Stephen W. Boyd, and Angelos D. Keromytis. In Proceedings of the *USENIX Annual Technical Conference,* pp. 149 - 161. April 2005, Anaheim, CA. *(Acceptance rate: 20.3%)*
67. *"Conversion and Proxy Functions for Symmetric Key Ciphers"*
Debra L. Cook and Angelos D. Keromytis. In Proceedings of the *IEEE International Conference on Information Technology: Coding and Computing (ITCC), Information and Security (IAS) Track,* pp. 662 - 667. April 2005, Las Vegas, NV.
68. *"The Effect of DNS Delays on Worm Propagation in an IPv6 Internet"*
Abhinav Kamra, Hanhua Feng, Vishal Misra, and Angelos D. Keromytis. In Proceedings of *IEEE INFOCOM,* vol. 4, pp. 2405 - 2414. March 2005, Miami, FL. *(Acceptance rate: 17%)*
69. *"MOVE: An End-to-End Solution To Network Denial of Service"*
Angelos Stavrou, Angelos D. Keromytis, Jason Nieh, Vishal Misra, and Dan Rubenstein. In Proceedings of the *12th Internet Society (ISOC) Symposium on Network and Distributed Systems Security (SNDSS),* pp. 81 - 96. February 2005, San Diego, CA. *(Acceptance rate: 12.9%)*
70. *"CryptoGraphics: Secret Key Cryptography Using Graphics Cards"*
Debra L. Cook, John Ioannidis, Angelos D. Keromytis, and Jake Luck. In Proceedings of the *RSA Conference, Cryptographer's Track (CT-RSA),* pp. 334 - 350. February 2005, San Francisco, CA.
71. *"The Dual Receiver Cryptogram and Its Applications"*
Ted Diament, Homin K. Lee, Angelos D. Keromytis, and Moti Yung. In Proceedings of the *11th ACM Conference on Computer and Communications Security (CCS),* pp. 330 - 343. October 2004, Washington, DC. *(Acceptance rate: 13.9%)*
72. *"Hydan: Hiding Information in Program Binaries"*
Rakan El-Khalil and Angelos D. Keromytis. In Proceedings of the *6th International Conference on Information and Communications Security (ICICS),* pp. 187 - 199. October 2004, Malaga, Spain. *(Acceptance rate: 16.9%)*
73. *"Recursive Sandboxes: Extending Systrace To Empower Applications"*
Aleksey Kurchuk and Angelos D. Keromytis. In Proceedings of the *19th IFIP International Information Security Conference (SEC),* pp. 473 - 487. August 2004, Toulouse, France. *(Acceptance rate: 22%)*
74. *"SQLrand: Preventing SQL Injection Attacks"*
Stephen W. Boyd and Angelos D. Keromytis. In Proceedings of the *2nd International Conference on Applied Cryptography and Network Security (ACNS),* pp. 292 - 302. June 2004, Yellow Mountain, China. *(Acceptance rate: 12.1%)*
75. *"CamouflageFS: Increasing the Effective Key Length in Cryptographic Filesystems on the Cheap"*
Michael E. Locasto and Angelos D. Keromytis. In Proceedings of the *2nd International Conference on Applied Cryptography and Network Security (ACNS),* pp. 1 - 15. June 2004, Yellow Mountain, China. *(Acceptance rate: 12.1%)*
76. *"A Pay-per-Use DoS Protection Mechanism For The Web"*



New Bay Capital, LLC
Ex.1015-Page 2655 of 3151

Angelos Stavrou, John Ioannidis, Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In Proceedings of the *2nd International Conference on Applied Cryptography and Network Security (ACNS),* pp. 120 - 134. June 2004, Yellow Mountain, China. *(Acceptance rate: 12.1%)*

77. *"Dealing with System Monocultures"*
Angelos D. Keromytis and Vassilis Prevelakis. In Proceedings (electronic) of the *NATO Information Systems Technology (IST) Panel Symposium on Adaptive Defense in Unclassified Networks.* April 2004, Toulouse, France.
78. *"Managing Access Control in Large Scale Heterogeneous Networks"*
Angelos D. Keromytis, Kostas G. Anagnostakis, Sotiris Ioannidis, Michael Greenwald, and Jonathan M. Smith. In Proceedings (electronic) of the *NATO NC3A Symposium on Interoperable Networks for Secure Communications (INSC).* November 2003, The Hague, Netherlands.
79. *"Countering Code-Injection Attacks With Instruction-Set Randomization"*
Gaurav S. Kc, Angelos D. Keromytis, and Vassilis Prevelakis. In Proceedings of the *10th ACM International Conference on Computer and Communications Security (CCS),* pp. 272 - 280. October 2003, Washington, DC. *(Acceptance rate: 13.8%)*
80. *"Using Graphic Turing Tests to Counter Automated DDoS Attacks Against Web Servers"*
William G. Morein, Angelos Stavrou, Debra L. Cook, Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In Proceedings of the *10th ACM International Conference on Computer and Communications Security (CCS),* pp. 8 - 19. October 2003, Washington, DC. *(Acceptance rate: 13.8%)*
81. *"EasyVPN: IPsec Remote Access Made Easy"*
Mark C. Benvenuto and Angelos D. Keromytis. In Proceedings of the *17th USENIX Large Installation Systems Administration (LISA) Conference,* pp. 87 - 93. October 2003, San Diego, CA. *(Acceptance rate: 25%)*
82. *"A Cooperative Immunization System for an Untrusting Internet"*
Kostas G. Anagnostakis, Michael B. Greenwald, Sotiris Ioannidis, Angelos D. Keromytis, and Dekai Li. In Proceedings of the *11th IEEE International Conference on Networks (ICON),* pp. 403 - 408. September/October 2003, Sydney, Australia.
83. *"Accelerating Application-Level Security Protocols"*
Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *11th IEEE International Conference on Networks (ICON),* pp. 313 - 318. September/October 2003, Sydney, Australia.
84. *"WebSOS: Protecting Web Servers From DDoS Attacks"*
Debra L. Cook, William G. Morein, Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In Proceedings of the *11th IEEE International Conference on Networks (ICON),* pp. 455 - 460. September/October 2003, Sydney, Australia.
85. *"TAPI: Transactions for Accessing Public Infrastructure"*
Matt Blaze, John Ioannidis, Sotiris Ioannidis, Angelos D. Keromytis, Pekka Nikander, and Vassilis Prevelakis. In Proceedings of the *8th IFIP Personal Wireless Communications (PWC) Conference,* pp. 90 - 100. September 2003, Venice, Italy.
86. *"Tagging Data In The Network Stack: mbuf_tags"*
Angelos D. Keromytis. In Proceedings of the *USENIX BSD Conference (BSDCon),* pp. 125 - 131. September 2003, San Mateo, CA.
87. *"The Design of the OpenBSD Cryptographic Framework"*
Angelos D. Keromytis, Jason L. Wright, and Theo de Raadt. In Proceedings of the *USENIX Annual Technical Conference,* pp. 181 - 196. June 2003, San Antonio, TX. *(Acceptance rate: 23%)*



New Bay Capital, LLC
Ex.1015-Page 2656 of 3151

88. *"Secure and Flexible Global File Sharing"*
Stefan Miltchev, Vassilis Prevelakis, Sotiris Ioannidis, John Ioannidis, Angelos D. Keromytis, and Jonathan M. Smith. In Proceedings of the *USENIX Annual Technical Conference, Freenix Track,* pp. 165 - 178. June 2003, San Antonio, TX.
89. *"Experience with the KeyNote Trust Management System: Applications and Future Directions"*
Matt Blaze, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *1st International Conference on Trust Management,* pp. 284 - 300. May 2003, Heraclion, Greece.
90. *"The STRONGMAN Architecture"*
Angelos D. Keromytis, Sotiris Ioannidis, Michael B. Greenwald, and Jonathan M. Smith. In Proceedings of the *3rd DARPA Information Survivability Conference and Exposition (DISCEX III),* volume 1, pp. 178 - 188. April 2003, Washington, DC.
91. *"Efficient, DoS-Resistant, Secure Key Exchange for Internet Protocols"*
William Aiello, Steven M. Bellovin, Matt Blaze, Ran Canetti, John Ioannidis, Angelos D. Keromytis, and Omer Reingold. In Proceedings of the *9th ACM International Conference on Computer and Communications Security (CCS),* pp. 48 - 58. November 2002, Washington, DC. *(Acceptance rate: 17.6%)*
92. *"Secure Overlay Services"*
Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In Proceedings of the *ACM SIGCOMM Conference,* pp. 61 - 72. August 2002, Pittsburgh, PA. Also available through the *ACM Computer Communications Review (SIGCOMM Proceedings),* vol. 32, no. 4, October 2002. *(Acceptance rate: 8.3%)*
93. *"Using Overlays to Improve Network Security"*
Angelos D. Keromytis, Vishal Misra, and Dan Rubenstein. In Proceedings of the *ITCom Conference,* special track on *Scalability and Traffic Control in IP Networks,* pp. 245 - 254. July/August 2002, Boston, MA. *(Invited paper)*
94. *"Designing an Embedded Firewall/VPN Gatweway"*
Vassilis Prevelakis and Angelos D. Keromytis. In Proceedings of the *International Network Conference (INC),* pp. 313 - 322. July 2002, Plymouth, England. **(Best Paper Award)**
95. *"A Study of the Relative Costs of Network Security Protocols"*
Stefan Miltchev, Sotiris Ioannidis, and Angelos D. Keromytis. In Proceedings of the *USENIX Annual Technical Conference, Freenix Track,* pp. 41 - 48. June 2002, Monterey, CA.
96. *"A Secure Plan (Extended Version)"*
Michael W. Hicks, Angelos D. Keromytis, and Jonathan M. Smith. In Proceedings of the *DARPA Active Networks Conference and Exposition (DANCE),* pp. 224 - 237. May 2002, San Francisco, CA. *(Extended version of the paper IWAN 1999 paper.)*
97. *"Fileteller: Paying and Getting Paid for File Storage"*
John Ioannidis, Sotiris Ioannidis, Angelos D. Keromytis, and Vassilis Prevelakis. In Proceedings of the *6th Financial Cryptography (FC) Conference,* pp. 282 - 299. March 2002, Bermuda. *(Acceptance rate: 25.6%)*
98. *"Offline Micropayments without Trusted Hardware"*
Matt Blaze, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *5th Financial Cryptography (FC) Conference,* pp. 21 - 40. February 2001, Cayman Islands.
99. *"Trust Management for IPsec"*
Matt Blaze, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *8th Internet Society (ISOC) Symposium on Network and Distributed Systems Security (SNDSS)* , pp. 139 - 151. February 2001, San Diego, CA. *(Acceptance rate: 24%)*



New Bay Capital, LLC
Ex.1015-Page 2657 of 3151

100. *"Implementing a Distributed Firewall"*
Sotiris Ioannidis, Angelos D. Keromytis, Steven M. Bellovin, and Jonathan M. Smith. In Proceedings of the *7th ACM International Conference on Computer and Communications Security (CCS),* pp. 190 - 199. November 2000, Athens, Greece. *(Acceptance rate: 21.4%)*

101. *"Implementing Internet Key Exchange (IKE)"*
Niklas Hallqvist and Angelos D. Keromytis. In Proceedings of the *USENIX Annual Technical Conference, Freenix Track,* pp. 201 - 214. June 2000, San Diego, CA.

102. *"Transparent Network Security Policy Enforcement"*
Angelos D. Keromytis and Jason Wright. In Proceedings of the *USENIX Annual Technical Conference, Freenix Track,* pp. 215 - 226. June 2000, San Diego, CA.

103. *"Cryptography in OpenBSD: An Overview"*
Theo de Raadt, Niklas Hallqvist, Artur Grabowski, Angelos D. Keromytis, and Niels Provos. In Proceedings of the *USENIX Annual Technical Conference, Freenix Track,* pp. 93 - 101. June 1999, Monterey, CA.

104. *"DHCP++: Applying an efficient implementation method for fail-stop cryptographic protocols"*
William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith. In Proceedings of the *IEEE Global Internet (GlobeCom),* pp. 59 - 65. November 1998, Sydney, Australia.

105. *"Automated Recovery in a Secure Bootstrap Process"*
William A. Arbaugh, Angelos D. Keromytis, David J. Farber, and Jonathan M. Smith. In Proceedings of the *5th Internet Society (ISOC) Symposium on Network and Distributed System Security (SNDSS),* pp. 155 - 167. March 1998, San Diego, CA. An older version is available as *University of Pennsylvania Technical Report MS-CIS-97-13.*

106. *"Implementing IPsec"*
Angelos D. Keromytis, John Ioannidis, and Jonathan M. Smith. In Proceedings of the *IEEE Global Internet (GlobeCom),* pp. 1948 - 1952. November 1997, Phoenix, AZ.

