



US008560705B2

(12) **United States Patent**  
**Larson et al.**

(10) **Patent No.:** **US 8,560,705 B2**  
(45) **Date of Patent:** **\*Oct. 15, 2013**

(54) **SYSTEM AND METHOD EMPLOYING AN AGILE NETWORK PROTOCOL FOR SECURE COMMUNICATIONS USING SECURE DOMAIN NAMES**

(75) Inventors: **Victor Larson**, Fairfax, VA (US);  
**Robert Dunham Short, III**, Leesburg, VA (US); **Edmond Colby Munger**, Crownsville, MD (US); **Michael Williamson**, South Riding, VA (US)

(73) Assignee: **VirnetX, Inc.**, Zephyr Cove, NV (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **13/342,795**

(22) Filed: **Jan. 3, 2012**

(65) **Prior Publication Data**

US 2012/0102206 A1 Apr. 26, 2012

**Related U.S. Application Data**

(63) Continuation of application No. 13/049,552, filed on Mar. 16, 2011, which is a continuation of application No. 11/840,560, filed on Aug. 17, 2007, now Pat. No. 7,921,211, which is a continuation of application No. 10/714,849, filed on Nov. 18, 2003, now Pat. No. 7,418,504, which is a continuation of application No. 09/558,210, filed on Apr. 26, 2000, now abandoned, which is a continuation-in-part of application No. 09/504,783, filed on Feb. 15, 2000, now Pat. No. 6,502,135, which is a continuation-in-part of application No. 09/429,643, filed on Oct. 29, 1999, now Pat. No. 7,010,604.

(60) Provisional application No. 60/106,261, filed on Oct. 30, 1998, provisional application No. 60/137,704, filed on Jun. 7, 1999.

(51) **Int. Cl.**  
**G06F 15/16** (2006.01)

(52) **U.S. Cl.**  
USPC ..... **709/227**

(58) **Field of Classification Search**

USPC ..... 709/223-227  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

2,895,502 A 7/1959 Roper et al.  
4,677,434 A 6/1987 Fascenda

(Continued)

**FOREIGN PATENT DOCUMENTS**

DE 19924575 12/1999  
EP 0838930 4/1988

(Continued)

**OTHER PUBLICATIONS**

ITU-T Recommendation H.323, "Infrastructure of Audiovisual Services—Systems and Terminal Equipment for Audiovisual Services. Packet-Based Multimedia Communications System," International Telecommunications Union, pp. 1-128, Feb. 1998.

(Continued)

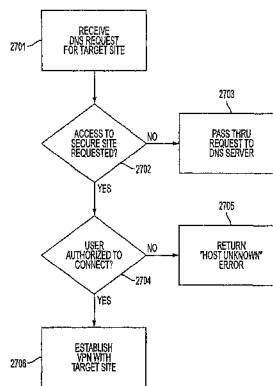

*Primary Examiner* — Krisna Lim

(74) *Attorney, Agent, or Firm* — McDermott Will & Emery LLP

(57) **ABSTRACT**

A client device comprises: (a) a memory, (b) an application program, and (c) a signal processing configuration. The memory is configured and arranged to facilitate a connection of the client device with a target device over a secure communication link created based on (i) an address request generated by the client device, and (ii) a determination as a result of the address request that the target device is a device with which a secure communication link can be established when the requested address is identified in an address lookup. The application program is configured and arranged so as to allow participation in audio/video communications with the target device over the secure communication link once the secure communication link is established. The signal processing configuration is arranged to execute the application program.

**30 Claims, 40 Drawing Sheets**



(56)

