

VIRTUAL DISPERSIVE NETWORKING SYSTEMS AND METHODS

CROSS-REFERENCE TO RELATED APPLICATION

[001] The present application is, for purposes of the U.S., a U.S. nonprovisional patent application of, and claims priority under 35 U.S.C. §119(e) to, U.S. provisional patent application 62/249,619, filed November 2, 2015, which provisional patent application is incorporated by reference herein.

COPYRIGHT STATEMENT

[002] All of the material in this patent document is subject to copyright protection under the copyright laws of the United States and other countries. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in official governmental records but, otherwise, all other copyright rights whatsoever are reserved.

BACKGROUND OF THE INVENTION

[003] The present invention generally relates to network routing and network communications.

[004] Conventional networks, such as the Internet, rely heavily on centralized routers to perform routing tasks in accomplishing network communications. The vulnerability and fragility of these conventional networks make entities feel insecure about using them.

[005] Further, network attacks are a common occurrence in today's cyber environment. Computers that are not protected by firewalls are under constant attack. Servers are especially vulnerable due to the fact that they must keep a port open on their firewall to enable a connection to a client. Hackers use this vulnerability to attack the server or computer with a Denial of Service (DoS) attack to deny use of the server.

[006] Cyber security is under siege from hackers. Since cyber security solutions can be bought, studied, and reverse engineered, hackers often have had a free hand attacking and stealing information from individuals, companies, and nations. One major issue is that, traditionally, the security industry creates "a single lock" for all of their customers and assumes that no one can break the lock. However, this paradigm is not working, as companies and nations have had information stolen. Recently, for example, Sony and the United States Office of Program Management have been the victim of hacks. Once a hacker figures out a technique for how to break a cyber defense tool such as a firewall, deep packet inspection, virus protection, intrusion detection system, IPS or others, the technique generally becomes part of their play book on how to hack.

[007] Further, when vendors develop red team tools that are used to test a company's defenses, these tools can simply be pointed at other companies and used by hackers to break encryption and hack firewalls and intrusion detection systems.

[008] Needs exist for improvement in cyber security. For example, needs exist for security mechanisms that are not algorithmic in nature and which can be modified in unique ways so that security can be achieved with simple software modifications. Needs exist for system configuration to

be leveraged to change operation functionality as well. One or more of these needs is addressed by one or more aspects of the present invention.

SUMMARY

[009] The present disclosure includes many aspects and embodiments. Moreover, while many aspects and embodiments relate to, and are described in, the context of network routing and network communications associated with the Internet, the present disclosure is not limited to use only in conjunction with the Internet and is applicable in other networked systems not associated with the Internet, as will become apparent from the following summaries and detailed descriptions of aspects, and one or more embodiments of the present invention.

[010] A first aspect relates to a method of communicating data which includes splitting into a first plurality of data streams, at endpoint software running on a first device, first data for communication to a destination device; selecting, at the first device by the endpoint software, a first plurality of deflects for use in communicating the first plurality of data streams; communicating each of the first plurality of data streams over a different one of the selected first plurality of deflects; splitting, at a first deflect, a first data stream of the first plurality of data streams into a second plurality of data streams; selecting, at the first deflect, a second plurality of deflects for use in communicating the second plurality of data streams; and communicating each of the second plurality of data streams over a different one of the selected second plurality of deflects;

receiving, at a second deflect from a first set of deflects, the second plurality of data streams; reassembling, at the second deflect, the second plurality of data streams into the first data stream, and communicating the first data stream onward to another device; receiving, at a second device from a second set of deflects, the first plurality of data streams including the reassembled first data stream; and reassembling, at the second device, the first plurality of data streams into the first data.

- [027] In an embodiment, each deflect represents deflect software loaded on a distinct physical computing device.
- [028] In an embodiment, two or more deflects represent deflect software loaded on the same physical computing device.
- [029] In an embodiment, the first device or the second device has both deflect software and endpoint software loaded thereon.
- [030] In an embodiment, one or more deflects represent deflect software loaded on a guest operating system running in a virtual machine.
- [031] In an embodiment, the method involves communication of a data stream between two virtual machines running on the same physical computing device.
- [032] In an embodiment, the method involves utilizing encryption for communications between nodes.
- [033] In an embodiment, the method involves utilizing node-to-node encryption for communications between nodes in combination with end-to-end encryption for the entire communication from the first device to the second device.
- [034] In an embodiment, application-level encryption is utilized.
- [035] In an embodiment, public-key encryption is utilized and a key is communicated utilizing a methodology involving splits inside of splits.
- [036] In an embodiment, interleaving is utilized inside of one or more splits.
- [037] In an embodiment, the insertion of red herring data is utilized within one or more splits.

[038] In an embodiment, selecting, at the first device by the endpoint software, the first plurality of deflects for use in communicating the first plurality of data streams comprises selecting the first plurality of deflects based at least in part on networking information.

[039] In an embodiment, endpoint software is loaded on a desktop computer.

[040] In an embodiment, endpoint software is loaded on a laptop computer.

[041] In an embodiment, endpoint software is loaded on a tablet.

[042] In an embodiment, endpoint software is loaded on a smartphone.

[043] In an embodiment, endpoint software is loaded on a smart watch.

[044] In an embodiment, endpoint software is loaded on a wearable computing device.

[045] In an embodiment, endpoint software is loaded on a smart appliance.

[046] In an embodiment, deflect software is loaded on a desktop computer.

[047] In an embodiment, deflect software is loaded on a laptop computer.

[048] In an embodiment, deflect software is loaded on a tablet.

[049] In an embodiment, deflect software is loaded on a smartphone.

[050] In an embodiment, deflect software is loaded on a smart watch.

[051] In an embodiment, deflect software is loaded on a wearable computing device.

[052] In an embodiment, deflect software is loaded on a smart appliance.

[053] Another embodiment relates to one or more computer readable media containing computer executable instructions for performing the disclosed method.

[054] Another embodiment relates to a system for performing the disclosed method.

[055] Another embodiment relates to software program for performing the disclosed method.

[056] In addition to the aspects and embodiments of the present disclosure, it should be noted that the present disclosure further encompasses the various possible combinations and subcombinations of such aspects and embodiments. Thus, for example, any aspect may be combined with an aforementioned embodiment in accordance with the present disclosure without requiring any other aspect or embodiment.

BRIEF DESCRIPTION OF THE DRAWINGS

[057] One or more preferred embodiments of the present invention now will be described in detail with reference to the accompanying drawings.

[058] FIG. 1 illustrates components of a VDR software client loaded onto a client device in accordance with an embodiment of the present invention.

[059] FIG. 2 illustrates how a VDR client gathers LAN routing information and queries an external network for backbone information and application-specific routing information in accordance with an embodiment of the present invention.

[060] FIG. 3 illustrates how data is added to the payload of a packet on each of a plurality of hops in accordance with an embodiment of the present invention.

[061] FIGS. 4A-C provide a simplified example of a VDR software response to a network attack in accordance with an embodiment of the present invention.

- [062] FIGS. 5A-C illustrate an exemplary VDR implementation in accordance with a preferred embodiment of the present invention.
- [063] FIG. 6 includes Table 1, which table details data stored by a node in the payload of a packet.
- [064] FIG. 7 illustrates a direct connection between two clients in accordance with one or more preferred implementations.
- [065] FIG. 8 illustrates an exemplary process for direct transfer of a file from a first client to a second client in accordance with one or more preferred implementations.
- [066] FIG. 9B illustrates an exemplary user interface for a Sharzing file transfer application in accordance with one or more preferred implementations.
- [067] FIG. 10 presents table 9, which illustrates potential resource reduction in accordance with one or more preferred implementations.
- [068] FIG. 11 illustrates client and server architectures in accordance with one or more preferred implementations.
- [069] FIGS. 12 and 13 illustrate exemplary processes for downloading of a file in accordance with one or more preferred implementations.
- [070] FIG. 14 illustrates the use of a dispersive virtual machine implemented as part of a software application that can be easily downloaded to a device such as a PC, smart phone, tablet, or server.
- [071] FIG. 15 illustrates a methodology in which data to be sent from a first device to another device is split up into multiple parts which are sent separately over different routes and then reassembled at the other device.
- [072] FIG. 16 illustrates how multiple packets can be sent over different deflects in a direct spreading of packets methodology.
- [073] FIG. 17 illustrates how multiple packets can be sent to different IP addresses and/or ports in a hopping IP addresses and ports methodology.
- [074] FIG. 18 illustrates an exemplary system architecture configured to allow clients in a task network to access the Internet through an interface server using virtual dispersive networking (VDN) spread spectrum protocols.
- [075] FIG. 19 illustrates an exemplary system architecture configured to enable a workstation with four independent connections to the internet to send traffic using virtual dispersive networking spread spectrum protocols to an interface server from each independent internet connection.
- [076] FIGS. 20-23 illustrate an exemplary scenario utilizing point of entry gateways.
- [077] FIGS. 24-27 illustrate a similar scenario as that illustrated in FIGS. 20-23, only one or more deflects utilized for some communications or connections.
- [078] FIG. 28 illustrates the overlapping of data by sending it to a mobile device via two different wireless networks.
- [079] FIG. 29 illustrates VDN routing from a first virtual thin client to another virtual thin client using two deflects that are configured simply to pass data through.

[080] FIG. 30 illustrates a system in which deflects are configured to reformat data to another protocol.

[081] FIG. 31A illustrates a first packet which includes data A, as well as a second packet which includes data B.

[082] FIG. 31B illustrates a superframe which has been constructed that includes both data A and data B.

[083] FIGS. 32-33 illustrate a process in which a superframe is sent from a first virtual thin client to a deflect.

[084] FIG. 34 illustrates a superframe which itself includes a superframe (which includes two packets) as well as a packet.

[085] FIG. 35 illustrates how communications are intercepted at the session layer and parsed out to a plurality of links.

[086] FIG. 36 illustrates a plurality of data paths from a first end device to a second end device.

[087] FIG. 37 illustrates several data paths that pass through multiple deflects.

[088] FIG. 38 illustrates how a connection from a first VTC to a second VTC through a second deflect is still possible even though it is not possible to connect through a first or third deflect.

[089] FIG. 39 illustrates a system including two end clients that are communicating over a network that includes three deflects.

[090] FIGS. 40-47 illustrate exemplary systems utilizing splitting of communications.