**Books/Book Chapters**

1. *"Voice over IP Security: A Comprehensive Survey of Vulnerabilities and Academic Research"*
Angelos D. Keromytis. Springer Briefs, ISBN 978-1-4419-9865-1, April 2011.
2. *"Buffer Overflow Attacks"*
Angelos D. Keromytis. In *Encyclopedia of Cryptography and Security, 2nd Edition.* Springer, 2011.
3. *"Network Bandwidth Denial of Service (DoS)"*
Angelos D. Keromytis. In *Encyclopedia of Cryptography and Security, 2nd Edition.* Springer, 2011.
4. *"Monitoring Technologies for Mitigating Insider Threats"*
Brian M. Bowen, Malek Ben Salem, Angelos D. Keromytis, and Salvatore J. Stolfo. In *Insider Threats in Cyber Security and Beyond,* Matt Bishop, Dieter Gollman, Jeffrey Hunker, and Christian Probst (editors), pp. 197 - 218. Springer, 2010.
5. *"Voice over IP: Risks, Threats, and Vulnerabilities"*
Angelos D. Keromytis. In *Cyber Infrastructure Security,* Tarek Saadawi and Louis Jordan (editors). Strategic Study Institute (SSI), 2010.
6. *Proceedings of the 2008 New Security Paradigms Workshop (NSPW)*
Angelos D. Keromytis, Anil Somayaji, and M. Hossain Heydari (editors).
7. *Proceedings of the 6th International Conference on Applied Cryptography and Network Security (ACNS)*



New Bay Capital, LLC
Ex.1015-Page 2658 of 3151

Steven M. Bellovin, Rosario Gennaro, Angelos D. Keromytis, and Moti Yung (editors). Lecture Notes in Computer Science (LNCS). Springer, 2008.

8. *"Insider Attack and Cyber Security: Beyond the Hacker"*
   Salvatore J. Stolfo, Steven M. Bellovin, Angelos D. Keromytis, Sara Sinclair, and Sean W. Smith (editors). Advances in Information Security Series, ISBN 978-0387773216. Springer, 2008.
9. *Proceedings of the 2007 New Security Paradigms Workshop (NSPW)*
   Kostantin Beznosov (Editor), Angelos D. Keromytis (editor), and M. Hossain Heydari (Editor).
10. *"The Case for Self-Healing Software"*
    Angelos D. Keromytis. In *Aspects of Network and Information Security: Proceedings NATO Advanced Studies Institute (ASI) on Network Security and Intrusion Detection, held in Nork, Yerevan, Armenia, October 2006,* E. Haroutunian, E. Kranakis, and E. Shahbazian (editors). IOS Press, 2007. *(By invitation, as part of the NATO ASI on Network Security, October 2005.)*
11. *"Designing Firewalls: A Survey"*
    Angelos D. Keromytis and Vassilis Prevelakis. In *Network Security: Current Status and Future Directions,* Christos Douligeris and Dimitrios N. Serpanos (editors), pp. 33 - 49. Wiley - IEEE Press, April 2007.
12. *"Composite Hybrid Techniques for Defending against Targeted Attacks"*
    Stelios Sidiroglou and Angelos D. Keromytis. In *Malware Detection,* vol. 27 of Advances in Information Security Series, Mihai Christodorescu, Somesh Jha, Douglas Maughan, Dawn Song, and Cliff Wang (editors). Springer, October 2006. *(By invitation, as part of the ARO/DHS 2005 Workshop on Malware Detection.)*
13. *"Trusted computing platforms and secure Operating Systems"*
    Angelos D. Keromytis. In *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft,* Markus Jakobsson and Steven Myers (editors), pp. 387 - 405. Wiley, 2006.
14. *"CryptoGraphics: Exploiting Graphics Cards for Security"*
    Debra Cook and Angelos D. Keromytis. Advances in Information Security Series, ISBN 0-387-29015-X. Springer, 2006.
15. *Proceedings of the 3rd Workshop on Rapid Malcode (WORM)*
    Angelos D. Keromytis (editor). ACM Press, 2005.
16. *Proceedings of the 3rd International Conference on Applied Cryptography and Network Security (ACNS)*
    John Ioannidis, Angelos D. Keromytis, and Moti Yung (editors). Lecture Notes in Computer Science (LNCS) 3531. Springer, 2005.
17. *"Distributed Trust"*
    John Ioannidis and Angelos D. Keromytis. In *Practical Handbook of Internet Computing,* Munindar Singh (editor), pp. 47/1 - 47/16. CRC Press, 2004.
18. *"Experiences Enhancing Open Source Security in the POSSE Project"*
    Jonathan M. Smith, Michael B. Greenwald, Sotiris Ioannidis, Angelos D. Keromytis, Ben Laurie, Douglas Maughan, Dale Rahn, and Jason L. Wright. In *Free/Open Source Software Development,* Stefan Koch (editor), pp. 242 - 257. Idea Group Publishing, 2004. Also re-published in *Global Information Technologies: Concepts, Methodologies, Tools, and Applications,* Felix B. Tan (editor), pp. 1587 - 1598. Idea Group Publishing, 2007.
19. *"STRONGMAN: A Scalable Solution to Trust Management in Networks"*
    Angelos D. Keromytis. Ph.D. Thesis, University of Pennsylvania, November 2001.



New Bay Capital, LLC
Ex.1015-Page 2659 of 3151

20. *"The Role of Trust Management in Distributed Systems Security"*
Matt Blaze, Joan Feigenbaum, John Ioannidis, and Angelos D. Keromytis. In *Secure Internet Programming: Issues in Distributed and Mobile Object Systems,* Jan Vitek and Christian Jensen (editors), pp. 185 - 210. Springer-Verlag Lecture Notes in Computer Science *State-of-the-Art* series, 1999.
21. *"Security in Active Networks"*
D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith. In *Secure Internet Programming: Issues in Distributed and Mobile Object Systems,* Jan Vitek and Christian Jensen (editors), pp. 433 - 451. Springer-Verlag Lecture Notes in Computer Science *State-of-the-Art* series, 1999.

**Workshops**

1. *"REASSURE: A Self-contained Mechanism for Healing Software Using Rescue Points"*
Georgios Portokalidis and Angelos D. Keromytis. To appear in the Proceedings of the *6th International Workshop on Security (IWSEC).* November 2011, Tokyo, Japan.
2. *"Taint-Exchange: a Generic System for Cross-process and Cross-host Taint Tracking"*
Angeliki Zavou, Georgios Portokalidis, and Angelos D. Keromytis. To appear in the Proceedings of the *6th International Workshop on Security (IWSEC).* November 2011, Tokyo, Japan.
3. *"The MINESTRONE Architecture: Combining Static and Dynamic Analysis Techniques for Software Security"*
Angelos D. Keromytis, Salvatore J. Stolfo, Junfeng Yang, Angelos Stavrou, Anup Ghosh, Dawson Engler, Marc Dacier, Matthew Elder, and Darrell Kienzle. In Proceedings of the *1st Workshop on Systems Security (SysSec).* July 2011, Amsterdam, Netherlands.
4. *"The SPARCHS Project: Hardware Support for Software Security"*
Simha Sethumadhavan, Salvatore J. Stolfo, David August, Angelos D. Keromytis, and Junfeng Yang. In Proceedings of the *1st Workshop on Systems Security (SysSec).* July 2011, Amsterdam, Netherlands.
5. *"Towards a Forensic Analysis for Multimedia Communication Services"*
Dimitris Geneiatakis and Angelos D. Keromytis. In Proceedings of the *7th International Symposium on Frontiers in Networking with Applications (FINA),* pp. 424 - 429. March 2011, Biopolis, Singapore.
6. *"Security Research with Human Subjects: Informed Consent, Risk, and Benefits"*
Maritza Johnson, Steven M. Bellovin, and Angelos D. Keromytis. In Proceedings of the *2nd Workshop on Ethics in Computer Security Research (WECSR).* March 2011, Saint Lucia.
7. *"Global ISR: Toward a Comprehensive Defense Against Unauthorized Code Execution"*
Georgios Portokalidis and Angelos D. Keromytis. In Proceedings of the *ARO Workshop on Moving Target Defense.* October 2010, Fairfax, VA.
8. *"Securing MANET Multicast Using DIPLOMA"*
Mansoor Alicherry and Angelos D. Keromytis. In Proceedings of the *5th International Workshop on Security (IWSEC),* pp. 232 - 250. November 2010, Kobe, Japan. *(Acceptance rate: 29%)*
9. *"Evaluating a Collaborative Defense Architecture for MANETs"*
Mansoor Alicherry, Angelos Stavrou, and Angelos D. Keromytis. In Proceedings (electronic) of the *IEEE Workshop on Collaborative Security Technologies (CoSec),* pp. 37 - 42. December 2009, Bangalore, India. *(Acceptance rate: 17.2%)*
10. *"Identifying Proxy Nodes in a Tor Anonymization Circuit"*
Sambuddho Chakravarty, Angelos Stavrou, and Angelos D. Keromytis. In Proceedings of the



New Bay Capital, LLC
Ex.1015-Page 2660 of 3151

*2nd Workshop on Security and Privacy in Telecommunications and Information Systems (SePTIS),* pp. 633 - 639. December 2008, Bali, Indonesia. *(Acceptance rate: 37.5%)*

11. *"Online Network Forensics for Automatic Repair Validation"*
Michael E. Locasto, Matthew Burnside, and Angelos D. Keromytis. In Proceedings of the *3rd International Workshop on Security (IWSEC),* pp. 136 - 151. November 2008, Kagawa, Japan. *(Acceptance rate: 19.1%)*
12. *"Return Value Predictability for Self-Healing"*
Michael E. Locasto, Angelos Stavrou, Gabriela F. Cretu, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *3rd International Workshop on Security (IWSEC),* pp. 152 - 166. November 2008, Kagawa, Japan. *(Acceptance rate: 19.1%)*
13. *"Asynchronous Policy Evaluation and Enforcement"*
Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *2nd Computer Security Architecture Workshop (CSAW),* pp. 45 - 50. October 2008, Fairfax, VA.
14. *"Race to the bottom: Malicious Hardware"*
Angelos D. Keromytis, Simha Sethumadhavan, and Ken Shepard. In Proceedings of the *1st FORWARD Invitational Workshop for Identifying Emerging Threats in Information and Communication Technology Infrastructures.* April 2008, Goteborg, Sweden. *(Invited paper)*
15. *"Arachne: Integrated Enterprise Security Management"*
Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *8th Annual IEEE SMC Information Assurance Workshop (IAW),* pp. 214 - 220. June 2007, West Point, NY.
16. *"Poster Paper: Band-aid Patching"*
Stelios Sidiroglou, Sotiris Ioannidis, and Angelos D. Keromytis. In Proceedings of the *3rd Workshop on Hot Topics in System Dependability (HotDep),* pp. 102 - 106. June 2007, Edinburgh, UK.
17. *"Data Sanitization: Improving the Forensic Utility of Anomaly Detection Systems"*
Gabriela F. Cretu, Angelos Stavrou, Salvatore J. Stolfo, and Angelos D. Keromytis. In Proceedings of the *3rd Workshop on Hot Topics in System Dependability (HotDep),* pp. 64 - 70. June 2007, Edinburgh, UK.
18. *"Bridging the Network Reservation Gap Using Overlays"*
Angelos Stavrou, David Michael Turner, Angelos D. Keromytis, and Vassilis Prevelakis. In Proceedings of the *1st Workshop on Information Assurance for Middleware Communications (IAMCOM),* pp. 1 - 6. January 2007, Bangalore, India.
19. *"Next Generation Attacks on the Internet"*
Evangelos Markatos and Angelos D. Keromytis. In Proceedings (electronic) of the *EU-US Summit Series on Cyber Trust: Workshop on System Dependability & Security,* pp. 67 - 73. November 2006, Dublin, Ireland. *(Invited paper)*
20. *"Dark Application Communities"*
Michael E. Locasto, Angelos Stavrou, and Angelos D. Keromytis. In Proceedings of the *New Security Paradigms Workshop (NSPW),* pp. 11 - 18. September 2006, Schloss Dagstuhl, Germany.
21. *"Privacy as an Operating System Service"*
Sotiris Ioannidis, Stelios Sidiroglou, and Angelos D. Keromytis. In Proceedings (electronic) of the *1st Workshop on Hot Topics in Security (HotSec).* July 2006, Vancouver, Canada.
22. *"PalProtect: A Collaborative Security Approach to Comment Spam"*
Benny Wong, Michael E. Locasto, and Angelos D. Keromytis. In Proceedings of the *7th Annual IEEE SMC Information Assurance Workshop (IAW),* pp. 170 - 175. June 2006, West Point, NY.
23. *"Adding a Flow-Oriented Paradigm to Commodity Operating Systems"*



New Bay Capital, LLC
Ex.1015-Page 2661 of 3151

Christian Soviani, Stephen A. Edwards, and Angelos D. Keromytis. In Proceedings of the *Workshop on Interaction between Operating System and Computer Architecture (IOSCA),* held in conjunction with the IEEE International Symposium on Workload Characterization, pp. 1 - 6. October 2005, Austin, TX.