References Cited

U.S. PATENT DOCUMENTS

4,761,334 A	8/1988	Sagoi et al.	6,158,011 A	12/2000	Chen et al.
4,885,778 A	12/1989	Weiss	6,168,409 B1	1/2001	Fare
4,920,484 A	4/1990	Ranade	6,173,399 B1	1/2001	Gilbrech
4,933,846 A	6/1990	Humphrey et al.	6,175,867 B1	1/2001	Taghadoss
4,952,930 A	8/1990	Franaszek et al.	6,178,409 B1	1/2001	Weber et al.
4,988,990 A	1/1991	Warrior	6,178,505 B1	1/2001	Schneider et al.
5,007,051 A	4/1991	Dolkas et al.	6,179,102 B1	1/2001	Weber et al.
5,164,988 A	11/1992	Matyas	6,182,141 B1	1/2001	Blum et al.
5,204,961 A	4/1993	Barlow	6,182,227 B1	1/2001	Blair et al.
5,276,735 A	1/1994	Boebert et al.	6,199,112 B1	3/2001	Wilson
5,303,302 A	4/1994	Burrows	6,202,081 B1	3/2001	Naudus
5,311,593 A	5/1994	Carmi	6,222,842 B1	4/2001	Sasyan et al.
5,329,521 A	7/1994	Walsh et al.	6,223,287 B1	4/2001	Douglas et al.
5,341,426 A	8/1994	Barney et al.	6,226,748 B1	5/2001	Bots et al.
5,345,439 A	9/1994	Marston	6,226,751 B1	5/2001	Arrow et al.
5,367,643 A	11/1994	Chang et al.	6,233,618 B1	5/2001	Shannon
5,384,848 A	1/1995	Kikuchi	6,243,360 B1	6/2001	Basilico
5,511,122 A	4/1996	Atkinson	6,243,749 B1	6/2001	Sitaraman et al.
5,548,646 A	8/1996	Aziz et al.	6,243,754 B1	6/2001	Guerin et al.
5,559,883 A	9/1996	Williams	6,246,670 B1	6/2001	Karlsson et al.
5,561,669 A	10/1996	Lenney et al.	6,256,671 B1	7/2001	Strentzsch et al.
5,588,060 A	12/1996	Aziz	6,262,987 B1	7/2001	Mogul
5,590,285 A	12/1996	Krause et al.	6,263,445 B1	7/2001	Blumenau
5,625,626 A	4/1997	Umekita	6,266,699 B1	7/2001	Sevcik
5,629,984 A	5/1997	McManis	6,269,099 B1	7/2001	Borella et al.
5,654,695 A	8/1997	Olnowich et al.	6,286,047 B1	9/2001	Ramanathan et al.
5,682,480 A	10/1997	Nakagawa	6,298,341 B1	10/2001	Mann et al.
5,689,566 A	11/1997	Nguyen	6,301,223 B1	10/2001	Hrastar et al.
5,689,641 A	11/1997	Ludwig et al.	6,308,213 B1	10/2001	Valencia
5,740,375 A	4/1998	Dunne et al.	6,308,274 B1	10/2001	Swift
5,757,925 A	5/1998	Faybishenko	6,311,207 B1	10/2001	Mighdoll et al.
5,764,906 A	6/1998	Edelstein et al.	6,314,463 B1	11/2001	Abbott et al.
5,771,239 A	6/1998	Moroney et al.	6,324,161 B1	11/2001	Kirch
5,774,660 A	6/1998	Brendel et al.	6,330,562 B1	12/2001	Boden et al.
5,787,172 A	7/1998	Arnold	6,332,158 B1	12/2001	Risley et al.
5,790,548 A	8/1998	Sistanizadeh et al.	6,333,272 B1	12/2001	McMillin et al.
5,796,942 A	8/1998	Esbensen	6,338,082 B1	1/2002	Schneider
5,805,801 A	9/1998	Holloway et al.	6,353,614 B1	3/2002	Borella et al.
5,805,803 A	9/1998	Birrell et al.	6,425,003 B1	7/2002	Herzog et al.
5,822,434 A	10/1998	Caronni et al.	6,430,155 B1	8/2002	Davie et al.
5,838,796 A	11/1998	Mittenthal	6,430,610 B1	8/2002	Carter
5,842,040 A	11/1998	Hughes et al.	6,487,598 B1	11/2002	Valencia
5,845,091 A	12/1998	Dunne et al.	6,496,867 B1	12/2002	Beser et al.
5,864,666 A	1/1999	Shrader	6,499,108 B1	12/2002	Johnson
5,867,650 A	2/1999	Osterman	6,502,135 B1	12/2002	Munger et al.
5,870,610 A	2/1999	Beyda et al.	6,505,232 B1	1/2003	Mighdoll et al.
5,878,231 A	3/1999	Baehr et al.	6,510,154 B1	1/2003	Mayes et al.
5,884,038 A	3/1999	Kapoor	6,549,516 B1	4/2003	Albert et al.
5,892,903 A	4/1999	Klaus	6,557,037 B1	4/2003	Provino
5,898,830 A	4/1999	Wesinger, Jr. et al.	6,560,634 B1	5/2003	Broadhurst
5,905,859 A	5/1999	Holloway et al.	6,571,296 B1	5/2003	Dillon
5,918,018 A	6/1999	Gooderum et al.	6,571,338 B1	5/2003	Shaio et al.
5,918,019 A	6/1999	Valencia	6,581,166 B1	6/2003	Hirst et al.
5,950,195 A	9/1999	Stockwell et al.	6,606,708 B1	8/2003	Devine et al.
5,950,519 A	9/1999	Anatoli	6,615,357 B1	9/2003	Boden et al.
5,960,204 A	9/1999	Yinger et al.	6,618,761 B2	9/2003	Munger et al.
5,996,016 A	11/1999	Thalheimer et al.	6,671,702 B2	12/2003	Munger et al.
6,006,259 A	12/1999	Adelman et al.	6,687,551 B2	2/2004	Kruglikov et al.
6,006,272 A	12/1999	Aravamudan et al.	6,687,746 B1	2/2004	Steindl
6,016,318 A	1/2000	Tomoiike	6,701,437 B1	3/2004	Shuster et al.
6,016,512 A	1/2000	Huitema	6,714,970 B1	3/2004	Hoke et al.
6,041,342 A	3/2000	Yamaguchi	6,717,949 B1	3/2004	Fiveash et al.
6,052,788 A	4/2000	Wesinger et al.	6,751,738 B2	4/2004	Boden et al.
6,055,574 A	4/2000	Smorodinsky et al.	6,752,166 B2	6/2004	Wesinger, Jr. et al.
6,061,346 A	5/2000	Nordman	6,757,740 B1	6/2004	Lull et al.
6,061,736 A	5/2000	Rochberger et al.	6,760,766 B1	6/2004	Parekh et al.
6,079,020 A	6/2000	Liu	6,813,777 B1 *	7/2004	Sahlqvist
6,081,900 A	6/2000	Subramaniam et al.	6,826,616 B2	11/2004	Weinberger et al. .... 725/76
6,092,200 A	7/2000	Muniyappa et al.	6,839,759 B2	1/2005	Larson et al.
6,101,182 A	8/2000	Sistanizadeh et al.	6,937,597 B1	8/2005	Larson et al.
6,119,171 A	9/2000	Alkhatib	7,010,604 B1	3/2006	Rosenberg et al.
6,119,234 A	9/2000	Aziz et al.	7,039,713 B1	5/2006	Munger et al.
6,131,121 A	10/2000	Mattaway et al.	7,072,964 B1	7/2006	Van Gunter et al.
6,147,976 A	11/2000	Shand et al.	7,133,930 B2	11/2006	Whittle et al.
6,157,957 A	12/2000	Berthaud	7,167,904 B1	1/2007	Munger et al.
			7,188,175 B1	3/2007	Devarajan et al.
			7,188,180 B2	3/2007	McKeeth
			7,197,563 B2	3/2007	Larson et al.
			7,353,841 B2	4/2008	Sheymov et al.
					Kono et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