DETAILED DESCRIPTION

[091] As a preliminary matter, it will readily be understood by one having ordinary skill in the relevant art (“Ordinary Artisan”) that the present invention has broad utility and application. Furthermore, any embodiment discussed and identified as being “preferred” is considered to be part of a best mode contemplated for carrying out the present invention. Other embodiments also may be discussed for additional illustrative purposes in providing a full and enabling disclosure of the present invention. Moreover, many embodiments, such as adaptations, variations, modifications, and equivalent arrangements, will be implicitly disclosed by the embodiments described herein and fall within the scope of the present invention.

[092] Accordingly, while the present invention is described herein in detail in relation to one or more embodiments, it is to be understood that this disclosure is illustrative and exemplary of the present invention, and is made merely for the purposes of providing a full and enabling disclosure of the present invention. The detailed disclosure herein of one or more embodiments is not intended, nor is to be construed, to limit the scope of patent protection afforded the present invention, which scope is to be defined by the claims and the equivalents thereof. It is not intended that the scope of patent protection afforded the present invention be defined by reading into any claim a limitation found herein that does not explicitly appear in the claim itself.

[093] Thus, for example, any sequence(s) and/or temporal order of steps of various processes or methods that are described herein are illustrative and not restrictive. Accordingly, it should be understood that, although steps of various processes or methods may be shown and described as being in a sequence or temporal order, the steps of any such processes or methods are not limited to being carried out in any particular sequence or order, absent an indication otherwise. Indeed, the steps in such processes or methods generally may be carried out in various different sequences and orders while still falling within the scope of the present invention. Accordingly, it is intended that the scope of patent protection afforded the present invention is to be defined by the appended claims rather than the description set forth herein.

[094] Additionally, it is important to note that each term used herein refers to that which the Ordinary Artisan would understand such term to mean based on the contextual use of such term herein. To the extent that the meaning of a term used herein—as understood by the Ordinary Artisan based on the contextual use of such term—differs in any way from any particular dictionary definition of such term, it is intended that the meaning of the term as understood by the Ordinary Artisan should prevail.

[095] Furthermore, it is important to note that, as used herein, “a” and “an” each generally denotes “at least one,” but does not exclude a plurality unless the contextual use dictates otherwise. Thus, reference to “a picnic basket having an apple” describes “a picnic basket having at least one apple” as well as “a picnic basket having apples.” In contrast, reference to “a picnic basket having a single apple” describes “a picnic basket having only one apple.”

[096] When used herein to join a list of items, “or” denotes “at least one of the items,” but does not exclude a plurality of items of the list. Thus, reference to “a picnic basket having cheese or crackers” describes “a picnic basket having cheese without crackers”, “a picnic basket having crackers without cheese”, and “a picnic basket having both cheese and crackers.” Finally, when used herein to join a list of items, “and” denotes “all of the items of the list.” Thus, reference to “a picnic basket having cheese and crackers” describes “a picnic basket having cheese, wherein the picnic basket further has crackers,” as well as describes “a picnic basket having crackers, wherein the picnic basket further has cheese.”

[097] Further, as used herein, the term server may be utilized to refer to both a single server, or a plurality of servers working together.

[098] Referring now to the drawings, one or more preferred embodiments of the present invention are next described. The following description of one or more preferred embodiments is merely exemplary in nature and is in no way intended to limit the invention, its implementations, or uses.

VDR

[099] Virtual dispersive routing (hereinafter, “VDR”) relates generally to providing routing capabilities at a plurality of client devices using virtualization. Whereas traditional routing calls for most, if not all, routing functionality to be carried out by centrally located specialized routing devices, VDR enables dispersed client devices to assist with, or even takeover, routing functionality, and thus

is properly characterized as *dispersive*. Advantageously, because routing is performed locally at a client device, a routing protocol is selected by the client based upon connection requirements of the local application initiating the connection. A protocol can be selected for multiple such connections and multiple routing protocols can even be utilized simultaneously. The fragile nature of the routing protocols will be appreciated, and thus virtualization is utilized together with the localization of routing to provide a much more robust system. Consequently, such dispersive routing is properly characterized as *virtual*.

[0100] More specifically, preferred VDR implementations require that a VDR software client be loaded on each client device to help control and optimize network communications and performance. Preferably, VDR is implemented exclusively as software and does not include any hardware components. Preferably, the basic components of a VDR software client include a routing platform (hereinafter, “RP”); a virtual machine monitor (hereinafter, “VMM”); a dispersive controller (hereinafter, “DC”); and an application interface (hereinafter, “AI”). FIG. 1 illustrates each of these components loaded onto a client device. Each of these components is now discussed in turn.

The Routing Platform (RP) and Multiple Routing Protocols

[0101] Despite eschewing the traditional routing model utilizing central points of control, VDR is designed to function with existing routing protocols. Supported routing protocols, together with software necessary for their use, are included in the routing platform component of the VDR software, which can be seen in FIG. 1. For example, the RP includes software to implement and support the Interior Gateway Routing Protocol (“IGRP”), the Enhanced Interior Gateway Routing Protocol (“EIGRP”), the Border Gateway Protocol (“BGP”), the Open Shortest Path First (“OSPF”) protocol, and the Constrained Shortest Path First (“CSPF”) protocol. It will be appreciated that in at least some embodiments, a port will be needed to allow conventional routing software to run on a chip core (for example, a core of an Intel chip) at a client device. Preferably, multi-core components are used to allow routing protocols to be run on multiple cores to improve overall performance.

[0102] Moreover, it will be appreciated that the ability to support multiple routing protocols allows VDR to meet the needs of applications having varying mobility requirements. Applications can be supported by *ad hoc* algorithms such as pro-active (table driven) routing, reactive (on-demand) routing, flow oriented routing, adaptive (situation aware) routing, hybrid (pro-active/reactive) routing, hierarchical routing, geographical routing, and power aware routing. Further, the use of multiple protocols supports broadcasting, multi-casting, and simul-casting. It will be appreciated that the use of multiple protocols provides support for multi-threaded networking as well.

The Virtual Machine Monitor (VMM) and Virtualization

[0103] It will be appreciated that virtualization is known in some computing contexts, such as virtualization of memory and processing. Virtualization enables the abstraction of computer resources and can make a single physical resource appear, and function, as multiple logical resources. Traditionally, this capability enables developers to abstract development of an application so that it

runs homogenously across many hardware platforms. More generally, virtualization is geared to hiding technical detail through encapsulation. This encapsulation provides the mechanism to support complex networking and improved security that is required to enable routing at client devices.

[0104] More specifically, a virtual machine (hereinafter, “VM”) is a software copy of a real machine interface. The purpose of running a VM is to provide an environment that enables a computer to isolate and control access to its services. The virtual machine monitor (VMM) component is used to run a plurality of VMs on a real machine and interface directly with that real machine. As an example, consider a VMM on a real machine that creates and runs a plurality of VMs. A different operating system is then loaded onto each VM. Each VM provides a virtual interface that would appear to each operating system to be a real machine. The VMM runs the plurality of VMs and interfaces with the real machine.

[0105] In a VDR implementation, a VMM is utilized to create a VM for each distinct connection. It is helpful to explain at this juncture that what comprises a connection can vary, but in general includes a transfer of data in the form of packets from a first end device to a second end device along a path (or route). It will be appreciated that a single application can require multiple connections, for example, an application may require multiple connections because of bandwidth application requirements and performance requirements; in this event each connection preferably interfaces with its own VM and each connection can utilize (sometimes referred to as being tied to) the same routing protocol or different routing protocols, even though the connections are themselves necessitated by the same application. Similarly, although two connections may at times travel along an identical path, the connections themselves are nevertheless distinct, and each will preferably still continue to interface with its own VM.

The Dispersive Controller (DC) and Optimizing Performance

[0106] When the client is in need of a new connection, a dispersive controller located between an operating system and a driver that controls network hardware (such as a NIC card) intercepts the request for a new connection and tells the VMM to spawn a new VM associated with the desired connection. The DC then queries the application interface and utilizes any information obtained to select a routing protocol from among those supported by the RP. This selected routing protocol, however, is currently believed to be generally useless without knowledge of the surrounding network. To this end, the DC allows each client to find other clients, interrogate network devices, and utilize system resources. Thus, each VDR client is “network aware”, in that routing information is gathered and maintained at each client by the DC.

[0107] FIG. 2 illustrates how a VDR client **201** gathers LAN routing information and queries an external network for backbone information and application-specific routing information. In response to these queries, routing information is returned. This returned routing information is cached, processed, data mined, compared to historical data, and used to calculate performance metrics to

gauge and determine the overall effectiveness of the network. This is possible because the resources available at a VDR client will typically be greater than those available at a conventional router.

[0108] In at least some embodiments, a VDR network functions in some ways similarly to a conventional network. In a conventional network, data, in the form of packets, is sent to a router to be routed according to a routing table maintained at the router. Similarly, in a VDR network, after utilizing gathered network information to generate a routing table, a client device utilizes this generated routing table to select a route and transmit a packet accordingly, which packet is then received by another client device and routed according to that client's routing table, and so on, until the packet reaches its destination.

[0109] However, rather than simply passing on received packets from client to client, in a manner akin to a traditional router, VDR, via the DC, instead takes advantage of the storage and processing resources available at each client, while still remaining compatible with existing network architecture, by attaching lower level protocol data to the payload of transmitted packets for subsequent client analysis.

[0110] More specifically, when a packet is received at a VDR client, a virtual machine intercepts the packet passed from the networking hardware (for example, a NIC card) and places it in memory. The VDR client then processes the packet data. When the data is subsequently passed on, this processed data is appended to the payload of the packet together with information relating to the VDR client for analysis at the destination. As can be seen in FIG. 3, the result of this process is that each hop causes additional information to be added to the payload of a packet, and thus results in a direct increase in payload size proportionate to the number of hops taken by the packet. Specifically, each hop is believed to result in an increase of 35 bytes for an IPv4 implementation, and 59 bytes for an IPv6 implementation. Table 1 of FIG. 6 details the information stored from each layer, along with the number of bytes allotted for each field. It will be appreciated that different or additional information could be stored in alternative embodiments.