24. *"Speculative Virtual Verification: Policy-Constrained Speculative Execution"* Michael E. Locasto, Stelios Sidiroglou, and Angelos D. Keromytis. In Proceedings of the *New Security Paradigms Workshop (NSPW),* pp. 119 - 124. September 2005, Lake Arrowhead, CA.
25. *"Application Communities: Using Monoculture for Dependability"* Michael E. Locasto, Stelios Sidiroglou, and Angelos D. Keromytis. In Proceedings of the *1st Workshop on Hot Topics in System Dependability (HotDep),* held in conjunction with the International Conference on Dependable Systems and Networks (DSN), pp. 288 - 292. June 2005, Yokohama, Japan.
26. *"Towards Collaborative Security and P2P Intrusion Detection"* Michael E. Locasto, Janak Parekh, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *6th Annual IEEE SMC Information Assurance Workshop (IAW),* pp. 333 - 339. June 2005, West Point, NY.
27. *"FlowPuter: A Cluster Architecture Unifying Switch, Server and Storage Processing"* Alfred V. Aho, Angelos D. Keromytis, Vishal Misra, Jason Nieh, Kenneth A. Ross, and Yechiam Yemini. In Proceedings of the *1st International Workshop on Data Processing and Storage Networking: towards Grid Computing (DPSN),* pp. 2/1 - 2/7. May 2004, Athens, Greece.
28. *"One Class Support Vector Machines for Detecting Anomalous Windows Registry Accesses"* Katherine Heller, Krysta Svore, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *ICDM Workshop on Data Mining for Computer Security,* held in conjunction with the *3rd International IEEE Conference on Data Mining,* pp. 2 - 9. November 2003, Melbourn, FL.
29. *"A Holistic Approach to Service Survivability"* Angelos D. Keromytis, Janak Parekh, Philip N. Gross, Gail Kaiser, Vishal Misra, Jason Nieh, Dan Rubenstein, and Salvatore J. Stolfo. In Proceedings of the *1st ACM Workshop on Survivable and Self-Regenerative Systems (SSRS),* held in conjunction with the *10th ACM International Conference on Computer and Communications Security (CCS),* pp. 11 - 22. October 2003, Fairfax, VA.
30. *"High-Speed I/O: The Operating System As A Signalling Mechanism"* Matthew Burnside and Angelos D. Keromytis. In Proceedings of the *ACM SIGCOMM Workshop on Network-I/O Convergence: Experience, Lessons, Implications (NICELI),* held in conjunction with the *ACM SIGCOMM Conference,* pp. 220 - 227. August 2003, Karlsruhe, Germany.
31. *"A Network Worm Vaccine Architecture"* Stelios Sidiroglou and Angelos D. Keromytis. In Proceedings of the *12th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), Workshop on Enterprise Security,* pp. 220 - 225. June 2003, Linz, Austria.
32. *"Design and Implementation of Virtual Private Services"* Sotiris Ioannidis, Steven M. Bellovin, John Ioannidis, Angelos D. Keromytis, and Jonathan M. Smith. In Proceedings of the *12th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), Workshop on Enterprise Security, Special Session on Trust Management in Collaborative Global Computing,* pp. 269 - 274. June 2003, Linz, Austria.



New Bay Capital, LLC
Ex.1015-Page 2662 of 3151

33. *"WebDAVA: An Administrator-Free Approach To Web File-Sharing"*
Alexander Levine, Vassilis Prevelakis, John Ioannidis, Sotiris Ioannidis, and Angelos D. Keromytis. In Proceedings of the *12th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), Workshop on Distributed and Mobile Collaboration,* pp. 59 - 64. June 2003, Linz, Austria.
34. *"Protocols for Anonymity in Wireless Networks"*
Matt Blaze, John Ioannidis, Angelos D. Keromytis, Tal Malkin, and Avi Rubin. In Proceedings of the *11th International Workshop on Security Protocols.* April 2003, Cambridge, England.
35. *"xPF: Packet Filtering for Low-Cost Network Monitoring"*
Sotiris Ioannidis, Kostas G. Anagnostakis, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *Workshop on High Performance Switching and Routing (HPSR),* pp. 121 - 126. May 2002, Kobe, Japan.
36. *"Toward Understanding the Limits of DDoS Defenses"*
Matt Blaze, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *10th International Workshop on Security Protocols,* Springer-Verlag Lecture Notes in Computer Science, vol. 2467. April 2002, Cambridge, England.
37. *"Toward A Unified View of Intrusion Detection and Security Policy"*
Matt Blaze, Angelos D. Keromytis, and Salvatore J. Stolfo. In Proceedings of the *10th International Workshop on Security Protocols,* Springer-Verlag Lecture Notes in Computer Science, vol. 2467. April 2002, Cambridge, England.
38. *"Efficient, DoS-resistant, Secure Key Exchange for Internet Protocols"*
William Aiello, Steven M. Bellovin, Matt Blaze, Ran Canetti, John Ioannidis, Angelos D. Keromytis, and Omer Reingold. In Proceedings of the *9th International Workshop on Security Protocols,* Springer-Verlag Lecture Notes in Computer Science, vol. 2133, pp. 40 - 48. April 2001, Cambridge, England.
39. *"Scalable Resource Control in Active Networks"*
Kostas G. Anagnostakis, Michael W. Hicks, Sotiris Ioannidis, Angelos D. Keromytis, and Jonathan M. Smith. In Proceedings of the *2nd International Workshop for Active Networks (IWAN),* pp. 343 - 357. October 2000, Tokyo, Japan.
40. *"A Secure Plan"*
Michael W. Hicks and Angelos D. Keromytis. In Proceedings of the *1st International Workshop for Active Networks (IWAN),* pp. 307 - 314. June - July 1999, Berlin, Germany. An extended version is available as *University of Pennsylvania Technical Report MS-CIS-99-14,* and was also published in the Proceedings of the *DARPA Active Networks Conference and Exposition (DANCE),* May 2002.
41. *"Trust Management and Network Layer Security Protocols"*
Matt Blaze, John Ioannidis, and Angelos D. Keromytis. In Proceedings of the *7th International Workshop on Security Protocols,* Springer-Verlag Lecture Notes in Computer Science, vol. 1796, pp. 103 - 108. April 1999, Cambridge, England.
42. *"The SwitchWare Active Network Implementation"*
D. Scott Alexander, Michael W. Hicks, Pankaj Kakkar, Angelos D. Keromytis, Marianne Shaw, Jonathan T. Moore, Carl A. Gunter, Trevor Jim, Scott M. Nettles, and Jonathan M. Smith. In Proceedings of the *ACM SIGPLAN Workshop on ML,* held in conjunction with the *International Conference on Functional Programming (ICFP),* pp. 67 - 76. September 1998, Baltimore, MD.
43. *"KeyNote: Trust Management for Public-Key Infrastructures"*
Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. In Proceedings of the *6th*



New Bay Capital, LLC
Ex.1015-Page 2663 of 3151

*International Workshop on Security Protocols,* Springer-Verlag Lecture Notes in Computer Science, vol. 1550, pp. 59 - 63. April 1998, Cambridge, England. Also available as *AT&T Technical Report 98.11.1.*

**Additional Publications**

1. *"Transport Layer Security (TLS) Authorization Using KeyNote"*
   Angelos D. Keromytis. *Request For Comments (RFC) 6042,* October 2010.
2. *"X.509 Key and Signature Encoding for the KeyNote Trust Management System"*
   Angelos D. Keromytis. *Request For Comments (RFC) 5708,* January 2010.
3. *"SSARES: Secure Searchable Automated Remote Email Storage"*
   Adam J. Aviv, Michael E. Locasto, Shaya Potter, and Angelos D. Keromytis. In the Columbia Computer Science Student Research Symposium, Fall 2006.
4. *"IP Security Policy Requirements"*
   Matt Blaze, Angelos D. Keromytis, Michael Richardson, and Luis Sanchez. *Request For Comments (RFC) 3586,* August 2003.
5. *"On the Use of Stream Control Transmission Protocol (SCTP) with IPsec"*
   Steven M. Bellovin, John Ioannidis, Angelos D. Keromytis, and Randal R. Stewart. *Request For Comments (RFC) 3554,* June 2003.
6. *"The Use of HMAC-RIPEMD-160-96 within ESP and AH"*
   Angelos D. Keromytis and Niels Provos. *Request For Comments (RFC) 2857,* June 2000.
7. *"DSA and RSA Key and Signature Encoding for the KeyNote Trust Management System"*
   Matt Blaze, John Ioannidis, and Angelos D. Keromytis. *Request For Comments (RFC) 2792,* March 2000.
8. *"The KeyNote Trust-Management System, Version 2"*
   Matt Blaze, Joan Feigenbaum, John Ioannidis, and Angelos D. Keromytis. *Request For Comments (RFC) 2704,* September 1999.

**Technical Reports/Works in Progress**

1. *"Symantec Report on Rogue Security Software, July 2008 - June 2009"*
   Marc Fossi, Dean Turner, Eric Johnson, Trevor Mack, Teo Adams, Joseph Blackbird, Mo King Low, David McKinney, Marc Dacier, Angelos D. Keromytis, Corrado Leita, Marco Cova, Jon Orbeton, and Olivier Thonnard. Symantec Technical Report, October 2009.
2. *"LinkWidth: A Method to Measure Link Capacity and Available Bandwidth using Single-End Probes"*
   Sambuddho Chakravarty, Angelos Stavrou, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-002-08,* January 2008.
3. *"Can P2P Replace Direct Download for Content Distribution?"*
   Alex Sherman, Angelos Stavrou, Jason Nieh, Cliff Stein, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-020-07,* March 2007.
4. *"A Model for Automatically Repairing Execution Integrity"*
   Michael E. Locasto, Gabriela F. Cretu, Angelos Stavrou, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-005-07,* January 2007.
5. *"Speculative Execution as an Operating System Service"*
   Michael E. Locasto and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-024-06,* May 2006.
6. *"Quantifying Application Behavior Space for Detection and Self-Healing"*



New Bay Capital, LLC
Ex.1015-Page 2664 of 3151

Michael E. Locasto, Angelos Stavrou, Gabriela F. Cretu, Angelos D. Keromytis, and Salvatore J. Stolfo. *Columbia University Computer Science Department Technical Report CUCS-017-06,* April 2006.

7. *"Bloodhound: Searching Out Malicious Input in Network Flows for Automatic Repair Validation"*
   Michael E. Locasto, Matthew Burnside, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-016-06,* April 2006.
8. *"Binary-level Function Profiling for Intrusion Detection and Smart Error Virtualization"*
   Michael E. Locasto and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-002-06,* January 2006.
9. *"A General Analysis of the Security of Elastic Block Ciphers"*
   Debra Cook, Moti Yung, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-038-05,* September 2005.
10. *"The Pseudorandomness of Elastic Block Ciphers"*
    Debra Cook, Moti Yung, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-037-05,* September 2005.
11. *"PachyRand: SQL Randomization for the PostgreSQL JDBC Driver"*
    Michael E. Locasto and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-033-05,* August 2005.
12. *"Elastic Block Ciphers: The Feistel Cipher Case"*
    Debra L. Cook, Moti Yung, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-021-04,* May 2004.
13. *"Collaborative Distributed Intrusion Detection"*
    Michael E. Locasto, Janak J. Parekh, Salvatore J. Stolfo, Angelos D. Keromytis, Tal Malkin, and Vishal Misra. *Columbia University Computer Science Department Technical Report CUCS-012-04,* March 2004.
14. *"Elastic Block Ciphers"*
    Debra L. Cook, Moti Yung, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-010-04,* February 2004.
15. *"Just Fast Keying (JFK)"*
    William Aiello, Steven M. Bellovin, Matt Blaze, Ran Canetti, John Ioannidis, Angelos D. Keromytis, and Omer Reingold. *IETF IPsec Working Group*, April 2002,.
16. *"CASPER: Compiler-Assisted Securing of Programs at Runtime"*
    Gaurav S. Kc, Stephen A. Edwards, Gail E. Kaiser, and Angelos D. Keromytis. *Columbia University Computer Science Department Technical Report CUCS-025-02,* 2002.
17. *"The 'suggested ID' extension for IKE"*
    Angelos D. Keromytis and William Sommerfeld. *IETF IPsec Working Group*, November 2001.
18. *"SPKI: ShrinkWrap"*
    Angelos D. Keromytis and William A. Simpson. *IETF SPKI Working Group*, September 1997.
19. *"Active Network Encapsulation Protocol (ANEP)"*
    D. Scott Alexander, Bob Braden, Carl A. Gunter, Alden W. Jackson, Angelos D. Keromytis, Gary J. Minden, and David Wetherall. *Active Networks Group, DARPA Active Networks Project*, August 1997.