7,418,504	B2	8/2008	Larson et al.	
7,461,334	B1	12/2008	Lu et al.	
7,490,151	B2	2/2009	Munger et al.	
7,493,403	B2	2/2009	Shull et al.	
7,584,500	B2 *	9/2009	Dillon et al.	726/3
7,764,231	B1	7/2010	Karr et al.	
7,852,861	B2 *	12/2010	Wu et al.	370/401
7,921,211	B2	4/2011	Larson et al.	
7,933,990	B2	4/2011	Munger et al.	
8,051,181	B2	11/2011	Larson et al.	
2001/0049741	A1	12/2001	Skene et al.	
2002/0002675	A1	1/2002	Bush	
2002/0004898	A1	1/2002	Droge	
2003/0196122	A1	10/2003	Wesinger, Jr. et al.	
2004/0199493	A1	10/2004	Ruiz et al.	
2004/0199520	A1	10/2004	Ruiz et al.	
2004/0199608	A1	10/2004	Rechterman et al.	
2004/0199620	A1	10/2004	Ruiz et al.	
2005/0055306	A1	3/2005	Miller et al.	
2005/0108517	A1 *	5/2005	Dillon et al.	713/150
2006/0059337	A1	3/2006	Poyhonen et al.	
2006/0123134	A1	6/2006	Munger et al.	
2007/0208869	A1	9/2007	Adelman et al.	
2007/0214284	A1	9/2007	King et al.	
2007/0266141	A1	11/2007	Norton	
2008/0005792	A1	1/2008	Larson et al.	
2008/0144625	A1 *	6/2008	Wu et al.	370/392
2008/0235507	A1	9/2008	Ishikawa et al.	
2009/0193498	A1 *	7/2009	Agarwal et al.	726/1
2009/0193513	A1 *	7/2009	Agarwal et al.	726/15
2009/0199258	A1 *	8/2009	Deng et al.	725/142
2009/0199285	A1	8/2009	Agarwal et al.	