[0111] Currently, 128-bit addressing provides support for IPv4 and IPv6 addressing, but support for additional addressing schemes is contemplated. It will be appreciated that for a typical communication over the Internet, i.e., one consisting of around 20 hops, the overhead appended to the payload will be around 700 bytes utilizing IPv4 and around 1180 bytes utilizing IPv6. It is believed that, in a worst case scenario, an extra IP datagram could be required for every datagram sent. Although some of this data may seem redundant at first blush, some repetition is tolerable and even necessary because network address translation ("NAT") can change source or destination fields. That being said, it is contemplated that some implementations use caching to lower this overhead. Additionally, in at least some implementations, the VDR client utilizes application specific knowledge to tailor the information that is appended to the needs of a specific application.

[0112] Conventionally, when a packet is received at a router, routing information is typically stripped off each packet by the router and disregarded. This is because each router has limited memory and

handles an enormous number of packets. When a packet is received at a destination VDR client, however, the destination client has sufficient resources to store and process the information delivered to it. Additionally, to the extent that client resources may be taxed, the VDR client need not always store this information in every packet received, as in at least some embodiments application knowledge provides the client with an understanding of which packets are important to applications running on the client. Regardless of whether some or all of this information delivered in the payload of each data packet is processed, the information that is processed is analyzed to create a “network fingerprint” of the nodes involved in the communication link. Thus, VDR software loaded on nodes along a path enables the nodes to append information regarding a path of a packet, which in turn enables the generation of a network fingerprint at the destination device, which network fingerprint represents a historical record that is stored and maintained for later forensic analysis. In addition to forensic analysis by the client, the maintenance of network information on the client enables forensic analysis by a server as well.

The Application Interface (AI) & Application Knowledge

[0113] One of the benefits of providing routing functionality at a client device is that the client is able to utilize its knowledge of the application initiating a connection to enhance routing performance for that application. This knowledge is provided to the DC via an application interface, as can be seen in FIG. 1. Utilizing application knowledge to enhance routing performance could be useful to a variety of applications, such, as for example, computer games including massively multiplayer online role playing games.

[0114] The virtualization of routing functionality at a client device, as described hereinabove, allows multiple routing protocols and algorithms to be run simultaneously on a client device. Thus, the DC utilizes the application interface to obtain required criteria for an application connection and then chooses from among the protocols and algorithms available via the RP.

[0115] For example, Application “A” may need to communicate very large amounts of data, and thus require a routing protocol that optimizes bandwidth, while Application “B” may only need to communicate very small amounts of data at very fast speeds, and thus require a routing protocol that minimizes latency irrespective of bandwidth. A traditional router cannot tell the difference between packets originating from Application “A” and those originating from Application “B”, and thus will utilize the same routing protocol for packets from each application. A VDR client, however, is aware of applications running locally, and thus can be aware, through the AI, of various connection criteria for each application. These connection criteria can then be utilized by the VDR client in selecting a routing protocol or algorithm. Furthermore, as described hereinabove, both the selected routing protocol and the originating application associated with a packet can be communicated to other client nodes via data appended to the payload of the packet. Thus, the protocol selected at a source client can be utilized to route the packet throughout its path to a destination client. Further, because

virtualization allows multiple routing protocols to be run on a single client, each application can utilize its own routing protocol.

[0116] Moreover, a VDR client can utilize knowledge of the path of a specific connection to further optimize performance. Because a network fingerprint can be gathered detailing the nodes in a communication path, a VDR client running on a client device can analyze each network fingerprint to determine whether the associated connection satisfies the connection criteria of the application desiring to utilize the connection. If the connection does not satisfy the connection criteria, then the client can attempt to find a connection that does satisfy the criteria by switching to a different protocol and/or switching to a different first node in its routing table. Combinations utilizing various protocols and selecting a variety of first nodes can be attempted, and the resultant paths evaluated until a path is found that does satisfy connection criteria. Additionally, combinations utilizing various protocols and selecting a variety of first nodes can be utilized to create route redundancy. Such route redundancy can provide to an application both higher bandwidth and controllable quality of service.

[0117] Although connection criteria for source and destination clients will often be identical, there are many situations where this will not be the case. For example, if one client is downloading streaming video from another client, then the connection requirements for each client will likely not be identical. In this and other situations, connections between two clients may be asymmetrical, i.e., client “A” transmits packets to client “B” over path 1, but client “B” transmits packets to client “A” over path 2. In each case, because path information gleaned from the payload of packets is stored and processed at the destination client, the evaluation of whether the path meets the required connection criteria is made at the destination client. In the example above, client “B” would determine whether path 1 satisfies its application’s connection criteria, while client “A” would determine whether path 2 satisfies its application’s connection criteria.

[0118] Perhaps the epitome of a connection that does not satisfy connection criteria is a broken, or failed, connection. In the event of a connection break, VDR enjoys a significant advantage over more traditional routing. Conventionally, recognition of a connection break would require a timeout at an upper level application, with either the path being re-routed subsequent to the timeout or a connection failure message being presented to a user. A VDR client, however, is aware of generally how long it should take to receive a response to a transmitted communication, and can utilize this awareness to speed up route convergence for additional network connections to insure application robustness and performance requirements, performance requirements being defined as criteria that must be met to allow the application to run properly, i.e., video conferencing can’t wait too long for packets to show up or else the audio “crackles” and the image “freezes.” For example, a VDR client may be aware that it should receive a response to a communication in 500 ms. If a response has not been received after 500 ms, the VDR client can initiate a new connection utilizing a different routing protocol and/or first node as outlined above with respect to finding a satisfactory connection path.

[0119] In addition to performance optimization, application knowledge can also be utilized to enhance network security. For example, an application may have certain security requirements. A VDR client aware of these requirements can create a “trusted network” connection that can be used to transfer information securely over this connection in accordance with the requirements of the application. A more traditional routing scheme could not ensure such a trusted connection, as it could not differentiate between packets needing this secure connection and other packets to be routed in a conventional manner.

[0120] But before elaborating on security measures that may be built in to a VDR implementation, it is worth noting that a VDR client is able to work in concert with an existing client firewall to protect software and hardware resources. It will be appreciated that conventional firewalls protect the flow of data into and out of a client and defend against hacking and data corruption. Preferably, VDR software interfaces with any existing client firewall for ease of integration with existing systems, but it is contemplated that in some implementations VDR software can include its own firewall. In either implementation, the VDR software can interface with the firewall to open and close ports as necessary, thereby controlling the flow of data in and out.

[0121] In addition to this firewall security, by utilizing application knowledge the VDR software can filter and control packets relative to applications running on the client. Thus, packets are checked not only to ensure a correct destination address, but further are checked to ensure that they belong to a valid client application.

[0122] One way VDR software can accomplish this is by utilizing “spiders” to thread together different layers of the protocol stack to enable data communication, thereby reducing delays and taking advantage of network topologies. Each spider represents software that is used to analyze data from different layers of the software stack and make decisions. These threaded connections can be used to speed data transfer in static configurations and modify data transfer in dynamic circumstances. As an example, consider a client device running a secure email application which includes a security identification code. Packets for this application include a checksum that when run will come up with this identification code. A spider would allow this upper level application security identification code to be connected to the lower layer. Thus, the lower layer could run a checksum on incoming packets and discard those that do not produce the identification code. It will be appreciated that a more complex MD5 hash algorithm could be utilized as well.

[0123] Moreover, because the VDR software is knowledgeable of the application requiring a particular connection, the software can adaptively learn and identify atypical behavior from an outside network and react by quarantining an incoming data stream until it can be verified. This ability to match incoming data against application needs and isolate any potential security issues significantly undermines the ability of a hacker to gain access to client resources.

[0124] Additionally, when such a security issue is identified, a VDR client can take appropriate steps to ensure that it does not compromise the network. Because a VDR client is network aware and keeps

track of other clients that it has been communicating with, when a security issue is identified, the VDR client can not only isolate the suspect connection, the VDR client can further initiate a new connection utilizing a different routing protocol and/or first node as outlined above with respect to finding a satisfactory connection path. Alternatively, or additionally, the VDR client could simply choose to switch protocols on the fly and communicate this switch to each client with which it is in communication.

[0125] FIGS. 4A-C provide a simplified example of such action for illustrative effect. In FIG. 4A, VDR client **403** is communicating with VDR client **405** over connection **440**. In FIG. 4B, external computer **411** tries to alter packet **491** transmitted from client **403** to client **405**. Client **405** runs a hashing algorithm on the received packet **491** and identifies that it has been corrupted. Client **405** then quarantines packets received via connection **440** and, as can be seen in FIG. 4C, establishes a new connection **450** with client **403**.

[0126] Upon discovery of an “attack” on a network or specific network connection, a VDR client can monitor the attack, defend against the attack, and/or attack the “hacker”. Almost certainly, a new, secure connection will be established as described above. However, after establishing a new connection, the VDR client can then choose to simply kill the old connection, or, alternatively, leave the old connection up so that the attacker will continue to think the attack has some chance of success. Because each connection is virtualized, as described hereinabove, a successful attack on any single connection will not spill over and compromise the client as a whole, as crashing the VM associated with a single connection would not affect other VMs or the client device itself. It is contemplated that a VDR client will attempt to trace back the attack and attack the original attacker, or alternatively, and preferably, communicate its situation to another VDR client configured to do so.

An Exemplary Implementation

[0127] Traditionally, wired and wireless networks have tended to be separate and distinct. Recently, however, these types of networks have begun to merge, with the result being that the routing of data around networks has become much more complex. Further, users utilizing such a merged network desire a high level of performance from the network regardless of whether they are connected wirelessly or are connected via a fixed line. As discussed hereinabove, VDR enables a client to monitor routing information and choose an appropriate routing protocol to achieve the desired performance while still remaining compatible with existing network architecture. VDR can be implemented with wired networks, wireless networks (including, for example, Wi-Fi), and networks having both wired and wireless portions.

[0128] FIG. 5A illustrates an exemplary local area network **510** (hereinafter, “LAN”) utilizing VDR. The LAN **510** includes three internal nodes **511,513,515**, each having VDR software loaded onto a client of the respective node. The internal nodes **511,513,515** can communicate with one another, and further can communicate with edge nodes **512,514,516,518**, each also having VDR software loaded onto a client of the respective node. The coverage area **519** of the LAN **510** is represented by a dotted

circle. It will be appreciated that the edge nodes **512,514,516,518** are located at the periphery of the coverage area **519**. The primary distinction between the internal nodes **511,513,515** and the edge nodes **512,514,516,518** is that the internal nodes **511,513,515** are adapted only to communicate over the LAN **510**, while the edge nodes **512,514,516,518** are adapted to communicate both with the internal nodes **511,513,515** and with edge nodes of other LANs through one or more wide area networks (hereinafter, “WANs”). As one of the nodes **511,513,515** moves within the LAN **510** (or, if properly adapted, moves to another LAN or WAN), VDR allows it to shift to *ad hoc*, interior, and exterior protocols. This ability to shift protocols allows the node to select a protocol which will provide the best performance for a specific application.