New Bay Capital, LLC
Ex.1015-Page 2665 of 3151

20. *"Creating Efficient Fail-Stop Cryptographic Protocols"*
Angelos D. Keromytis and Jonathan M. Smith. *University of Pennsylvania Technical Report MS-CIS-96-32*, December 1996.



New Bay Capital, LLC
Ex.1015-Page 2666 of 3151

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | | |
|---|---|---|
| In re *Inter Partes* Reexamination of: | ) | |
| | ) | |
| Victor Larson et al. | ) | Control No.: 95/001,788 |
| | ) | |
| U. S. Patent No. 7,418,504 | ) | Group Art Unit: 3992 |
| | ) | |
| Issued: August 26, 2008 | ) | Examiner: Roland Foster |
| | ) | |
| For: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | ) | Confirmation No. 5823 |

Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

**DECLARATION OF DR. ROBERT DUNHAM SHORT III**

I, Robert Dunham Short III, declare as follows:

1. I have been the Chief Technology Officer of VirnetX Inc. ("VirnetX") since June 2010 and the Chief Scientist for VirnetX since May 2006. Prior to joining VirnetX, from 1994 to April 2006, I held various positions including Assistant Vice President and Division Manager at Science Applications International Corporation ("SAIC"). Prior to SAIC, I worked at ARCO Power Technologies Inc., Sperry Corporate Technology Center, and Sperry Research Center. I have a Ph.D. in Electrical Engineering from Purdue University as well as a M.S. in Mathematics and a B.S. in Electrical Engineering from Virginia Tech.

2. I am one of the named inventors of U.S. Patent No. 7,418,504 ("the '504 patent"), which I understand is the subject of the above-identified reexamination proceeding. I am familiar with the '504 patent, including its claims.

3. Prior to and at the time of the inventions claimed in the '504 patent, there was a significant and increasing concern with the security of computer network communication. The widespread connectivity between computers that was enabled by the swift increase in network access in homes and businesses also led to many security breaches as well as concerns regarding the safety of confidential information sent over computer networks. This problem received significant attention from the research and development community. Practical experience showed that there was a need for a system that could be easily and correctly used to enable secure communications, because a system that made it difficult for an end-user to enable secure communications would likely lead to a

lack of use or incorrect use. The inventions disclosed and claimed in the '504 patent and other patents in this family met this need. For instance, the inventions disclosed and claimed in the '504 patent include a domain name service for establishing secure communication links. As an example, independent claim 1 recites "[a] system for providing a domain name service for establishing a secure communication link, the system comprising[] a domain name service system configured . . . to comprise an indication that the domain name service system supports establishing a secure communication link." ('504 patent 55:49-56.) Dependent claim 8 recites that the domain name service system is connectable to a virtual private network (VPN) through the communication network and dependent claim 9 recites that the virtual private network is one of a plurality of secure communication links in a hierarchy of secure communication links. ('504 patent 56:5-10.) Further, dependent claim 16 recites that the "domain name service system is configured to support establishing a secure communication link between [a] first location and [a] second location." ('504 patent 56:40-43.) As another example, claim 27 recites that the domain name service system is configured to enable establishment of a secure communication link between a first location and a second location transparently to a user at the first location. ('504 patent 57:13-16.) The inventions combine both ease-of-use and security aspects without sacrificing one or the other.

4. As one example of the manifestation of the long-felt need, the Defense Advanced Research Projects Agency ("DARPA") funded various research programs to further the science and technology of information assurance and survivability. DARPA programs, such as the "Information Assurance" and "Dynamic Coalitions" programs, were focused on the need to provide easy-to-enable secure communications. These projects received significant funding to be spent developing technologies that could solve this need. For example, one such project entitled "Next Generation Internet" received funding in fiscal year 1998 of approximately $39.3 million, in fiscal year 1999 of approximately $49.5 million, and in fiscal year 2000 of approximately $40 million. (Ex. B-1 at VNET00219302, 319-321.) Another program funded by DARPA, "Dynamic Coalitions," was created to address the ability of the Department of Defense to quickly and easily enable secure communications over the Internet. (*See, e.g.*, Ex. B-2 at VNET00219244, 284, 298-299, 593, 625.)

5. According to DARPA officials at the time, "existing group membership protocols d[id] not support the security needs of multidimensional organizations. The overarching challenge [wa]s creating secure groups rapidly. This [wa]s a significant issue when countries [we]re faced with an operation that require[d] immediate multinational attention." (Ex. B-3 at 1.) DARPA contracted with some of the most skilled organizations in the area of secured communications in an effort to meet its security needs (e.g., NAI Labs, a division of PGP Security, Network Associates Incorporated, Los



New Bay Capital, LLC
Ex.1015-Page 2668 of 3151

Angeles, and the Microelectronics Center of North Carolina, Research Triangle Park, North Carolina, as well as Johns Hopkins University, Baltimore; Northeastern University, Boston; and Veridian-PSR, Arlington, Virginia). (*Id.* at 1.) In all, more than 15 organizations were researching the various components that made up the programs initiated by the Department of Defense. (*Id.*) However, none of these prestigious institutions came up with a solution, during the relevant time frame, close to what is disclosed and claimed in the '504 patent. (*Id.* at 1-4.) That is, they did not develop a solution that provided a domain name service for establishing a secure communication link.

6. As a second example of the long-felt need for the inventions of the '504 patent, In-Q-Tel, which is a venture capital firm that invests in companies developing cutting edge technology aimed at supporting the United States intelligence community, including the Central Intelligence Agency (CIA), funded the original development of the technology with approximately $3.4 million. In-Q-Tel's willingness to enter into a relationship with SAIC (the original assignee of the application that led to the '504 patent) for the development of this technology further evidences a long-felt need for technology that made it easy and convenient to enable secure communications.

7. A third example was the extent to which SAIC internally funded the research and development of the technology. When I was employed at SAIC, its business model was to sell hours to the federal government. SAIC was not structured to bring products to the market, which typically requires significant internal investments in research and development. In an average year during the development of the technology that led to the '504 patent, SAIC would spend approximately $2 million on internal research and development efforts. In the case of the technology claimed in the '504 patent, SAIC invested $1.7 million, which represents almost the entirety of SAIC's internal research and development budget for one whole year. A technology review committee also approved our team's patent development efforts and costs on an ongoing basis. A third party (Cambridge Strategic Management Group or CSMG) also substantiated the value of the technology. Moreover, a significant percentage of all of SAIC's patent development efforts have focused on this technology. I understand that SAIC spent one-third of its total patent portfolio efforts on our patent portfolio at that time.

8. In fact, as demonstrated in an article written before the claimed inventions of the '504 patent, it was widely recognized that providing secure remote access to a LAN or WAN was extremely difficult for IT support desks. (Ex. B-4 at 1.) In that time period, remote access was "a nightmare for support desks. Staffers never kn[e]w what combination of CPU, modem, operating system and software configuration they [were] going to have to support," and adding the commercially-available VPN software only made matters worse. (*Id.*)



New Bay Capital, LLC
Ex.1015-Page 2669 of 3151

9. This article precisely captured the computer and Internet security industry's attitude toward the tradeoff between the ease of use of a secure system, such as a VPN system, for the average computer user and the security that the VPN system provided. The article recognized that the "ease of installation isn't always a good thing: In many cases, the easier the client is to install, the less secure it is." (*Id.* at 2.) The claimed inventions of the '504 patent, which provide a domain name service for establishing a secure communication link (for example, a VPN communication link), combine both ease of use and security aspects without sacrificing one or the other.

10. Moreover, many others before and around the time of the inventions claimed in the '504 patent have attempted to solve the need of easy-to-use methods of enabling secure communications over the Internet. But, as discussed above, many of these attempts have failed. For example, despite investing enormous amounts of money and enlisting the resources of numerous prestigious institutions and their talented employees, DARPA's projects still fell far short of the claimed inventions of the '504 patent. (*See* ¶¶ 4-5, *supra.*)

11. Additionally, as discussed above, no one had yet achieved the results of the claimed inventions of the '504 patent in that time period, because remote access was "a nightmare" for support desks to handle, and adding the commercially-available VPN software was even more difficult. In fact, at this time, the security industry generally viewed ease of use and VPN security as mutually exclusive. (*See* ¶¶ 8-9, *supra.*) By providing a domain name service for establishing a secure communication link, the inventions of the '504 patent provided a system for easily establishing secure communication links without sacrificing security, thereby succeeding where others failed.

12. The claimed inventions of the '504 patent have been commercially successful, for example, through the licensing revenues they have generated for VirnetX. In July 2002, SafeNet, a leading provider of Internet security technology that is the de facto standard in the VPN industry, entered into a portfolio license with SAIC to incorporate features into SafeNet's underlying VPNs. SafeNet licensed the patents because of features disclosed and claimed in the patents, including those in the '504 patent. Microsoft has also entered into a similar license that includes the '504 patent. Microsoft entered into its license with VirnetX after it was found to have infringed two other VirnetX patents in the same family, resulting in a damages award of over one hundred million dollars, leading ultimately to a license agreement of two hundred million dollars.

13. The claimed inventions of the '504 patent were also contrary to the accepted wisdom at the time of the inventions. For example, there was a general understanding that reliable security could only be achieved through difficult-to-provision VPNs and easy-to-set-up connections could not be secure. This belief was reinforced by the IT offices of many large companies and institutions,



New Bay Capital, LLC
Ex.1015-Page 2670 of 3151

whose livelihood depended on the need for highly-trained specialists to arrange secure network connections.

14. The industry had long accepted as a fact that secure systems, such as VPN systems, would be difficult to set up, and the secure communication modes could not be easily and conveniently enabled. In a 1999 article entitled "CEOs Chew the VPN Fat" that predicted what the future held for the start-up companies that developed VPNs, the wish list did not even address the type of solutions provided by the '504 patent, such as a domain name service for establishing secure communication links. (Ex. B-5 at 1-2.)

15. The technology of the '504 patent was also met with skepticism by those skilled in the art who learned of our inventions. Sami Saydjari, a program manager for DARPA, informed Edmund Munger, a co-inventor of the '504 patent, that our technology would never be adopted. Moreover, the IT offices of many large companies and institutions expressed skepticism that secure connections could ever be enabled easily by regular computer users.

16. Several events also demonstrate praise for the inventions in the '504 patent by those in the field. As discussed above, SAIC invested a disproportionately large percentage of its internal resources in the technology. SafeNet and Microsoft have both licensed the technology of the '504 patent. A study done by CSMG also praised the inventions. Jim Rutt at Network Solutions, which was acquired by Verisign, praised and expressed significant interest in the technology and would have invested but for a change in circumstances at his company.

17. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further, that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the '504 patent.

Dated: March 29, 2012

By: /Robert Dunham Short III/

Robert Dunham Short III



New Bay Capital, LLC
Ex.1015-Page 2671 of 3151

## APPENDIX - LIST OF EXHIBITS

| EXHIBIT | DESCRIPTION |
|---|---|
| A-1 | Verdict Form from *VirnetX, Inc. v. Microsoft Corp.*, No. 6:07-CV-80 (E.D. Tex.) |
| A-5 | Press Release, VirnetX, VirnetX and Aastra Sign a Patent License Agreement (May 3, 2012), http://virnetx.com/?p=1301 |
| A-6 | D. Eastlake, Domain Name System Security Extensions, RFC 2535 (March 1999) |
| B-1 | Excerpt from Department of Defense FY 2000/2001 Biennial Budget Estimates, Feb. 1999 |
| B-2 | Collection of Reports and Presentations on DARPA Projects |
| B-3 | Maryann Lawlor, *Transient Partnerships Stretch Security Policy Management*, SIGNAL Magazine (Sept. 2001), http://www.afcea.org/signal/articles/anmviewer.asp?a=494&print=yes |
| B-4 | Joel Snyder, *Living in Your Own Private Idaho*, Network World (January 26, 1998), http://www.networkworld.com/intranet/0126review.html |
| B-5 | Tim Greene, *CEOs Chew the VPN Fat*, CNN.com (June 17, 1999), http://www.cnn.com/TECH/computing/9906/17/vpnfat.ent.idg/index.html?iref=allsearch |

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | | |
|---|---|---|
| In re *Inter Partes* Reexamination of: | ) | |
| | ) | |
| Victor Larson et al. | ) | Control No.: 95/001,851 |
| | ) | |
| U. S. Patent No. 7,418,504 | ) | Group Art Unit: 3992 |
| | ) | |
| Issued: August 26, 2008 | ) | Examiner: Roland G. Foster |
| | ) | |
| For: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | ) | Confirmation No. 1688 |

Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Commissioner:

**PETITION SEEKING WAIVER OF 37 C.F.R. § 1.943 FOR PATENT OWNER'S RESPONSE TO OFFICE ACTION OF MARCH 1, 2012**

Pursuant to 37 C.F.R. § 1.183, Patent Owner VirnetX Inc., ("VirnetX") requests that the Director waive the requirement of 37 C.F.R. § 1.943(b) limiting patent owner's responses to 50 pages in length. Specifically, VirnetX requests that the Office accept its 71-page response to the March 1, 2012, Office Action ("Office Action").[1] VirnetX is submitting this petition concurrently with its response.