[0129] FIG. 5B illustrates an exemplary path between node **513** in LAN **510** and node **533** in LAN **530**. It will be appreciated that an “interior” protocol is utilized for communications inside each LAN, and an “exterior” protocol is utilized for communications between edge nodes of different LANs. Thus, it will likewise be appreciated that each edge node must utilize multiple protocols, an interior protocol to communicate with interior nodes, and an exterior protocol to communicate with other edge nodes of different LANs. Further, at any time an *ad hoc* protocol could be set up which is neither a standard interior nor exterior protocol.

[0130] In FIG. 5B, LAN **510** and LAN **530** are both using CSPF as an interior protocol, while LAN **520** and LAN **540** are utilizing EIGRP as an interior protocol. All edge nodes of each of the LANs **510,520,530** are connected to a WAN utilizing BGP to communicate between edge nodes.

[0131] The exemplary path between node **513** and node **533** includes node **515**, edge node **518**, edge node **522**, node **521**, node **523**, node **525**, edge node **528**, edge node **534**, and node **531**. Further, because a particular protocol was not selected and propagated by the transmitting node, this connection utilizes CSPF for internal communications within LAN **510** and LAN **530**, EIGRP for internal communications within LAN **520**, and BGP for external communications between edge nodes. At one or both end nodes, the VDR software can analyze this information and determine whether the combination of protocols along this path is satisfactory for the communicating application. It will be appreciated that the VDR software can further analyze the information gathered and determine whether the path meets application requirements for throughput, timing, security, and other important criteria.

[0132] In a static environment, this path may represent a connection that meets application requirements and thus no further adjustment would be needed. However, if a network outage were to occur, a network or a node were to move, or another dynamic event was to occur, the path could need to be altered.

[0133] For example, if LAN **520** were to move out of range, node **533** might analyze the path information appended to a packet received after the movement and determine that increased latency resulting from this movement rendered this path unsuitable per application requirements. Node **533** would then attempt to establish a new connection utilizing a different route that would satisfy

application requirements. FIG. 5C illustrates such a new connection, which remains between node 513 and node 533, but rather than being routed through LAN 520 as with the path illustrated in FIG. 5B, the path is instead routed through LAN 540.

[0134] It will be appreciated that the ability to influence path selection based on client application needs significantly enhances the performance, flexibility, and security of the network.

[0135] It will further be appreciated from the above description that one or more aspects of the present invention are contemplated for use with end, client, or end-client devices. A personal or laptop computer are examples of such a device, but a mobile communications device, such as a mobile phone, or a video game console are also examples of such a device. Still further, it will be appreciated that one or more aspects of the present invention are contemplated for use with financial transactions, as the increased security that can be provided by VDR is advantageous to these transactions.

Network Data Transfer

[0136] It will be appreciated that the transmission of data over the Internet, or one or more similar networks, often utilizes precious server processing, memory, and bandwidth, as the data is often delivered from, or processed at, a server. In implementations in accordance with one or more preferred embodiments of the present invention, some of this server load is mitigated by use of a direct connection between two end-user devices, such as, for example two end-user devices having virtualized routing capabilities as described hereinabove. Preferably, packets are then routed between the two end-user devices without passing through a conventional server.

[0137] Notably, however, although transferred data packets do not pass through a server, a server may still be utilized to establish, monitor, and control a connection, as illustrated in FIG. 7. Specifically, FIG. 7 illustrates two clients and an IP server which determines that the clients are authorized to communicate with one another, and which passes connection information to the clients that is utilized to establish a direct connection between the clients. Importantly, the IP server is not involved in this direct connection, i.e. data transferred via this direct connection is not routed through or processed by the IP server, which would require the use of additional resources of the IP server.

[0138] It will be appreciated that, in some networks, a firewall may be setup to prevent an end-user device from accepting connections from incoming requests. There are three basic scenarios that can occur. In a first case, there is no firewall obstruction. In the first case, either client can initiate the connection for the direct connect. In a second case, a single client has a firewall obstructing the connection. In this case, the client that is obstructed from accepting the connection is instructed by the IP Server to initiate the connection to the client that is unobstructed by the firewall. In a third case, both clients have firewalls obstructing the connection. In this case, a software router, or software switch, is used to accept the connection of the two clients and pass the packets through to the clients directly. Notably, this software router truly acts as a switch, and does not modify the payload as it passes the packet through. In a preferred implementation, a software router is implemented utilizing

field programmable gate arrays (FPGAs) or other specific hardware designed to implement such cross-connect functionality.

[0139] A preferred system for such a described direct connection includes one or more end-user devices having client software loaded thereon, an IP server, or control server, having server software loaded thereon, and one or more networks (such as, for example Internet, Intranet or Extranet supported by Ethernet, Mobile Phone data networks, e.g. CDMA, WiMAX, GSM, WCDMA and others, wireless networks, e.g. Bluetooth, WiFi, and other wireless data networks) for communication.

[0140] In a preferred implementation, client software installed at an end-user device is configured to communicate with an IP server, which associates, for example in a database, the IP address of the end-user device with a unique identification of a user, such as an email address, telephone number, or other unique identification. The client then periodically “checks in” with the IP server and conveys its IP address to the server, for example by providing its IP address together with the unique identification of the user. This checking in occurs when the client is “turned on”, e.g. when the end-user device is turned on or when the client software is loaded, as well as when the IP address has changed, or upon the occurrence of any other network event that would change the path between the client and server, or in accordance with other configured or determined events, times, or timelines, which may be user-configurable.

[0141] By collecting, and updating, the current IP address of a user, other users may communicate with that user as the user moves from place to place. The IP server thus acts as a registry providing updated IP addresses associated with users. This capability also enables multiple device delivery of content to multiple end-user devices a user designates or owns.

[0142] Preferably, such functionality is utilized in combination with virtualized routing capability as described hereinabove. Specifically, it will be appreciated that, currently, Internet communications utilize sessions, and that upon being dropped, e.g. due to a lost connection, a new session must be initialized.

[0143] In a preferred implementation, however, rather than having to re-initiate a new session, for example upon obtaining a new IP address, a new session is created and data is transferred from the old session to the new session while maintaining the state of the old session. In this way, a near-seamless transition is presented to a user between an old session and a new session. For example, a user might be connected via their mobile device to a Wi-Fi connection while they are on the move. They might move out of range of the Wi-Fi connection, but still be in range of a cellular connection. Rather than simply dropping their session, a new session is preferably created, and data from the old session copied over, together with the state of the old session. In this way, although the end-user device is now connected via a cellular connection, rather than via a Wi-Fi connection, the user’s experience was not interrupted.

One Client to One Client – File Transfer Implementation

[0144] In a preferred implementation, direct connections between end-user devices having virtualized routing capabilities are utilized in a file transfer context, such as, for example, with a file sharing application.

[0145] FIG. 8 illustrates an exemplary file transfer use scenario between two clients. As described above, each client is in communication with an IP server, for example to communicate its IP address to the IP server. Such communications are exemplified by steps **1010** and **1020**.

[0146] In use, a first client communicates to an IP server a request to connect to a particular client, user, or end-user device at step **1030**. The IP server, or control server, determines whether or not the other client, user, or end-user device is available, e.g. online, and, if so, looks up the current IP address or addresses associated with the specified client, user, or end-user device. If the client, user, or end-user device is either not online or has left the network, a connection failure message is sent. If the client, user, or end-user device is online, the IP server will take action based upon a pre-selected preference setting. Preferably, each user may choose to accept connection requests automatically, require a confirmation of acceptance, or require some other authentication information, such as an authentication certificate, in order to accept a connection request. If the connection request is accepted, either automatically or manually, the IP server enables the transfer, e.g. by communicating to a second client that the first client has a file for transfer, as exemplified by step **1040**.

[0147] Preferably, the IP server notifies each client involved in the transfer of required security levels and protocols, such as, for example, hashing algorithms used to confirm that received packets have not been altered. The IP server also insures that the client software at each end-user device has not been tampered, altered, or hacked.

[0148] The clients complete a messaging “handshake”, and then begin transfer of a file. More specifically, the second client requests a connection with the first client at step **1050**, the first client notifies the IP server of its status, e.g. that it is beginning a transfer, at step **1060**, the first client grants the second client’s request at step **1070**, and the second client notifies the IP server of its status, e.g. that its connection request was accepted, at step **1080**. The file transfer begins at step **1090**.

[0149] Periodically, both clients will update the server on the status of the download, as illustrated by exemplary steps **1100** and **1110**. The server will keep track of the file transfer and compare the information received from both clients for completeness and security. Once the file transfer is completed, at step **1120**, a status is sent of each client is sent to the IP server at steps **1130** and **1140**, and the connection is terminated at step **1150**. The clients continue to update their availability with the IP server for future downloads, as illustrated by exemplary steps **1160** and **1170**.

[0150] It will be appreciated that because one of the problems with the TCP/IP protocol is that significant timing delays can occur between communications, using a virtual machine advantageously allows messages to be sent at the lowest levels of the stack between virtual machines of different clients, thus helping insure that communications are not delayed. Further, the inclusion of local routing capabilities enables each client to setup another communication link, if needed, for example to

continue a stalled download. Further still, as preferably both clients include such routing capability, either client can reinitiate a separate communication to the other client, thus helping insure that TCP/IP packet delay timeouts do not draw out the communication.

[0151] Additionally, to facilitate more robust transfers, one of the clients can instruct the other to open other TCP/IP connections independent of the server link. For example, a first client may receive an IP address for a second client via the IP server, and the second client could then communicate additional IP addresses to the first client and communicate duplicate packets via connections established with these additional IP addresses, thus increasing the reliability of the link. Additionally, the client could send multiple packets over separate IP addresses to insure a different starting point for transmission, and thus insure unique paths. It will be appreciated that this might advantageously allow for the continuing transfer of packets even if one of the connection paths fails. Notably, each path is closed upon completion of the transmission.