To the extent that entry and consideration of this petition requires suspension of any rules, suspension is requested pursuant to 37 C.F.R. § 1.183. In addition, a petition fee of $400 is being submitted with this petition. If there is any other fee due in connection with the filing of this petition, please charge the fee to Deposit Account 06-0916.

Rule 1.943(b) states that "[r]esponses by the patent owner and written comments by the third party requester [cannot] exceed 50 pages in length, excluding . . . reference materials." 37 C.F.R. § 1.943(b). VirnetX seeks entry and consideration of this petition so that it can comprehensively address all the issues raised by the Examiner in the Office Action. Specifically, the Office Action adopted twenty-one grounds of rejections based on nineteen different combinations of references proposed by the third-party requester, Cisco Systems, Inc. ("Cisco"). In doing so, the Office Action

[1] The listed page and word count excludes the pages and words that constitute the "amendments, appendices of claims, and reference materials" as the Office has interpreted that language of 37 C.F.R. § 1.943(b).

relied on and incorporated by reference corresponding portions of the 39 pages of Cisco's request and 328 pages of accompanying claim charts. VirnetX seeks adequate opportunity to comprehensively address the issues raised in the Office Action.

VirnetX, therefore, requests that the Director waive the page-limit requirements of § 1.943(b) and permit VirnetX to submit an Office Action response containing 71 pages and two supporting declarations.

## I. BACKGROUND

On December 13, 2011, Cisco Systems, Inc. initiated an *inter partes* reexamination of all claims 1-60 of the '504 patent. In its request, Cisco proposed twenty-one rejections based on nineteen different combinations of references. The Office granted Cisco's request for reexamination of all claims 1-60 of the '504 patent and assigned it control no. 95/001,851 ("'1,851 proceeding"). (*See* Order Granting/Denying Request for *Inter Partes* Reexamination.) On March 1, 2012, the Office issued an Office Action, adopting each of the twenty-one rejections proposed by Cisco in its request for reexamination. (*See* 03/01/12 Office Action, "Office Action" or "OA" at 4-18.)

In adopting Cisco's proposed rejections, the Office Action "adopted and incorporated by reference" the corresponding portions of the 39 pages of Cisco's request and the accompanying claim charts F1-F3. (*See, e.g., id.* at 4, 11, 16, 18.)

VirnetX's response to the Office Action is 71 pages long. In addition, VirnetX is submitting two declarations, one from Dr. Robert Dunham Short III (one of the inventors of the '504 patent) that was previously submitted in control no. 95/001,788 and another one from Angelos D. Keromytis Ph.D. (an expert). The declaration of Dr. Short presents facts regarding secondary considerations and the declaration of Dr. Keromytis discusses how one of ordinary skill in the art would have understood the references cited in the Office Action.

## II. ARGUMENT

Under 37 C.F.R. § 1.943(b), "[r]esponses by the patent owner and written comments by the third party requester [cannot] exceed 50 pages in length, excluding . . . reference materials." Because of the numerous issues raised in the Office Action and the incorporation by reference of 39 pages of Cisco's request and 328 pages of accompanying claim charts, VirnetX requests that the Office accept its response that has 71 pages. VirnetX has made every effort to pare down its response, but submits that limiting its response to 50 pages would severely compromise its ability to fully address the issues raised in the Office Action.

VirnetX's response seeks to comprehensively address the rejections adopted by the Examiner. In adopting each of the twenty-one grounds of rejection proposed by Cisco, the Examiner



New Bay Capital, LLC
Ex.1015-Page 2674 of 3151

incorporated by reference corresponding portions of the 39 pages of Cisco's request and 328 pages of accompanying claim charts. VirnetX's response to the Office Action seeks to address all the issues raised by the Examiner. Thus, given all of the issues, justice requires that the Office allow VirnetX to file its response, which contains 71 pages.

As noted above, with its response, VirnetX is also submitting two declarations, one by Dr. Short, one of the inventors, and another by Dr. Keromytis, an expert. The declaration of Dr. Short presents facts regarding secondary considerations and the declaration of Dr. Keromytis discusses how one of ordinary skill in the art would have understood the references cited in the Office Action. Given their content, VirnetX does not believe that either declaration counts towards the page limit. Nevertheless, should the Office decide to include portions of either declaration in the page count for the response, VirnetX requests that the Office waive the requirements of Rule 1.943(b) and permit it to submit these declarations with its response. Indeed, even if the Office were to count the declarations towards the page limit, the total number of pages representing the response and the declarations would still be substantially less than the 39 pages of Cisco's request and 328 pages of accompanying claim charts relied upon and incorporated by reference in the Office Action.

## III. CONCLUSION

For the foregoing reasons, VirnetX requests that the Office grant this petition and accept its Office Action response, which contains 71 pages and exceeds the page count limitations imposed by 37 C.F.R. § 1.943(b).

Respectfully submitted,
FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: June 1, 2012

By: /Joseph E. Palys/
Joseph E. Palys
Reg. No. 46,508



New Bay Capital, LLC
Ex.1015-Page 2675 of 3151

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | | |
|---|---|---|
| In re *Inter Partes* Reexamination of: | ) | |
| | ) | |
| Victor Larson et al. | ) | Control No.: 95/001,851 |
| | ) | |
| U. S. Patent No. 7,418,504 | ) | Group Art Unit: 3992 |
| | ) | |
| Issued: August 26, 2008 | ) | Examiner: Roland G. Foster |
| | ) | |
| For: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | ) | Confirmation No. 1688 |
| | ) | |

Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Commissioner:

**CERTIFICATE OF SERVICE**

Pursuant to 37 C.F.R. §§ 1.248 and 1.903 and M.P.E.P. § 2666.06, the undersigned attorney for the Patent Owner certifies that copies of the following documents:

1. Transmittal Letter (2 pages);
2. Patent Owner's Response to Office Action (71 pages);
3. Declaration of Angelos D. Keromytis, Ph.D. (31 pages) with appended *curriculum vitae*;
4. Declaration of Dr. Robert Dunham Short III in control no. 95/001,788 (5 pages);
5. Appendix - List of Exhibits (1 page);
6. Exhibits Listed on Appendix; and
7. Petition Seeking Waiver of 37 C.F.R. § 1.943 for Patent Owner's Response to Office Action of March 1, 2012 (3 pages); and
8. Certificate of Service (2 pages)

were served by first-class mail on June 1, 2012 on counsel for the third party Requester at the following address:

Haynes and Boone, LLP
IP Section
2323 Victory Avenue, Suite 700
Dallas, TX 75219

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: June 1, 2012

By: /Joseph E. Palys/
Joseph E. Palys
Reg. No. 46,508



New Bay Capital, LLC
Ex.1015-Page 2677 of 3151

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

| | | |
|---|---|---|
| In re patent of Larson et al. | § | *Inter Partes* Reexamination |
| | § | Control No. 95/001,851 |
| U.S. Patent No. 7,418,504 | § | |
| | § | Group Art Unit: 3992 |
| Issued: August 26, 2008 | § | |
| | § | Examiner: Roland Foster |
| Title: AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES | § | Confirmation No.: 1688 |

**COMMENTS BY THIRD PARTY REQUESTER**

**PURSUANT TO 37 C.F.R. § 1.947**

Mail Stop *Inter Partes* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

On June 1, 2012, the Patent Owner filed the Patent Owner's Response to Office Action ("Response") regarding the Office Action mailed March 1, 2012 ("the March 1 Office Action") in connection with the above-identified *inter partes* reexamination proceeding, which was initiated by a Request for *Inter Partes* Reexamination filed December 13, 2011 ("the Request").

It is respectfully requested, for the reasons identified below, that the Examiner:

(i) maintain his rejection of, and issue an action closing prosecution for, the original claims 1-60, and

(ii) deem the arguments advanced by the Patent Owner in the Response to be erroneous, improper, and/or unpersuasive.

In the context of this *inter partes* reexamination, the standard provided in MPEP § 2111 for claim interpretation during patent examination may be applied whereas a different standard may be used by a court in litigation. The Patent Office is not required to interpret claims in the same manner as a court would interpret claims in an infringement suit.

## TABLE OF CONTENTS


I. Comments on the Patent Owner Response of June 1, 2012 ..... 1

A. Response to Patent Owner's Arguments Regarding the Rejections Based on Lendenmann ..... 1

1. Overview of Lendenmann ..... 1

2. Response to Patent Owner's Arguments Regarding the Rejection of Claims 1-3, 5, 6, 14-30, 33-54, & 57-60 Under 35 U.S.C. § 102(b) based on *Lendenmann* (Issue 1) ..... 2

3. Response to Patent Owner's Arguments Regarding Rejection of Claims 1-3, 5, 6, 14-30, 33-54, and 57-60 Under 35 U.S.C. § 103(a) Based on Lendenmann (Issue 2) ..... 11

4. Response to Patent Owner's Arguments Regarding Rejection of Claim 7 Based on Lendenmann in view of Wesinger (Issue 3) ..... 13

5. Response to Patent Owner's Arguments Regarding Rejection of Claims 8 & 9 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz (Issue 4) ..... 13

6. Response to the Patent Owner's Argument Regarding the Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz and Schneider (Issue 5) ..... 14

7. Response to the Patent Owner's Argument Regarding the Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz and Martin (Issue 6) ..... 15

8. Response to the Patent Owner's Argument Regarding the Rejection of Claims 12 & 13 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz and RFC 793 ..... 16

9. Response to the Patent Owner's Argument Regarding the Rejection of Claims 31, 32, 55, & 56 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Ludwig and RFC 793 ..... 17

B. Aziz ..... 17

1. Overview of Aziz ..... 17

2. Response to Patent Owner's Arguments Regarding the Rejection of Claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 102(b) Based on Aziz (Issue 9) ..... 17




New Bay Capital, LLC
Ex.1015-Page 2679 of 3151

3. Response to Patent Owner's Arguments Regarding the Rejection of Claims 1, 2, 5-9, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 103(a) Based on Aziz (Issue 10) ..... 26

4. Response to Patent Owner's Arguments Regarding the Rejection of Claims 3, 4, & 26 Under 35 U.S.C. § 103(a) Based on Aziz in View of Lawton (Issue 11) ..... 27

5. Response to Patent Owner's Arguments Regarding the Rejection of Claim 9 Under 35 U.S.C. § 103(a) Based on Aziz in View of Franaszek (Issue 12) ..... 28

6. Response to Patent Owner's Arguments Regarding the Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on Aziz in View of Schneier (Issue 13) ..... 31

7. Response to Patent Owner's Arguments Regarding the Rejection of Claims11-13 Under 35 U.S.C. § 103(a) Based on Aziz in View of Martin (Issue 14) ..... 31

8. Response to Patent Owner's Arguments Regarding the Rejection of Claims 29-32 & 53-56 Under 35 U.S.C. § 103(a) Based on Aziz in View of Ludwig (Issue 15) ..... 33

C. Kiuchi and Pfaffenberger ..... 33

1. Overview of Kiuchi ..... 33

2. Overview of Pfaffenberger ..... 34

3. Response to Patent Owner's Arguments Regarding the Rejection of Claims 1-4, 6, 8-10, 12-19, 22, 24-30, 33, 34, 36-43, 46, 48-54, and 57-60 Under 35 U.S.C. § 103(a) Based on Kiuchi in View of Pfaffenberger (Issue 16) ..... 34

4. Response to Patent Owner's Arguments Regarding the Rejection of Claims 5, 23, & 47 Under 35 U.S.C. § 103(a) Based on Kiuchi in View of Pfaffenberger and Rivest (Issue 17) ..... 46

5. Response to Patent Owner's Arguments Regarding the Rejection of Claim 7 Under 35 U.S.C. § 103(a) Based on Kiuchi in View of Pfaffenberger and Borella (Issue 18) ..... 46

6. Response to Patent Owner's Arguments Regarding the Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on Kiuchi in View of Pfaffenberger and Martin (Issue 19) ..... 46

7. Response to Patent Owner's Arguments Regarding the Rejection of Claims 20, 21, 35, 44, & 45 Under 35 U.S.C. § 103(a) Based on Kiuchi in View of Pfaffenberger and Broadhurst (Issue 20) ..... 47

8. Response to Patent Owner's Arguments Regarding the Rejection of Claim 31, 33, 35, & 56 Under 35 U.S.C. § 103(a) Based on Kiuchi in View of Pfaffenberger and Ludwig (Issue 21) ..... 47



New Bay Capital, LLC
Ex.1015-Page 2680 of 3151

D. Response to Patent Owner's Argument That Secondary Considerations Demonstrate Non-Obviousness ..... 48

II. Conclusion ..... 49



New Bay Capital, LLC
Ex.1015-Page 2681 of 3151

## LIST OF EXHIBITS


Exhibit G:[1] Malkin, Gary, "Dial-in Virtual Private Networks Using Layer 3 Tunneling"

Exhibit H: Ortiz Jr., Sixto, "Virtual Private Networks: Leveraging the Internet"


[1] Exhibits A-F were presented in the Request for Reexamination.