[0152] FIG. 9 illustrates a user interface for an exemplary file sharing application in accordance with a preferred implementation. To initiate a transfer, a user clicks on an application icon to open the user's Friend's List and a "Sharzing" window. Bold texted names identify on-line contacts, while grey texted names indicate off-line contacts. When the blades of the graphical connection representation on the right side of the window, i.e. the Sharzing, are shut, the Sharzing is inactive. Clicking on an on-line contact opens the blades and establishes a Sharzing connection. The user may then "drag and drop" a file onto the open Sharzing.

[0153] Once a Sharzing connection is established, multiple files can be transferred in either direction. Further, multiple Sharzings can be opened simultaneously to different users. Preferably, when a Sharzing is connected, wallpaper of the opposite PC that is being connected to is displayed. As a file is "dragged and dropped" on the Sharzing, the Sharzing displays the progress of the file transfer. Using a Sharzing skin, a Sharzing depiction can take on identities such as, for example, a futuristic StarGate motif. In the case where such a StarGate motif is used, flash wormhole turbulence may begin when a file is placed in the Sharzing, and, subsequently, an opening at the end of the wormhole may emerge to display an image of the file and/or the recipient's desktop wallpaper. Preferably, when the transferred file is visible on the destination desktop, the transfer is complete.

Many Clients to Many Clients - Video and Audio Conferencing Implementation

[0154] In another preferred implementation, direct connections between end-user devices having virtualized routing capabilities are utilized in a telecommunications context, such as, for example, in an audio and video conferencing application.

[0155] It will be appreciated that in traditional audio and video conferencing applications, one or more conventional servers act to collate and process audio and video streams from a plurality of clients, and then distribute the processed audio and video streams to the clients. By way of contrast, in a preferred implementation, an end-user device can instead establish a direct connection with another end-user device, and communicate audio and video directly to the other end-user device,

rather communicating through a conventional server. In such implementations, this transmitted audio and video can be directly processed by either a communicating end-user device, a receiving end-user device, or both, rather than by a conventional server.

[0156] As described above, via the use of virtualization, a first end-user device can establish a direct connection with not just one other end-user device, but with multiple other end-user devices. The first end-user device provides each other end-user device with its video and audio stream, thus effectively acting as a server by “serving” its video and audio stream to each other end-user device. Each of the other end-user devices involved in a video conference will receive such video and audio streams served by this first end-user device; however, each other end-user device will additionally serve its own video and audio streams. Thus, each end-user device can be characterized as functioning as both a server and a client, possibly at the same time, i.e. as a multiplexed client/server.

[0157] Notably, however, although the end-user devices assume some functionality more traditionally assumed by a conventional server in video conferencing applications, a control server is preferably still used to oversee the establishment and maintenance of connections between end-user devices. Unlike in a traditional implementation, however, it is preferred that little to no audio or video processing is handled at this control server, and that the audio and video streams between end-user devices are not routed through the control server.

[0158] Instead, the control server primarily provides authentication and security functionality. Preferably, the control server keeps track of a unique identification of each end-user device, software loaded on each end-user device, and an IP address of each end-user device. Additionally, the control server preferably controls which end-user devices can communicate, and at what times they may communicate. For example, the control server preferably provides functionality allowing a moderator to “talk” over every other user at any given time.

[0159] Each end-user device preferably continually provides information to the control server, including: a status of the end-user device, whether the end-user device is receiving audio, whether the end-user device has lost its connection, an application status, application data, whether software at the end-user device has been tampered with, a rate of one or more communication links, and a packet error rate of one or more communication links.

Use with Conventional Servers - Media Server Implementation

[0160] In some preferred implementations, direct connections between end-user devices having virtualized routing capabilities are utilized in combination with one or more conventional file servers, such as, for example, in a media server application. Specifically, it will be appreciated that the conventional downloading of data, such as a video file, from a server is an intensive process that utilizes precious server processing, memory, and bandwidth. In preferred implementations, some of the strain of this process is offloaded from such a conventional server to one or more end-user devices having virtualized routing capabilities. This architecture decreases the processing, memory requirements and bandwidth loads on a media server. Table 2 of FIG. 10 shows the relationship for a

media file that is being downloaded from a server when some of the strain of multiple download requests is transferred off of the media server in accordance with such preferred implementations.

[0161] In a preferred implementation, a plurality of end-user devices comprise VDR clients, and a control server comprises a VDR server, each respectively including the architecture illustrated in FIG. 11.

[0162] FIG. 12 illustrates an exemplary process for downloading media content to two VDR clients that was originally stored on a customer video server, i.e. a media server. Notably, the process involves not just the two VDR clients and the media server, but also a VDR server as well.

[0163] The process begins when the first VDR client requests download of media content from the media server at step **1210**, followed by a corresponding TCP/IP handshake at step **1215**. Subsequently, the media server alerts the VDR server of the download at step **1220**. The VDR records the activity of the first VDR client along with necessary identification and contact information for the first VDR client. The media server follows the typical download procedure and begins the download to the first VDR client at step **1225**.

[0164] Thereafter, a second VDR client requests the same media content from the media server at step **1230**, followed by a corresponding TCP/IP handshake at step **1235**. At step **1240**, the media server alerts the VDR server to the download request by the second VDR client. The VDR server determines that a VDR client is active, gathers addressing information for the second VDR client, and notifies the media server that it will handle the download at step **1245**. Notably, a VDR client is active as long as its connection is active. It will be appreciated that several methodologies may be used to determine how long a client stays active. In at least some implementations, a client is shut down, i.e. rendered inactive, immediately after a file is transferred, which may represent the most efficient use of resources. In a preferred implementation, a timer is utilized, and the client remains active a user-specified number of minutes following activity. Alternatively, a client's connection could be left open until the user wants to close it, or until a network timeout occurs.

[0165] At step **1250**, the VDR server communicates to the first VDR client and configures it for download capability to the second VDR client, e.g. using the obtained addressing information for the second VDR client. The second VDR client initiates communication with the first VDR client for download of the media content from the first VDR client at step **1260**, followed by a corresponding TCP/IP handshake at step **1265**, and the download then begins at step **1270**.

[0166] Notably, the first VDR client, like most clients, has bandwidth available on the uplink when downloading content. It is believed that a typical personal computer, as of the filing date of this application, can handle 3-5 uploads without significant burdening or performance degradation.

[0167] Communication between the first and second VDR clients is accomplished between "Thin Virtual Machines" (TVM) of each VDR client. Each TVM is characterized as a "thin" virtual machine, as each preferably generally includes only functionality necessary to support virtualized networking, and, preferably, optimizes the resources needed to support the virtualization of the

Network Interface Card (NIC). As will be appreciated from the description hereinabove, each TVM enables each application to have a separate virtual interface to the NIC. This functionality enables customized security capabilities that can be added to each application interface individually.

[0168] At steps **1280** and **1285**, the VDR clients convey status information, e.g. concerning the download, to the VDR server. At step **1290**, the first VDR client completes its download of the media content from the media server. The first VDR client continues the download to the second VDR client, however. While the download continues, status information is sent to the VDR server from each VDR client as exemplified by step **1300**. Further, the first and second VDR clients continue to communicate via the virtual machine interface to detect connection issues and reroute packets.

[0169] The download continues at step **1310**. At step **1320**, the second VDR client completes its download, and each VDR client notifies the server of such success, as exemplified by step **1330**. The VDR server, in turn, notifies the media server that the download to the second VDR client was completed successfully at step **1340**.

[0170] If, instead of being completed successfully, the second VDR client's download of the media content had not completed, the second VDR client would have contacted the VDR server for another download opportunity.

[0171] FIG. 13 illustrates another exemplary process where, rather than downloading media content from one other VDR client, media content is downloaded from a plurality of VDR clients, thus increasing the speed of download.

[0172] More specifically, a media file is broken into fragments, and each fragment is downloaded to a target VDR client from a different source VDR client using a different connection. In FIG. 13, the process begins when, at step **1410**, a first VDR client communicates a download request to a media server, followed by a corresponding TCP/IP handshake at step **1415**. At step **1420**, the media server alerts the VDR server that a download has been requested. The VDR server determines that multiple VDR clients are available to download the requested media content from, and, at step **1430**, the VDR server informs the media server that it will handle the download request. At steps **1440**, **1450**, and **1460**, respectively, the VDR server communicates to second, third, and fourth VDR clients and passes addressing information corresponding to the first VDR client to each. The VDR server assigns each VDR client the portion of the media content that that VDR client will download to the first VDR client.

[0173] The first VDR client then downloads, at steps **1445**, **1455**, and **1465** respectively, the assigned portions of the media content from each of the other VDR clients. As exemplified by illustration of steps **1470**, **1472**, **1474** and **1476**, each VDR client reports to the VDR server status information on any downloads it is a part of, to insure each download is progressing as planned. If a connection is lost, the VDR server can act to correct the problem. Once the first VDR client has completed the download of the media content, it communicates such completion to the VDR server and each other

VDR client, as exemplified by illustration of step **1480**. Subsequently, the VDR server notifies the media server that the download was completed at step **1490**.

MMORPG Implementation

[0174] In another preferred implementation, direct connections between end-user devices having virtualized routing capabilities are utilized in a gaming context, such as, for example, in a massively multiplayer online role playing game application.

[0175] It will be appreciated that traditional MMORPGs handle the majority of processing for a game world at conventional servers. In a preferred implementation, some of this processing work is offloaded to end-user devices having virtualized routing capabilities. For example, each end-user device preferably functions as a server for serving an avatar associated with a user to other end-user devices whose users are disposed, in the game world, in close proximity. In this way, the processing associated with such avatars is largely offloaded from the server.

[0176] This offloading, and other similar offloading, reduces the resources required by an MMORPG server. Notably, however, although the end-user devices assume some functionality more traditionally assumed by a conventional server in MMORPG applications, a control server is preferably still used to oversee the establishment and monitor connections between end-user devices.

[0177] The control server preferably provides authentication and security functionality. Preferably, the control server keeps track of a unique identification of each end-user device, software loaded on each end-user device, and an IP address of each end-user device. Additionally, the control server preferably controls what actions each client can take.

[0178] Each end-user device preferably continually provides information to the control server, including: a status of the end-user device, whether the end-user device has lost its connection, an application status, application data, whether software at the end-user device has been tampered with, a rate of one or more communication links, a packet error rate of one or more communication links, a game state, a character state, and coordinates of the character's location in the game world.

[0179] It will be appreciated that voice conferencing can be an important part of the massive multiplayer experience, and, in accordance with one or more preferred embodiments, functionality and implementation similar to that outlined in an audio conferencing context is utilized in a massively multiplayer gaming context as well.

[0180] Notably, in such implementations, a client both receives information from other clients, for example in the form of avatar information, and additionally receives information from a content server, which may also comprise control server functionality.

Security

[0181] Preferably, in secure implementations, clients at end-user devices are alerted by a control server of an impending transfer and utilize a secure protocol such as public key encryption, AES (Advanced Encryption Standard), or SSL (Secure Socket Layer). Packets to be transferred are preferably intercepted by a virtual machine of a first client, prior to being sent to the network interface

of that client, and encrypted. Following receipt, the packets coming out of the network interface are then intercepted by a virtual machine of the other client and decrypted.

[0182] Preferably, strong security is achieved by employing a single encryption key that is passed between the two end-user devices controlled at layer 2 and 3 of the OSI (Open Source Interface) stack model. Regardless of whether the file is transported via Ethernet, WiFi, mobile phone data networks, or other wired or wireless technologies, the file is protected since it is decrypted at the router level of the destination before the data is passed to the application.

[0183] Although systems, methods, and apparatus are described herein largely in the context of end-user devices having virtualized routing capabilities, it will be appreciated that at least some implementations may be practiced in the absence of such virtualized routing capabilities.

[0184] Notably, virtualized routing capabilities, such as, for example, those presented by a VDR client, may be advantageous even in communicating with a client that does not enjoy such capabilities, e.g. a non-VDR client. In a preferred method, a VDR client in communication with a non-VDR client searches incoming packets for viruses or other anomalies, and, if such other anomalies are found, the VDR client can break off communications and re-establish a new connection.

Deflects and Spread Spectrum Networking

[0185] FIG. 14 illustrates the use of a dispersive virtual machine implemented as part of a software application that can be easily downloaded to a device such as a PC, smart phone, tablet, or server. In accordance with a preferred methodology, data to be sent from a first device utilizing such software to another device is split up into multiple parts which are sent separately over different routes and then reassembled at the other device.

[0186] FIG. 15 illustrates an exemplary such methodology utilizing one or more such exemplary software applications in which a message sent from a first user's (Joe's) computer to a second user's (Mary's) computer is split up into three parts, each part is transmitted to a different device also running such an exemplary software application, and each different device then retransmits such received part to the second user's (Mary's) computer. The message is then reassembled and displayed on the second user's (Mary's) computer. In this scenario, each of the three different devices that the message is communicated to for retransmission can be characterized as a "deflect", in that rather than sending the message directly to the second user's (Mary's) computer, the message is first intentionally "deflected" through such different devices. Notably, although complete control of routing over one or more networks a message will pass through may not be available, by choosing to utilize one or more such deflects, multiple paths can be selected and utilized.

[0187] It will be appreciated that splitting a message into three parts and sending such parts over three separate routes makes it more difficult to intercept such message. Leveraging the deflect capability of virtual dispersive routing, packets can be sent independent routes from one another, thus ensuring packets cannot be captured by copying packets off the Internet from a single router.

[0188] Different parts of a message could be sent all at once over different network paths, or, in at least some preferred implementations, a network path could be changed over time, for example, a device could be configured to utilize a different deflect after a certain period of time or after a certain number of packets are communicated.

[0189] Additionally or alternatively, however, in one or more preferred implementations, rather than utilizing multiple deflects to send packets over multiple network paths and/or change network paths, network communications utilize one or more network addresses and/or ports. Preferably, this includes concurrent communication of packets to multiple network addresses and/or ports of another device, and/or changing or shifting of the network address and/or port that packets are communicated to. Such spreading of packets across multiple IP addresses and ports can be characterized as spread spectrum networking (SSN).

[0190] In one or more preferred implementations, by constantly changing the IP address and/or ports that a computer device utilizes to talk to another computing device (e.g. P2P (peer to peer) or V2V (virtual machine to virtual machine)), it becomes much more difficult to track the device. In one or more preferred implementations, virtual machines are utilized to provide signaling and measure delay between hops.

[0191] FIG. 16 illustrates how multiple packets can be sent over different deflects in a direct spreading of packets methodology, and FIG. 17 illustrates how multiple packets can be sent to different IP addresses and/or ports in a hopping IP addresses and ports methodology.

[0192] FIG. 18 illustrates an exemplary system architecture configured to allow clients in a task network to access the Internet through an interface server using virtual dispersive networking (VDN) spread spectrum protocols.

[0193] The exemplary system illustrated in FIG. 18 includes a task network with ten VDN enabled client devices, a dispersive presence server, four deflect servers, and an interface server. In a preferred implementation, the task network client machines and the interface server will be hosted by a first entity at a particular location or facility, and the presence server and deflects will be hosted by another entity at disparate geographic locations.

[0194] Each of the task network client machines has installed thereon VDN software that includes a Xen hypervisor with CentOS host operating system containing virtual dispersive networking components. Each client machine further includes two Windows guest operating systems. The VDN components are configured to provide connectivity for one of the guest operating systems to the Interface Server through multiple deflects. This interface server provides access to internet sites such as Google for one of the guest operating systems for each of the task network clients. The VDN components provide connectivity for the other guest operating system only to the internal task network.

[0195] The interface server has installed thereon VDN software that includes a Xen hypervisor with CentOS host operating system containing virtual dispersive networking components. The interface

server further includes software configured to provide connectivity to the internet from the interface server. Dispersive presence server software and deflect server software is installed on geographically dispersed hosted servers.

[0196] All of the client machines and servers are configured to provide desired connectivity using virtual dispersive networking. Configuration of software is preferably adjusted as required based on performance metrics developed through testing.

[0197] FIG. 19 illustrates an exemplary system architecture configured to enable a workstation with four independent connections to the internet to send traffic using virtual dispersive networking spread spectrum protocols to an interface server from each independent internet connection.

[0198] The exemplary system illustrated in FIG. 19 includes a workstation configured with four independent connections to the internet and including VDN software configured to simultaneously utilize the independent connections. The system further includes an interface server, a dispersive presence server (DPS), and four deflect servers.

[0199] In this exemplary system, the servers are configured as described hereinabove with respect to FIG. 18, while the workstation is loaded with VDN software that includes a Xen hypervisor with CentOS host operating system containing Virtual Dispersive Networking components. The workstation also has loaded thereon guest operating systems. The VDN components provide connectivity for the guest operating systems to the Interface Server through multiple deflects and through each of the independent connections to the internet.

Point of Entry Gateways

[0200] In one or more preferred implementations, VDN systems use a DPS for connection processing and collecting network information on a client population. The connection processing on the DPS communicates with the virtual networking machines on the end clients to make connections between the two end-points of the segments that make up a VDN communication (that is, two end clients).

[0201] In one or more preferred implementations, however, some of the functionality handled by a DPS is spawned out to intermediate network computing points (e.g. servers), which can then be in communication with the DPS. In one or more preferred implementations, one or more point of entry gateways are utilized in this manner. These point of entry (PoE) gateways, as their name implies, serve as a point of entry for devices (such as end clients) seeking to communicate with the DPS, by receiving communications intended for the DPS and forwarding them on (either with or without processing at the PoE gateway).

[0202] In one or more preferred implementations, PoE gateways are utilized in combination with a DPS behind a firewall. That is, the DPS remains behind a firewall, while the PoE gateways are exposed. Thus, by spawning processes out to intermediate network computing points (e.g. the PoE gateways), this enables the database and call processing software control of the DPS to remain behind the firewall. The data needed by the PoE gateways will be specific to a particular VDN and also provide other generic PoE gateway devices to diversify the network connection.

[0203] FIGS. 20-23 illustrate an exemplary scenario utilizing PoE gateways.

[0204] An end client **2020** configured for VDN (behind a firewall **2022**) seeks to establish a connection with another end client **2030** (behind a firewall **2032**). The end client **2020** communicates a message requesting information for such a connection to a PoE gateway **2040** of a plurality of PoE gateways, as illustrated in FIG. 20. The PoE gateway **2040** then communicates the message on to the DPS **2010** (behind a firewall **2012**), as illustrated in FIG. 21. Thereafter, the DPS **2010** communicates connection information to both the end client **2020** and the end client **2030**, as illustrated in FIG. 22. The end clients **2020,2030** can utilize this information to set up a connection between one another, as illustrated in FIG. 23.

[0205] Although these communications and connections are illustrated as being generally direct (albeit preferably utilizing VDN, as described herein), in one or more preferred implementations, one or more deflects are utilized for any or all of these communications or connections. Specifically, in one or more preferred implementations, split traffic (e.g. VDN traffic with deflects) can be used between two network devices (e.g. DPS and PoE gateway or PoE gateway and VTC (Virtual Thin Client)). For example, FIGS. 24-27 illustrate the same scenario as outlined above, only with one or more deflects utilized for each communication or connection.

[0206] In one or more preferred implementations, such PoE gateways are utilized to ensure that all devices with important data and network information are behind a firewall. Further, it is believed that use of PoE gateways allows a DPS address to be obscured from network detection.

[0207] As illustrated in one or more preferred implementations, more than one PoE gateway can be utilized. In at least some preferred implementations, many PoE gateways can be used, reused and discarded easily.

[0208] The use of PoE gateways is further believed to allow system connections to be spread out over more computers, thus making a network or system harder to understand (and thus hack). Additionally, the use of PoE gateways is believed, in one or more preferred implementations, to alleviate throughput and connection limitations on a DPS, improve scaling, enable multiple points of entry, and provide points of entry that do not contain user data.

[0209] Specifically, with regard to this last point, as noted above, network attacks are a common occurrence in today's cyber environment, and computers that are not protected by firewalls are under constant attack. Servers are especially vulnerable due to the fact that they must keep a port open on a firewall to enable a connection to a client. Using computers or devices, such as PoE gateways, that do not have valuable information on them as interfaces is believed to be a useful way to protect user data.

Virtual Dispersive Hand-Off

[0210] Traditionally, network connections between different systems typically don't hand off connections smoothly or quickly. For example, as a mobile device with a WiFi connection moves from a first area, covered by a first WiFi network, to a second area, covered by a second WiFi

connection, there is typically not a smooth transition from one network to the other, and instead connectivity will typically be disrupted as the first WiFi network is lost.