New Bay Capital, LLC
Ex.1015-Page 2682 of 3151

## COMMENTS

These comments are based on an interpretation of the claims appropriate to this proceeding. In the context of this *inter partes* reexamination, the standard provided in MPEP § 2111 for claim interpretation during patent examination may be applied whereas a different standard may be sued by a court in litigation. The Patent Office is not required to interpret claims in the same manner as a court would interpret claims in an infringement suit.

### I. Comments on the Patent Owner Response of June 1, 2012

The Patent Owner's response to the Office Action consists, almost entirely, of flawed arguments that repeatedly commit several errors. First, the Patent Owner makes numerous attempts to rewrite the claims—adding and changing limitations—solely through attorney argument. These creative reinterpretations of the claims are wholly improper, especially since the Patent Owner had an opportunity to amend the claims to recite additional imitations but chose not to do so. Second, the Patent Owner selectively ignores the core teachings of the prior art references that were relied upon in the Examiner's rejections. For example, the Patent Owner repeatedly focuses on each prior art reference's description of basic name server functionality—such as returning a network address for a requested domain name—and then concludes that the prior art references were "distinguished" in the '504 patent's description of the admitted prior art. In doing so, the Patent Owner conveniently ignores the additional security-related teachings that the Examiner relied upon in rejecting the claims.

The Patent Owner also makes numerous bare assertions that the claims are patentable over the prior art of record, but in general fails to "specifically point out each supposed error in the examiner's action." (MPEP § 2666.) The Patent Owner frequently resorts to parroting the claim language along with a "general allegation that the claims define a patentable invention, without specifically pointing out how the language of the claims patentably distinguishes them over the references." (*Id.*) As such, the Patent Owner provides no basis on which the Examiner's rejections might be revisited. Nonetheless, Requester provides the following detailed analysis to demonstrate the correctness of the Examiner's rejections, which should be made final.

#### A. Response to Patent Owner's Arguments Regarding the Rejections Based on Lendenmann

##### 1. Overview of Lendenmann

"Understanding OSF DCE 1.1 for AIX and OS/2" by Ralf Lendenmann



New Bay Capital, LLC
Ex.1015-Page 2683 of 3151

("Lendenmann"), was published in October 1995. Publicly available over a year before the earliest claimed priority date (Oct. 30, 1998), Lendenmann is prior art under 35 U.S.C. § 102(b).

Lendenmann describes the Open Software Foundation (OSF) Distributed Computing Environment (DCE) software system that provides a broad set of name resolution and security features to support communications across computer networks. These services are tightly integrated, with the name server itself performing some security functions, such as checking a requester's authorization to access information associated with a particular domain name. (*See* Lendenmann at 34.) Lendenmann also describes setting aside a specific portion of the namespace for secure servers. (*See* Lendenmann at 28.)

### 2. Response to Patent Owner's Arguments Regarding the Rejection of Claims 1-3, 5, 6, 14-30, 33-54, & 57-60 Under 35 U.S.C. § 102(b) based on *Lendenmann* (Issue 1)

#### a. Independent Claim 1

The Examiner correctly determined that Lendenmann discloses "a domain name service system configured to … comprise an indication that the domain name service system supports establishing a secure communication link," because Lendenmann discloses:

- A domain name service system: "The directory service component that controls names inside a cell is called the Cell Directory Service (CDS)." (Lendenman at 21.)
- Several different indications that the system supports establishing a secure communication link:
    - The Cell Directory Service (CDS) returns the network address corresponding to a secure domain (Lendenmann at 21) and binding information for communicating with a server managing the security namespace. (Lendenmann at 27.)
    - The Cell Directory Service (CDS) is "integrated into the security service" and includes "ACL [access control list] management software, incorporated into all CDS clerks and servers, [that] performs access checking for incoming requests." (Lendenmann at 34.)
    - The Cell Directory Service (CDS) is accessible via a remote procedure call (Lendenmann at 173, 9), which includes facilities for authenticating and encrypting communications. (Lendenmann at 192.)
    - The Cell Directory Service (CDS) maintains and reports its status with respect to



New Bay Capital, LLC
Ex.1015-Page 2684 of 3151

security failures, including an indication that it has *not* suffered from any security errors. (Lendenmann at 37.)

- The Cell Directory Service (CDS) is part of a broader Distributed Computing Environment that includes online documentation describing its secure communications capabilities.

Accordingly, Lendenmann discloses the claimed "indication" in variety of ways, any one of which teaches the limitation. Thus, Lendenmann discloses "a domain name service system configured to ... comprise an indication that the domain name service system supports establishing a secure communication link" as recited in the claims.

**(i) Lendenmann's Cell Directory Service for "Returning the Network Address Corresponding to a Secure Domain Name" Discloses an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link."**

Patent Owner argues, on pages 6-8, that Lendenmann teaches "a CDS lookup process that simply returns a name, taking no measures to provide any indication that the CDS supports establishing a secure communication link." (Response at 7.) The Patent Owner's argument is incorrect because it fails to take into account the full teachings of Lendenmann with respect to the Cell Directory Service (CDS).

Lendenmann teaches that the Cell Directory Service (CDS) includes numerous secure communication features that go beyond merely resolving a name into a network address. For example, Lendenmann's directory service stores not just ordinary name entries, but also "[s]ome services, such as the Security Service (*sec*) ..., connect into the name space by means of specialized CDS entries, called ***junctions***. A junction entry contains binding information that enables a client to connect to a directory server outside of the directory service." (Lendenmann at 27.) The "binding handles are annotated with security information" (*Id.* at 185). Thus, Lendenmann teaches that the Cell Directory Service (CDS) includes a "junction" for the security namespace, and the junction is annotated with security information. Lendenmann further teaches that the "security namespace is managed by the registry service of the DCE security component" (*Id.*), and requesters "communicate with the registry server via authenticated RPC." (*Id.* at 49.)

In summary, Lendenmann teaches that specialized entries in the Cell Directory Service,



New Bay Capital, LLC
Ex.1015-Page 2685 of 3151

called junctions, enable the establishment of a secure communication link with registry service that manages security-related domain names. A junction is therefore an "indication that the domain name service system supports establishing a secure communication link."

The Patent Owner then continues by comparing Lendenmann's disclosure to the admitted prior art in the '504 patent specification. Absent from this comparison, however, is analysis comparing Lendenmann's disclosure to the claim limitation at issue. Lendenmann teaches not only name resolution capabilities, but also an indication of support for establishing a security communication link (e.g., returning network address and binding information corresponding to a secure domain). Patent Owner's analysis of the two prior art systems—without addressing the Lendenmann's teaching of an "indication"—is not relevant.

In summary, the Patent Owner fails to show how any distinction between the teachings of Lendenmann and the limitations recited in Claim 1. The Examiner's rejection is proper and should be made final.

**(ii) Lendenmann's Cell Directory Service "Integrated with the Security Services" Discloses an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link."**

The Patent Owner argues that the "Security Service's gatekeeping function, ... has no bearing on the operations of the alleged domain name service system, CDS." (Response at 8.) The Patent Owner's argument is that because Lendenmann's Security Service also plays a role in enabling the establishment of a secure communication link, the Cell Directory Service (CDS) must not play any role. The Patent Owner's argument is incorrect because it fails to take into account the full teachings of Lendenmann with respect to the Cell Directory Service (CDS).

Lendenmann teaches that only authorized users can access name information from the Cell Directory Service (CDS). The CDS *itself*, determines whether a user is authorized to access requested data: "ACL [access control list] management software, ***incorporated into all CDS clerks and servers***, [that] performs access checking for incoming requests." (Lendenmann at 34 (emphasis added).) Thus, it will not operate for unauthorized users. This aspect of the Cell Directory Service (CDS) is an "indication that the domain name service system supports establishing a secure communication link."

In addition, Lendenmann teaches that the "security namespace is managed by the registry service of the DCE security component." (Lendenmann at 27.) In other words, Lendenmann's



New Bay Capital, LLC
Ex.1015-Page 2686 of 3151

security service also provides domain name services and is properly considered part of Lendenmann's disclosure of a "domain name service system."

The Patent Owner also argues, at pages 9-10, that security integration like that taught by Lendenmann was already known in "conventional" domain name systems. But in fact, the '504 Patent expressly states that the *lack* of authorization checks was a deficiency in such systems: "The conventional scheme suffers from certain drawbacks. For example, any user can perform a DNS request." ('504 Patent, 39:43-45.) Lendenmann teaches how using a Cell Directory Service (CDS), which integrates certain security features, can solve this problem by "perform[ing] access checking for incoming requests." (Lendenmann at 34.) The Examiner's rejection was proper.

**(iii) Lendenmann's Binding Handles Disclose an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link."**

The Patent Owner concedes that Lendenmann teaches that "binding handles provide a network address and other identification-related information" and that "binding handles are annotated with security information." (Response at 10.) The Patent Owner then argues that Lendenmann fails to indicate support for establishing a secure communication link because the Cell Directory Service (CDS) "only provides partial binding information." (*Id.*)

Once again, the Patent Owner's argument is incorrect because it fails to take into account the full teachings of Lendenmann. Lendenmann teaches that "well-known endpoints are stored in CDS" so that "clients obtain *fully bound handles*." (Lendenmann at 186.) Even though a "server does not need to register well-known end points ... by registering well-known endpoints, the server ensures that clients can always obtain them." (Lendenmann at 200.) Lendenmann further states that binding handles provided by the Cell Directory Service (CDS) are "*usually* incomplete" (Lendenmann at 208 (emphasis added)), thus expressly contemplating that some binding handles are "complete." Furthermore, Lendenmann teaches that "specialized CDS entries, called ***junctions***... contain[] binding information that enables a client to connect to a directory server outside of the directory service." (Lendenmann at 27.) An example of a junction is a reference to the Security Service. (*Id.*)

In summary, Lendenmann teaches that entries in the Cell Directory Service (CDS) can provide either partially-bound or fully-bound handles (which are "annotated with security information" (Lendenmann at 185)) for a requested domain name. The Patent Owner's argument to the contrary is incorrect, and the Examiner's rejection is proper.



New Bay Capital, LLC
Ex.1015-Page 2687 of 3151

### (iv) Lendenmann's "Authentication Challenge" Discloses an "Indication That the Domain Name Service System Supports Establishing a Secure Communication Link."

The Patent Owner argues, on pages 11-12, that the use of "'mutual authentication' for running authenticated RPCs" is unrelated to the Cell Directory Service (CDS).

This argument fails to take into account Lendenmann's disclosure that communications with the Cell Directory Service (CDS) occur *via* remote procedure calls (RPCs): "Other DCE services use RPCs; they are also client/server applications. RPC is the basis for DCE." (Lendenmann at 173.) As Lendenmann illustrates in Fig. 3, the remote procedure call is a foundational block on which other services, including the Cell Directory Service (CDS), rely.

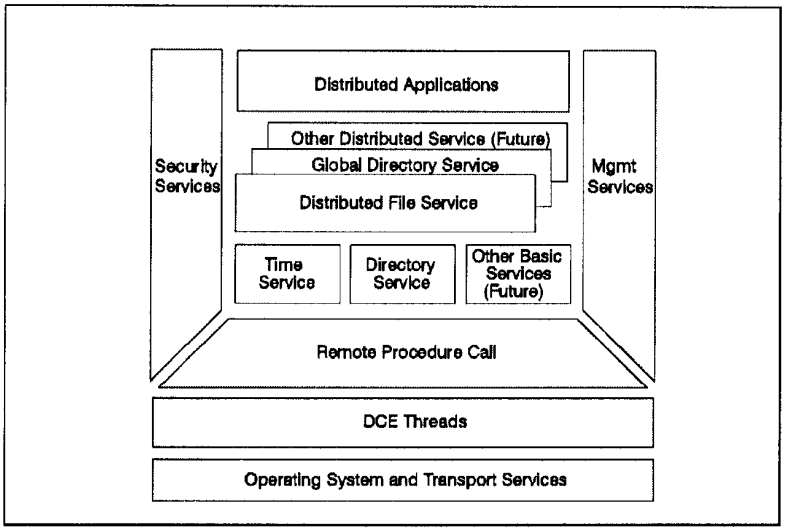


**LENDENMANN FIG. 3.**

Lendenmann further teaches that the Cell Directory Service (CDS) supports authentication and authorization of users' requests for information. (*See* Lendenmann at 34.) Accordingly, Lendenmann's description of authenticated RPC communications applies to all authenticated RPC communications, including those between a user and the Cell Directory Service (CDS):

> The client RPC runtime encrypts the RPC call with the session key and sends it to the server's runtime together with the ticket. The server immediately challenges the client by sending it a randomly generated number which the client has to encrypt with the session key and return to the server.

(Lendenmann at 194.) Thus, in the context of verifying a client's authenticity to access a domain



New Bay Capital, LLC
Ex.1015-Page 2688 of 3151

name entry, the server that issues a challenge is the Cell Directory Service (CDS). The Examiner properly determined that this challenge is an "indication that the domain name service system supports establishing a secure communication link." The Examiner's rejection is proper.

**(v) Lendenmann's "Server Status" Counters Provide an "Indication That the Domain Name Service System Supports Establishing a Secure Communications Link."**

The Patent Owner argues, on pages 12-13, that the security status of the Cell Directory Service (CDS) server is inapplicable because "Lendenmann provides for a separate Security Service to handle any steps for establishing security measures in communications." (Response at 13.)

The Patent Owner is incorrect. Lendenmann teaches that the Cell Directory Service (CDS) plays a pivotal role in establishing secure communications. For example, Lendenmann teaches that authorization decisions are the responsibility of each server in the Distributed Computing Environment: "Once the server obtains the client's authorization information, [it] ***is up to it what to do with this information***. To perform an authorization check.... ***The server can also implement an ACL manager*** that allows a security administrator to maintain permission in a standardized way...." (Lendenmann at 193 (emphasis added).) Lendenmann further teaches that the Cell Directory Service (CDS) uses its own ACL manager to make security authorization decisions: "***CDS ACL management software, incorporated into all CDS clerks and servers, performs access checking for incoming requests***. When a principal requests an operation on a CDS name or a privileged operation on a CDS clerk or server, ACL management software examines the ACL entry associated with that name or principal name and grants or denies the operation." (Lendenmann at 34 (emphasis added).)

Thus, Lendenmann teaches that the Cell Directory Service (CDS) itself provides security components, including at least the access control list (ACL) management software. The server status report, including a count of "Security Failures," is therefore a further teaching by Lendenmann of an "indication that the domain name service system supports establishing a secure communication link" as recited in Claim 1.

**(vi) Lendenmann's "Online Documentation" Discloses an "Indication That the Domain Name Service System Supports Establishing a Secure Communications Link."**

The Patent Owner argues, at pages 13-15, that the online documentation included with



New Bay Capital, LLC
Ex.1015-Page 2689 of 3151

Lendenmann's DCE software system "does not indicate anything about the CDS, much less that the CDS is configured to support establishing a secure communication link." (Response at 14.)

This argument fails because it does not take into account the full teachings of Lendenmann. Lendenmann teaches that the online documentation for DCE is installed onto a Cell Directory Service (CDS) server. A specific example describes installing both the CDS Server and Online Documentation onto a single computer: "This example shows the selection of the DCE Client, Security Server and CDS Server components. In fact, we also selected ... the Online Documentation for our installation." (Lendenmann at 118.) It is understood that the online documentation for the CDS Server would describe its features, such as authenticated and encrypted remote procedure calls, that support the establishment of a secure communication link. Thus, the online documentation stored on the CDS server is an "indication that the domain name service system supports establishing a secure communication link."

In summary, the Patent Owner fails to show any distinction between the teachings of Lendenmann and the limitations recited in Claim 1. The Examiner's rejection is proper and should be made final.

**b. Independent Claims 36 and 60**

Regarding independent claims 36 and 60, the Patent Owner's relies on the same improper and erroneous arguments related to claim 1 that are refuted above. Accordingly, the rejections of claims 36 and 60 should be reaffirmed and made final.

**c. Dependent Claims 5, 23, and 47**

Claim 5 depends from claim 1 and further recites that "the domain name service system is configured to authenticate the query using a cryptographic technique." Claims 23 and 47 recite similar limitations.

The Examiner properly rejected these claims as anticipated by Lendenmann because Lendenmann teaches:

- Communications with the Cell Directory Service (CDS), such as queries for a network address, occur via remote procedure call (RPC) (Lendenmann at 173).
- To authenticate a client request received via remote procedure call (RPC), a server "challenges the client by sending it a randomly generated number which the client has to encrypt with the session key and return to the server." (Lendenmann at 194.)
- At the server, the "random number is decrypted, and if it matches, everything is set for



New Bay Capital, LLC
Ex.1015-Page 2690 of 3151

authenticated RPC." (Lendenmann at 194.)

Accordingly, Lendenmann teaches that (i) Cell Directory Service (CDS) server receives requests to resolve domain names via remote procedure call; (ii) the Cell Directory Service (CDS) server issues an encryption challenge to the client to prove its identity; and (iii) the client returns an encrypted response; and (iv) the Cell Directory Service (CDS) server decrypts the client response and compares the decrypted value to the challenge value initially sent to the client, thereby verifying the client's query. The Examiner correctly determined that Lendenmann teaches "to authenticate the query using a cryptographic technique" as recited in claim 5 and the similar limitations in claims 23 and 47.

The Patent Owner argues that "the Security Service—not the CDS—performs the alleged authentication process in acting as a gatekeeper to the information in the CDS." (Response at 16.) This argument fails for at least two reasons. First, the Patent Owner is attempting to add a new limitation through attorney argument, *e.g.*, "wherein the authentication is performed without reference to any external security service or gatekeeper." It is wholly improper to attempt to rewrite the claim in this fashion.

Second, Lendenmann expressly teaches that the "DCE Security Service ... has been integrated with the other DCE services" (Lendenmann at 9), such as the Cell Directory Service (CDS). Lendenmann further teaches that "The CDS ... is ***integrated*** into the security service." (Lendenmann at 194.) Thus, even accepting the Patent Owner's argument that a software component of the Security Service performs the cryptographic steps to decrypt and verify a client's response, these steps are taken because the Cell Directory Service (CDS) directed the security service to do so.

The Patent Owner further argues that the "separate Security Service component may permit or deny access to the CDS." (Response at 16.) This argument also fails for numerous reasons. First, the Patent Owner is once again attempting to improperly add a limitation to the claim by attorney argument. Claim 5 does not recite any limitation about permitting or denying access—or even any limitation about verifying a user's authorization. Instead, Claim 5 recites "to *authenticate* the query," but does not recite any limitation relating to *authorization* as the Patent Owner argues.

Second, Lendenmann explicitly states that it is the access control list (ACL) management software—part of the Cell Directory Service (CDS)—that verifies a client's authorization: "CDS



New Bay Capital, LLC
Ex.1015-Page 2691 of 3151

ACL management software, ***incorporated into all CDS clerks and servers***, performs access checking for incoming requests." (Lendenmann at 34 (emphasis added).) The Patent Owner's assertion that authorization is performed by a separate Security Service is incorrect.

**d. Dependent Claims 16, 17, 27, 33, 40, 41, 51, & 57**

With respect to claims 16, 17, 40, and 41, the Patent Owner relies on its arguments about claim 1 that, as shown above, are improper and erroneous. The Examiner properly rejected claims 16, 17, 40, and 41, and those rejections should be reaffirmed and made final.

Claims 27, 33, 51, and 57 recite that the domain name service system is configured to "enable establishment of a secure communication link." Regarding these claims, the Patent Owner argues that "CDS merely returns incomplete binding handles containing server identification information" and "performs no functions beyond those ... characterizing a conventional domain name service." (Response at 18.) These arguments fail for several reasons. First, the Patent Owner is attempting to add a limitation by attorney argument, e.g., "wherein the domain name service system is unconventional." Accordingly, this argument is improper. Second, Lendenmann teaches that the Cell Directory Service (CDS) can be configured to return *complete* binding handles. (*See* Lendenmann at 186 and discussion, *supra*, p. 5.) Finally, Lendenmann discloses that the Cell Directory Service (CDS) does much more than other prior art domain name systems, including that the Cell Directory Service (CDS) itself is capable of engaging in a secure communication link. Lendenmann teaches that the Cell Directory Service (CDS) communicates via an authenticated remote procedure call (RPC), which may further be encrypted for security. (*See* Lendenmann at 192.) Thus, Lendenmann teaches to "enable establishment of a secure communication link" as recited in claims 27, 33, 51, and 57.

The Examiner's rejection was proper and should be made final.

**e. Dependent Claims 24 & 48**

The Examiner properly rejected claims 24 and 48 because Lendenmann teaches:

- The Cell Directory Service (CDS) includes entries with "security specific names." (Lendenmann at 28.) For example, name "/.:/subsys/dce/sec/master" is "the server entry for the master security server." (Lendenmann at 28.)
- The Cell Directory Service (CDS) assists clients by resolving names to network addresses. (*See* Lendenmann at 21.)
- "Not all DCE names are stored directly in the DCE Directory Service.... [S]ome



New Bay Capital, LLC
Ex.1015-Page 2692 of 3151

services, such as the Security Service (*sec*) ..., connect into the name space by means of specialized CDS entries, called ***junctions***. A junction entry contains binding information that enables a client to connect to a directory server outside of the directory service." (Lendenmann at 27.) The "binding handles are annotated with security information" (*Id.* at 185).

- "The security namespace is managed by the registry service of the DCE security component...." (Lendenmann at 27.)
- A master security server provides the "Registry Service (RS) – A replicated service which maintains the cell's security database." (Lendenmann at 45.) "Applications communicate with the registry server via authenticated RPC." (Lendenmann at 49.)

Accordingly, Lendenmann teaches that (i) the Cell Directory Service (CDS) stores the names of servers with security-specific names, (ii) the Cell Directory Service (CDS) provides the network address of a security server upon request, (iii) the Cell Directory Service (CDS) stores and provides a "junction" with binding information for the registry service that manages the security namespace; and (iv) all communications with the registry service are authenticated. Thus, Lendenmann's security-specific names and junction to the registry service each teach that a "domain name[] includes an indication that the domain name service system supports the establishment of a secure communication link."

The Examiner's rejection was proper and should be made final.

**f. Dependent Claims 2, 3, 6, 14, 15, 18-22, 25, 26, 28-30, 34, 35, 37-39, 42-46, 49, 50, 52-54, 58 & 59**

The Patent Owner does not raise any argument specific to claims 2, 3, 6, 14, 15, 18-22, 25, 26, 28-30, 34, 35, 37-39, 42-46, 49, 50, 52-54, 58 and 59, thereby conceding that Lendenmann teaches the additional limitations recited in these claims. Because these claims' respective parent claims are anticipated by Lendenmann, as shown above, the rejection of these claims is also proper and should be made final.

**3. Response to Patent Owner's Arguments Regarding Rejection of Claims 1-3, 5, 6, 14-30, 33-54, and 57-60 Under 35 U.S.C. § 103(a) Based on Lendenmann (Issue 2)**

The Examiner properly rejected claims 1-3, 5, 6, 14-30, 33-54, and 57-60 because Lendenmann renders obvious the limitations in these claims. Although the Patent Owner does not point to any specific limitation regarding the obviousness rejections, in refuting the



New Bay Capital, LLC
Ex.1015-Page 2693 of 3151

corresponding anticipation rejections the Patent Owner repeatedly asserts that Lendenmann describes a "separate Security Service component that handles the aspects of any desired security measures." (Response at 8.) In short, the Patent Owner effectively concedes that Lendenmann teaches all of the claim limitations, but argues that the claimed domain name service system requires functionality that, in Lendenmann, is provided by two servers.

As explained above in the Requester's Comments regarding anticipation, Lendenmann teaches that the Cell Directory Service (CDS) is "integrated into the security service" and includes "ACL [access control list] management software, incorporated into all CDS clerks and servers, [that] performs access checking for incoming requests." (Lendenmann at 34.) "The security namespace is managed by the registry service of the DCE security component." (Lendenmann at 27.) Thus, Lendenmann provides a disclosure that integrates the functionality of the Cell Directory Service (CDS) and the Security Service. To the extent that Lendenmann does not anticipate the claims, Lendenmann provides an express teaching and suggestion to combine the functionalities of its name and security servers. In view of this teaching, it would have been obvious to one of ordinary skill in the art to combine the functionality of Lendenmann's Cell Directory Service (CDS) and the Security Service, for example, by installing both services to a single server.