[0211] In one or more preferred implementations, however, VDN is utilized for a smooth hand-off. For example, with respect to the previously outlined scenario, the use of multiple virtual network connections would allow the device to connect to both WiFi networks, and enable a smooth transition as the device moves from the first area to the second area.

[0212] In one or more preferred implementations, in order to figure out which access points to connect to, several mechanisms are used. Several exemplary such mechanisms and methodologies will now be described, although it will be appreciated that this is not an exhaustive list.

[0213] In one or more preferred implementations, GPS is utilized to determine the speed and direction of a device and determine network devices in close proximity for connection.

[0214] In one or more preferred implementations, the device detects RF signal and connects to as many devices as feasible. Once connection is achieved, network presence is used to determine the connection point and move the network connection to the new access point. In one or more preferred implementations, devices can maintain a connection by using Spread Spectrum Protocol setups to “roll” the connection and continue streaming data while moving from access point to access point.

[0215] In one or more preferred implementations, by setting up deflects, routing paths and data can be forced to the next access point that a mobile device is moving toward. That is, deflects can be utilized to force traffic to a particular access point.

[0216] In some preferred implementations, data is overlapped (e.g. sent to or from the mobile device via multiple connections) to insure that communicated information has been or will be received (either by the mobile device or by a device it is being sent to by the mobile device), as illustrated in FIG. 28. In one or more preferred implementations, data intended for a mobile device can be preloaded to a deflect and extracted when the mobile device connects to the access point. In some preferred implementations, data is given a timeframe to exist on the deflect, and is removed or archived once this timeframe expires, so that the deflect is not overloaded with “old” data.

[0217] Such hand-off functionality is believed to allow for the use of WiFi devices to stream data to mobile devices while the mobile devices are moving, allow for the ability to increase the data throughput of a system, enable the use of inexpensive mobile carrier platforms (e.g. WiFi hotspots and public WiFi), maintain a connection for mobile devices across networks and access points, and force traffic (via the use of deflects) to an access point for connection to a mobile device.

Multi-Protocol Dispersion

[0218] In one or more preferred implementations utilizing VDN, a deflect simply passes data through without reformatting it. In some cases, then, VDN traffic can be detected and identified by tracking the data in and out of a deflect. In one or more preferred implementations, a deflect is configured to reformat data to another protocol (e.g. from UDP to a TCP connection) so as to obviate this potential

issue. Such a change in protocol is believed to make it much more difficult for external hackers to detect VDN data and try to put it back together.

[0219] As an example, FIGS. 29 and 30 both illustrate exemplary scenarios utilizing a system which includes, inter alia, DPS **2110**, PoE gateway **2140**, two virtual thin clients (VTCs) **2120,2130**, and two deflects **2150,2160**. Specifically, FIG. 29 illustrates VDN routing from virtual thin client (VTC) **2120** to VTC **2130** using two deflects **2150,2160** that are configured simply to pass data through. In contrast to this, FIG. 30 illustrates a system in which the deflects **2150,2160** are configured to reformat data to another protocol. In the illustrated example, the deflect **2150** reformats data from HTTP to TCP, and the deflect **2160** reformats data from UDP to RTP.

[0220] This is believed to make it very difficult for a hacker or monitoring technology to detect VDN. Additionally, this allows traffic of one type to be made to look like another.

Superframes

[0221] It will be appreciated that one or more VDN methodologies disclosed herein can be characterized as resulting in VDN traffic that goes in and out of deflects in a 1:1 relationship. That is, in accordance with such a methodology, each packet received at a deflect is subsequently routed out of the deflect. Notably, however, such VDN traffic can be detected in and out of deflects due to the 1:1 nature of the traffic.

[0222] In one or more preferred implementations, “superframes” are utilized, which represent a single packet or communication that contains two packets or communications therein, optionally together with instructions as to where each packet is bound. In one or more preferred implementations, such a superframe further includes instructions as to what device should split a superframe up.

[0223] For example, FIG. 31A illustrates a first packet **3120** which includes data A, as well as a second packet **3130** which includes data B. FIG. 31B illustrates a superframe which has been constructed that includes both data A and data B. In one or more preferred implementations, such superframe also includes instructions to break apart this data, e.g. instructions as to where to send each set of data, and, preferably even instructions as to where this data should be broken apart (e.g. at a particular deflect). FIGS. 32-33 illustrate a process in which a superframe is sent from a first virtual thin client (VTC) **3220** to a deflect **3260**. The second deflect breaks apart the superframe and communicates some data onward toward the second VTC **3230**, and some data onward toward another deflect **3270**. The deflect **3270** may then route the packet towards the second VTC **3230**, or somewhere else. In one or more preferred implementations, a superframe includes instructions as to where the data contained therein is to be sent. Further, in one or more preferred implementations, a superframe includes instructions as to where it is to be broken apart.

[0224] Although in the illustrated example the data packets the superframe was broken into were communicated onward towards different devices, in one or more preferred implementations a superframe is broken up but both packets continue to be routed along the same path.

[0225] Further, although illustrated and described with reference to a superframe containing two packets or sets of data, in one or more preferred implementations a superframe may contain three or more packets or sets of data, which may be bound for different destinations or the same destination, and may travel the same route or path, or may travel different routes or paths. Further, in one or more preferred implementations, a superframe might include other superframes nested therein. For example, FIG. 34 illustrates a superframe **3410** which itself includes a superframe **3420** (which includes two packets **3430,3440**) as well as a packet **3440**.

[0226] Thus, a deflect can be used as a tool to split a large frame or pack into multiple packets and forward the split packets to other deflects and/or one or more final destination computing devices.

[0227] The use of superframes is believed to facilitate making it very difficult for a hacker or monitoring technology to detect VDN.

[0228] Additionally, superframes can be utilized to facilitate the broadcasting of data to other devices. For example, this could be utilized in support of one to many type applications.

Application Split Traffic

[0229] Typically, applications need to keep track of connection state to operate and suffer when network connections are poor. In one or more preferred implementations, virtualization is utilized to insulate an application from IP address and network connection duties, which enables the splitting of traffic over multiple links. Preferably, communications are intercepted at the session layer and parsed out to each link, as illustrated in FIG. 35. For example, traffic from an application can be split up and parsed out, at the session layer, to a plurality of different virtual connections. Data received over these connections can similarly be reconstructed and passed up to the application. Notably, a connection can be converted to other protocols for diversity.

[0230] This not only provides the ability to split traffic of a connection, but additionally provides the ability to setup channels to roll a VDN spread spectrum connection, as well as the ability to hide networking information.

Priority Through Time Measurement

[0231] Conventionally, priority of packets through a network is primarily accomplished by marking or labeling a packet high priority. Unfortunately, this does not guarantee a fast transfer of data from end to end. For example, not all routers are able to handle or have been setup to handle special packet markings.

[0232] In one or more preferred implementations, VDN is utilized to provide better end to end quality of service (QoS) without resorting to labeling or marking packets (although it will be appreciated that methodologies disclosed herein could be utilized in combination with labeling or marking packets).

[0233] In one or more preferred implementations, a system that can message at the lowest level in the stack can determine the time it takes for packets to move from end to end. For example, as illustrated in FIG. 36, first and second end devices **3610,3620** can determine that it takes 100 ms to communicate

from one end to another through a first deflect, 600 ms to communicate from one end to another through a second deflect, and 500 ms to communicate from one end to another through a third deflect.

[0234] Thus, in this scenario, higher priority traffic from the first end device **3610** to the second end device **3620** could be given priority by being routed via the first deflect, while lower priority traffic could be routed via the second and third deflects.

[0235] More generally, if multiple paths are available, end devices can find network connections that give faster end to end throughput and use of such routes and connections can be allocated, prioritized, and/or reserved to devices that can be designated to require higher throughput, such as, for example, emergency management, police, fire, military or other high priority customers and decision makers. Users without as high of priority could have a threshold set of a mid-range time as a minimum thereby keeping the faster routes available for the higher priority calls but still enabling the system to enable connections. By choosing connections with shorter time to send and receive packets, the parallel connection can actually speed up data being sent between the end points.

[0236] In one or more preferred implementations, end devices are designated with a particular priority level, which might represent an arbitrary designation or could simply represent a minimum and/or maximum threshold time.

[0237] In one or more preferred implementations, an end device determines the time it takes packets to travel to a particular destination over several different routes, and selects a route in accord with its priority level. For example, returning to the example illustrated in FIG. 1, the end device **3610** might have a priority level setting a minimum threshold at 300 ms, and would accordingly limit its communications to use of the second and third deflects (at 600 ms and 500 ms respectively), thereby reserving use of the first deflect (corresponding to a 100 ms travel time) to higher priority devices/communications.

[0238] In one or more preferred implementations, priority designations or minimum or maximum thresholds may be set at an application or connection level rather than on a per device basis, and/or particular applications and/or connections may deviate from a device level priority designation or minimum or maximum threshold.

[0239] In one or more preferred implementations, priority provision based on time measurement is believed to provide the ability to give priority to packets without requiring a change to current packets, the ability to give service to lower priority connections without jeopardizing higher speed connections, the ability to dynamically set time thresholds, the ability to give individual devices extremely high priority, allow for measurement corresponding to end to end priority and quality of service rather than simply measuring the speed through an individual router.

Multi-Layer Deflects

[0240] Systems in which a dispersive presence server (DPS) downloads a spread spectrum protocol (SSP) to a plurality of virtual thin clients (VTCs) are disclosed herein. In exemplary such systems, a VTC can connect to a deflect that is either specified or chosen from a pool of deflects. Such a deflect

can be behind or in front of a firewall. The deflect then connects to a destination or final VTC. Further, in one or more preferred implementations, a VTC can itself be a deflect (for example, a third VTC can be utilized as a deflect for communications from a first VTC to a second VTC).

[0241] In one or more preferred implementations, rather than simply using one deflect along any communication route from a source end point to a destination end point, a SSP setup calls out for multi-layers of deflects. For example, FIG. 37 illustrates several data paths that pass through multiple deflects, including a first communication path that passes through deflect device 1 and deflect device 3, and a second communication path that passes through deflect device 2 and deflect device 5. Although only two layers of deflects are illustrated, in one or more preferred implementations more than two layers of deflects may be utilized.