Lendenmann further provides an example of such an installation in Fig. 39, which is configured with "*ev1* as the primary security server, [and] the initial CDS server," and node ev2 is configured "to become a secondary CDS server, [and] a secondary security server." (Lendenmann at 105.) Thus, in this example, Lendenmann teaches that both ev1 and ev2 are servers that provide by name resolution through the Cell Directory Service (CDS) and security through the Security Service.

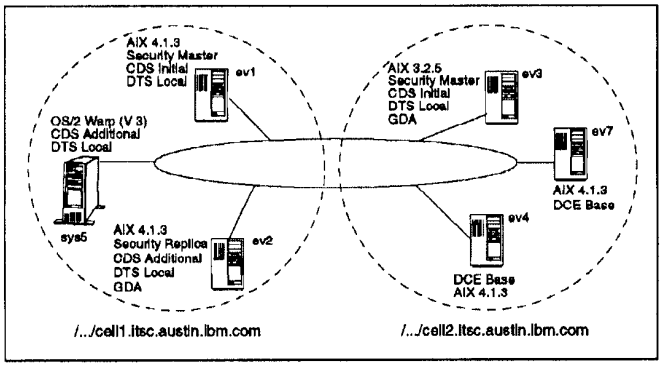


**LENDENMANN FIG. 39**



New Bay Capital, LLC
Ex.1015-Page 2694 of 3151

Accordingly, it would have been obvious to combine Lendenmann's teachings regarding the Cell Directory Service (CDS) and Security Service, thus rendering obvious claims 1-3, 5, 6, 14-30, 33-54, and 57-60. The Examiner's rejection is proper and should be made final.

#### 4. Response to Patent Owner's Arguments Regarding Rejection of Claim 7 Based on Lendenmann in view of Wesinger (Issue 3)

The Patent Owner does not contest Wesinger's teaching that a domain name service system may be on an edge router, as recited in claim 7. Accordingly, the rejection of claim 7 is proper and should be made final.

The Patent Owner argues, at pages 20-21, that Wesinger fails to teach the "indication" limitation of parent claim 1. However, "[o]ne cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) Here, the Examiner properly relied on Lendenmann—not Wesinger—to teach the "indication" limitation. The Patent Owner's argument is without merit.

#### 5. Response to Patent Owner's Arguments Regarding Rejection of Claims 8 & 9 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz (Issue 4)

The Patent Owner argues, at page 21, that Gaspoz fails to teach the "indication" limitation of parent claim 1. However, "[o]ne cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) Here, the Examiner properly relied on Lendenmann—not Gaspoz—to teach the "indication" limitation. The Patent Owner's argument is without merit.

##### a. Claim 8

The Examiner properly rejected claim 8 because Lendenmann and Gaspoz teach:

- The Cell Directory Service (CDS) server is part of a network cell in a Distributed Computing Environment. (Lendenmann at 1, 7, & Fig. 3.)
- Computers in separate network cells may communicate with each other over a communication network: "[A]ccess of the foreign cell is established over the Internet...." (Lendenmann at 23.)
- The Cell Directory Service (CDS) stores the information needed for intercell communication: "intercell communication can now be defined within CDS." (Lendenmann at 25; *see also id.* at 68.)
- A Virtual Private Network service can be implemented on a communication network



New Bay Capital, LLC
Ex.1015-Page 2695 of 3151

using the Distributed Computing Environment. (Gaspoz, Abstract.)

Accordingly, Lendenmann and Gaspoz teach that (i) the Cell Directory Service (CDS) server is connectable to a communications network, and (ii) the communications network may form, in part, a virtual private network. Thus, Lendenmann and Gaspoz render obvious that "the domain name service system is connectable to a virtual private network through the communication network" as recited in claim 8.

The Patent Owner argues, at pages 21-23, that "the Security Server—not the CDS—manages any security measure for subsequent [intercell] communications, including any authentication, authorization, or encryption." (Response at 21.) This argument fails because it is unrelated to the limitation of claim 8. The Patent Owner is attempting to insert additional limitations into the claim via attorney argument. The claim does not recite any limitation related to the Patent Owner's argument, e.g., "wherein the domain name service system manages the security measures for the virtual private network."

The Patent Owner also argues that "[s]imply because a CDS is a component of a cell does not mean that the entire cell is a 'domain name service system.'" (Response at 22.) With this argument, the Patent Owner improperly attempts to rewrite claim 8 to recite a "domain name service ~~system~~ server" through attorney argument. Lendenmann expressly discloses that the Cell Directory Service (CDS) is a required part of every cell, and the cells can securely communicate with each other through a communication network such as the Internet.

The Examiner's rejection is proper and should be made final.

**b. Claim 9**

The Patent Owner does not raise any argument specific to claim 9, thereby conceding that Lendenmann and Gaspoz render obvious the additional limitation of claim 9. Accordingly, the Examiner's rejection is proper and should be made final.

**6. Response to the Patent Owner's Argument Regarding the Rejection of Claim 10 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz and Schneider (Issue 5)**

Regarding claim 10, the Patent Owner argues that Schneider fails to teach the "indication" limitation of claim 1 and the "connectable to a virtual private network" limitation of claim 8. The Examiner's rejection was proper, however, because Lendenmann and Gaspoz—not Schneider—were relied upon to teach these limitations. "One cannot show nonobviousness by



New Bay Capital, LLC
Ex.1015-Page 2696 of 3151

attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) The Examiner's rejection is proper and should be made final.

### 7. Response to the Patent Owner's Argument Regarding the Rejection of Claim 11 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz and Martin (Issue 6)

The Examiner properly rejected claim 11 because the prior art teaches:

- When a client has multiple equivalent destination servers to choose from, it may select a destination server at random. (Lendenmann at 185.)
- A client can choose at random from multiple source addresses. (Martin at 9.)

Accordingly, the prior art teaches that when a client is initiating a new communication link to a server, the client can randomly choose both the source ("from") address and the destination ("to") address to be used. The addresses used can be different for each connection. Thus, the prior art teaches a "network address hopping regime that is used to pseudorandomly change network addresses in packets" as recited in the claim 11.

The Patent Owner argues, on pages 24-25, that Martin fails to teach the "indication" limitation of claim 1. However, "[o]ne cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) Here, the Examiner properly relied on Lendenmann—not Martin—to teach the "indication" limitation. The Patent Owner's argument is without merit.

The Patent Owner concludes, without explanation, that a regime for changing network addresses based on randomly selecting source and destination addresses, as taught by Lendenmann and Martin, is somehow different than "pseudorandomly chang[ing] network addresses in packets" as recited in claim 11. Randomly using different network addresses for each connection teaches that the source and destination network addresses in the packets transiting the virtual private network randomly change. To the extent that the Patent Owner believes that the claim has a different meaning, the Patent Owner has neither explained that meaning nor provided any reasoning for it. (See 37 CFR 1.111(b).)

The Examiner's rejection is proper and should be made final.



New Bay Capital, LLC
Ex.1015-Page 2697 of 3151

### 8. Response to the Patent Owner's Argument Regarding the Rejection of Claims 12 & 13 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Gaspoz and RFC 793

#### a. Claim 12

The Examiner properly rejected claim 12 because the prior art teaches:

- Secure communications can occur via Transmission Control Protocol (TCP). (Lendenmann at 179.)
- TCP requires that incoming packets include a sequence number that is within a "window" range of expected values. (RFC 793 at 24.)

Accordingly, the prior art renders obvious that (i) secure communication among nodes in a virtual private network can occur via TCP, and (ii) TCP requires comparing sequence values in each packet to a window of valid values. Thus, the prior art renders obvious that "the virtual private network is based on comparing a value in each data packet transmitted between a first device and a second device to a moving window of valid values" as recited in claim 12.

The Patent Owner argues, at pages 25-26, that RFC 793 does not disclose the "indication" limitation of claim 1 or the "connectable" limitation of claim 8. However, "[o]ne cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) Here, the Examiner properly relied on Lendenmann and Gaspoz—not RFC 793—to teach these limitations. The Patent Owner's argument is without merit.

The Examiner's rejection is proper and should be made final.

#### b. Claim 13

The Patent Owner does not argue for the separate patentability of claim 13, relying instead on its arguments regarding parent claims 1 and 8. The Patent Owner also argues that RFC 793 fails to disclose the "indication" limitation of claim 1, but "nonobviousness by attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) Here, the Examiner properly relied on Lendenmann—not RFC 793—to teach the limitations of claim 1. Since the Patent Owner's arguments are without merit, the rejection of claim 13 is proper and should be made final.



New Bay Capital, LLC
Ex.1015-Page 2698 of 3151

### 9. Response to the Patent Owner's Argument Regarding the Rejection of Claims 31, 32, 55, & 56 Under 35 U.S.C. § 103(a) Based on Lendenmann in view of Ludwig and RFC 793

The Patent Owner does not argue for the separate patentability of claims 31, 32, 55, and 56, relying instead on its arguments regarding parent claims 1 and 36. The Patent Owner also argues that Ludwig and RFC 793 fail to disclose the "indication" limitations of claims 1 and 36, but "nonobviousness by attacking references individually where the rejections are based on combinations of references." (MPEP § 2145 (IV).) Here, the Examiner properly relied on Lendenmann—not Ludwig or RFC 793—to teach the limitations of claims 1 and 36. Since the Patent Owner's arguments are without merit, the Examiner's rejection of claims 31, 32, 55, and 56 is proper and should be made final.

## B. Aziz

### 1. Overview of Aziz

U.S. Pat. No. 6,119,234, "Method and apparatus for client-host communication over a computer network," to Aziz, Jr., et al. ("Aziz") was filed on June 27, 1997, and issued September 12, 2000. As a patent issuing on an application filed before the '504 Patent's earliest claimed priority date of October 30, 1998, Aziz is prior art under 35 U.S.C. § 102(e).

Aziz describes a specialized Domain Name Server that is used to establish secure network links between computers over the Internet global network. (Aziz, 5:61-64.) To facilitate secure communications, Aziz describes configuring the Domain Name Server to respond to requests with special records that include information needed for secure communications. (Aziz, 4:9-16.) The Domain Name Server also includes such special records for itself, thus enabling itself to establish secure communication links. (Aziz, 8:66-9:2.)

### 2. Response to Patent Owner's Arguments Regarding the Rejection of Claims 1, 2, 5-8, 14-25, 27, 28, 33-52, and 57-60 Under 35 U.S.C. § 102(b) Based on Aziz (Issue 9)

The Examiner's rejections are proper and should be reaffirmed for the reasons below.

#### a. Independent Claim 1

The Examiner properly rejected claim 1 as anticipated by Aziz because Aziz teaches:

- A specialized domain name server that "is configured to return … an SX record, in response to requests for information needed for secure communications with protected hosts." (Aziz 4:9-13.)



New Bay Capital, LLC
Ex.1015-Page 2699 of 3151

- Configuring the name server includes a step of "defining an SX resource record type and adding appropriate records to the name server database for outside NS 120." (Aziz, 8:67-9:1.)
- "The data in the SX record is used ... to update information used by a client for secure communications with protected hosts." (Aziz, 6:57-60.)
- The domain name server uses "secure DNS" security extensions (Aziz, 5:61-64), which are well known to those of skill in the art from their definition in RFC 2065. (Aziz, 6:11-15.)
- The RFC 2065 secure DNS extensions define a bit that "indicates that this key is valid for use in conjunction with" the IPSEC secure communications standard. Therefore, the "key could be used in connection with secured communication." (RFC 2065, p. 12.)

Accordingly, Aziz teaches that the domain name server (i) stores and provides information that enables clients to establish secure communication links; and (ii) stores and provides keys that show that it supports establishing a secure communication link using the IPSEC secure communications standard.

The Patent Owner advances arguments, on pages 28-29, that would improperly rewrite claim 1 through attorney argument. Specifically, the Patent Owner asserts that the claimed "domain name service system" must be embodied in a single *server*, not in a "system" that may be comprised of multiple components. The Patent Owner provides no reasoning or justification, merely making the conclusory statement that "an expansive reading is not consistent with the plain and ordinary meaning, and the broadest reasonable interpretation, of the term 'domain name service system.'" (Response at 28-29.) While a patentee is entitled to act as its own lexicographer, the Patent Owner points to nothing in the specification to support its position that it defines the word "system" to mean "one computer" or "one server." To the contrary, the specification repeatedly uses the word "system" to refer to a collection of multiple components. For example, Figure 26 (showing at least 5 separate computers) is described as showing a "system employing a DNS proxy server with transparent VPN creation." ('504 Patent, 9:7-8.) Another example is Figure 33, which illustrates "a system block diagram 3300 of a computer network" ('504 Patent, 49:12-13) and includes numerous components. The system in Fig. 33 includes, among other things, computers 3301 and 3304, network 3302, website 3308, and secure domain name service 3313. Notably, the Patent Owner does not point to, and the Requester



New Bay Capital, LLC
Ex.1015-Page 2700 of 3151