[0242] Moreover, multi-layer deflect implementations can utilize the same deflect multiple times. In one or more preferred implementations, a loop is even set up between two or more deflects whereby data is continually looped between such deflects. In one or more preferred implementations, data is stored in a network by being looped continuously between deflects. In theory, such data could be stored indefinitely in the network.

[0243] In preferred implementations, VTCs are utilized as deflects in order to make it more difficult to separate traffic originating or terminating at that VTC from traffic for which that VTC serves as a deflect.

[0244] Further, the use of deflects, and multi-layer deflects, is believed to facilitate obscuring of source and destination IP address information.

VDN Casting

[0245] Most communication paths and networks are not 100% reliable. Some communication paths will not be available all of the time. In one or more preferred implementations, VDN casting is utilized to find a way to reach a VTC over varied communication paths.

[0246] At times, a destination VTC may be in an area that has limited communication or sporadic communication. In this and similar situations, another VTC can test deflects for connection to the destination VTC. A virtual dispersive network (VDN) can be rather large, but since the members of a VDN keep track of connection information (even though a dispersive presence server may lose connectivity), the network can still communicate as long as the network has not change significantly and the destination VTC address had not changed.

[0247] FIG. 38 illustrates how a connection from a first VTC to a second VTC through a second deflect is still possible even though it is not possible to connect through a first or third deflect.

[0248] In one or more preferred implementations, VDN casting is believed to facilitate making a network delay tolerant, and allow network devices to be found at later points in time.

Splits Inside of Splits

[0249] In accordance with one or more preferred implementations, dispersive networking which splits IP traffic between a source and destination significantly reduces the ability to hack an IP based

communication connection or packet stream by forcing independent paths using deflects. The spreading of these packets over deflects that are strategically placed on a network (such as the internet) forces a hacker to collect the packets and figure out how to associate the packets from the independent paths before the hacker can start to break the encryption. A communication that is split and communicated over three different paths (e.g. using deflects) is more difficult to break than a communication that occurs over a single path, but it is still possible for a hacker to collect packets and work to put the packets back together.

[0250] For example, FIG. 39 illustrates a system including two end clients that are communicating over a network that includes three deflects. Each connection (or hop) is labeled with an A to indicate that it is associated with the same split. That is, communication has been split into three streams that are each associated with the same split.

[0251] Although the system illustrates the use of three deflects, the number of deflects could be increased or decreased based on processing power, physical hardware, and security requirements. Preferably, the deflects can be used and movement of the traffic through new deflects can occur while the session communication stream is in process. Preferably, a direct channel (without any deflects) could also be used as a communication stream along with communication over deflects if an administrator enables it.

[0252] In accordance with one or more preferred implementations, to increase security, IP traffic can be split again either at a deflect and/or at an end client. For example, FIG. 40 illustrates a second splitting of communications from one of the deflects to the second end client. Each of the connections for the second split (which can be characterized as being “inside” of the first split) is labeled with a B to indicate that it is associated with the second split.

[0253] Preferably, such use of an additional split inside of a split operates to change packet headers, thereby confusing any hacker who is monitoring traffic. The use of additional splits can also increase reliability through a network and/or speed up throughput. Splits inside of splits can be used to help get traffic through a very congested portion of a network thereby increasing throughput and reliability of communications.

[0254] The use of a split inside of a split also extends the number of hops in a network changing the characteristics of the system to confuse hackers. In accordance with one or more preferred implementations involving splits inside of splits, the number of streams forming a second split inside of a first split is independent of the number of streams in the first split.

[0255] In accordance with one or more preferred implementations, a system can utilize multiple splits inside of splits, as illustrated in FIG. 41. In accordance with one or more preferred implementations, each network connection from an end client to a deflect or a deflect to an end client is a candidate for additional splitting. These connections are illustrated as, and can be visualized as, line segments of a communication.

[0256] In accordance with one or more preferred implementations, each line segment can be encrypted separately as well as the end to end encryption of each communication from the source and destination (which represent client devices running endpoint software). In accordance with one or more preferred implementations, split inside of a split methodologies are utilized in systems in which encryption is used at the application level and/or the public key for that application does not transmit in the open. In accordance with one or more preferred implementations, a key is communicated utilizing splits inside of a split.

[0257] In accordance with one or more preferred implementations, techniques such as interleaving and insertion of "red herring" data (false data) can be used within a split.

[0258] In accordance with one or more preferred implementations, endpoint and deflect software can be loaded on to the same computing device. Further, in accordance with one or more preferred implementations, multiple instances of endpoint and deflect software can be loaded on to the same computing device, as illustrated in FIG. 42, in which several nodes represent devices having both endpoint and deflect software installed thereon. Since endpoint and deflect software can be loaded on to the same computing device along with multiple instances, additional computing machines are typically not required to do the splitting to host the multiple instances of software.

[0259] FIG. 43 illustrates an exemplary system in which several devices have both endpoint and deflect software loaded thereon. Because multiple instances of deflect software can be loaded on the same device, different streams can utilize different instances of deflect software loaded on the same device, as illustrated in FIG. 44. Further, in accordance with one or more preferred implementations, different streams may even utilize the same instance of deflect software.

[0260] Additionally, in accordance with one or more preferred implementations, streams passing through a device may be received at or communicated from multiple, different software instances loaded on that device, and combined or split at software at the device, as illustrated in FIG. 45.

[0261] In accordance with one or more preferred implementations, additional security is achieved by utilizing a virtual machine running its open guest operating system (e.g. MAC OS, Windows, Linux, or Android), and loading endpoint and/or deflect client software on the guest operating system. By loading client software (e.g. endpoint and/or deflect) on a single computer, communications between such software may not be available to hackers (e.g. available over the public internet). As only a portion of the total communications would be publicly available to intercept, this would serve to dramatically increase the difficulty of trying to reverse engineer splits and assemble packets. FIGS. 46 and 47 illustrate exemplary systems involving use of a virtual machine having endpoint and/or deflect client software loaded on a guest operating system.

[0262] In accordance with one or more preferred implementations, running multiple instances of endpoint or deflect software on a single machine complicates the ability of a hacker to "camp" at a choke point in a network, copy packets and attempt to decode the headers and put the message back together.

[0263] In accordance with one or more preferred implementations, only a portion of any single split appears on a network (e.g. the internet) if client software is run on a single machine.

[0264] In accordance with one or more preferred implementations involving splits inside of splits, similar to use of parentheses in a mathematical equation, inside or inner splits must be put back together before outer or larger splits.

[0265] In accordance with one or more preferred implementations, splits are not algorithmically chosen so extensive processing power does not help to steal information, although in at least one or more preferred implementations splits may be algorithmically chosen. In one or more preferred implementations, splits are chosen based on networking information.

[0266] Based on the foregoing description, it will be readily understood by those persons skilled in the art that the present invention is susceptible of broad utility and application. Many embodiments and adaptations of the present invention other than those specifically described herein, as well as many variations, modifications, and equivalent arrangements, will be apparent from or reasonably suggested by the present invention and the foregoing descriptions thereof, without departing from the substance or scope of the present invention. Accordingly, while the present invention has been described herein in detail in relation to one or more preferred embodiments, it is to be understood that this disclosure is only illustrative and exemplary of the present invention and is made merely for the purpose of providing a full and enabling disclosure of the invention. The foregoing disclosure is not intended to be construed to limit the present invention or otherwise exclude any such other embodiments, adaptations, variations, modifications or equivalent arrangements, the present invention being limited only by the claims appended hereto and the equivalents thereof.

What is claimed is:

1. A method of communicating data, the method comprising:
 - (a) splitting into a first plurality of data streams, at endpoint software running on a first device, first data for communication to a destination device;
 - (b) selecting, at the first device by the endpoint software, a first plurality of deflects for use in communicating the first plurality of data streams;
 - (c) communicating each of the first plurality of data streams over a different one of the selected first plurality of deflects;
 - (d) splitting, at a first deflect, a first data stream of the first plurality of data streams into a second plurality of data streams;
 - (e) selecting, at the first deflect, a second plurality of deflects for use in communicating the second plurality of data streams;
 - (f) communicating each of the second plurality of data streams over a different one of the selected second plurality of deflects;
 - (g) receiving, at a second deflect from a first set of deflects, the second plurality of data streams;
 - (h) reassembling, at the second deflect, the second plurality of data streams into the first data stream, and communicating the first data stream onward to another device;
 - (i) receiving, at a second device from a second set of deflects, the first plurality of data streams including the reassembled first data stream; and
 - (j) reassembling, at the second device, the first plurality of data streams into the first data.
2. The method of claim 1, wherein each deflect represents deflect software loaded on a distinct physical computing device.
3. The method of claim 1, wherein two or more deflects represent deflect software loaded on the same physical computing device.
4. The method of claim 1, wherein the first device or the second device has both deflect software and endpoint software loaded thereon.
5. The method of claim 1, wherein one or more deflects represent deflect software loaded on a guest operating system running in a virtual machine.
6. The method of claim 1, wherein the method involves communication of a data stream between two virtual machines running on the same physical computing device.
7. The method of claim 1, wherein the method involves utilizing encryption for communications between nodes.

8. The method of claim 1, wherein the method involves utilizing node-to-node encryption for communications between nodes in combination with end-to-end encryption for the entire communication from the first device to the second device.
9. The method of claim 1, wherein application-level encryption is utilized.
10. The method of claim 1, wherein public-key encryption is utilized and a key is communicated utilizing a methodology involving splits inside of splits.
11. The method of claim 1, wherein interleaving is utilized inside of one or more splits.
12. The method of claim 1, wherein the insertion of red herring data is utilized within one or more splits.
13. The method of claim 1, wherein selecting, at the first device by the endpoint software, the first plurality of deflects for use in communicating the first plurality of data streams comprises selecting the first plurality of deflects based at least in part on networking information.
14. The method of claim 1, wherein endpoint software is loaded on a desktop computer, a tablet, a smartphone, a smart watch, a wearable computing device, or a smart appliance.
15. The method of claim 1, wherein deflect software is loaded on a desktop computer, a tablet, a smartphone, a smart watch, a wearable computing device, or a smart appliance.
16. One or more computer readable media containing computer executable instructions for performing the method of claim 1.
17. A system for performing the method of claim 1.
18. A software program for performing the method of claim 1.