FOREIGN PATENT DOCUMENTS

EP	0814589	12/1997
EP	836306	4/1998
EP	0858189	8/1998
GB	2317792	4/1998
GB	2334181	8/1999
GB	2340702	2/2000
JP	62-214744	9/1987
JP	04-117826	4/1992
JP	04-363941	12/1992
JP	09-018492	1/1997
JP	09-270803	10/1997
JP	10-070531	3/1998
JP	10-111848	4/1998
JP	10-215244	8/1998
WO	WO9827783	6/1998
WO	WO9843396	10/1998
WO	WO9855930	12/1998
WO	WO9859470	12/1998
WO	WO9911019	3/1999
WO	WO9938081	7/1999
WO	WO9948303	9/1999
WO	WO0017775	3/2000
WO	WO0070458	11/2000
WO	WO01/61922	2/2001
WO	WO0116766	3/2001
WO	WO0150688	7/2001

OTHER PUBLICATIONS

ITU-T Recommendation H.225.0, "Infrastructure of Audiovisual Services—Transmission Multiplexing and Synchronization. Call Signaling Protocols and Media Stream Packetization for Packet-Based Multimedia Communication systems," International Telecommunication Union, pp. 1-155, Feb. 1998.

ITU-T Recommendation H.235, "Infrastructure of Audiovisual Services—Systems Aspects. Security and Encryption for H-Series (H.323 and other H.245-based) Multimedia Terminals," International Telecommunication Union, pp. 1-39, Feb. 1998.

ITU-T Recommendation H.245, "Infrastructure of Audiovisual Services—Communication Procedures. Control Protocol for Multimedia Communication," International Telecommunication Union, pp. 1-280, Feb. 1998.

Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 8,051,181).

Transmittal Letters (Patent No. 8,051,181).

Exhibit X5, Droms, R., RFC 2131, "Dynamic Host Configuration Protocol," 1987.

Cisco Comments and Petition for Reexamination 95/001,679 dated Jun. 14, 2012.

Exhibit S, Declaration of Nathaniel Polish, Ph.D.

Exhibit R, Excerpts from Patent Owner & Plaintiff VimetX Inc.'s First Amended P.R. 3-1 and 3-2 Disclosure of Asserted Claims and Infringement Contentions.

Extended European Search Report dated Mar. 26, 2012 from Corresponding European Application No. 11005793.2.

Bergadano, et al., "Secure WWW Transactions Using Standard HTTP and Java Applets," Proceedings of the 3rd USENIX Workshop on Electronic Commerce, 1998.

Alexander Invalidity Expert Report dtd May 22, 2012 with Exhibits. Deposition of Peter Alexander dtd Jul. 27, 2012.

Cisco '151 Comments by Third Party Requester dtd Aug. 17, 2012 with Exhibits.

Cisco '151 Petition to Waive Page Limit Requirement for Third Party Comments dtd Aug. 17, 2012.

Deposition of Stuart Stubblebine dtd Aug. 22, 2012.

Defendants' Motion for Reconsideration of the Construction of the Term "Secure Communication Link," 7 pages, Jun. 2012.

Green, "Cisco Leverages Altiga Technology for VPN's," 2 pages, 2000 <http://www.crn.com/news/channel-programs/18807923/cisco-leverages-altiga-technology-for-vpns.htm>.

Altiga Networks Archived at <http://web.archive.org/web/20000823023437/http://www.altiga.com/products/> 1999 and Retrieved by the Wayback Machine.

Kiuchi, "C-HTTP The Development of a Secure, Closed HTTP-Based Network on the Internet," Department of Epidemiology and Biostatistics, Faculty of Medicine, University of Tokyo, Japan.

Third Party Requester Comments dated Jun. 25, 2012—After Non Final Office Action (U.S. Appl. No. 95/001,788).

Reexam Affidavit/Declaration/Exhibit Filed by 3rd Party on Jun. 25, 2012 (U.S. Appl. No. 95/001,788).

Hopen Transcript dated Apr. 11, 2012.

VirnetX Claim Construction Opinion.

Declaration of Angelos D. Keromytis, Ph.D.

Declaration of Dr. Robert Dunham Short III.

Exhibit A-1, Verdict Form from *VirnetX, Inc. v. Microsoft Corp.*, No. 6:07-CV-80 (E.D. Tex.).

Exhibit A-3, Declaration of Jason Nieh, Ph.D. (U.S. Appl. No. 95/001,269).

Exhibit A-4, Redacted Deposition of Chris Hopen from *VirnetX, Inc. v. Cisco Systems, Inc.*, No. 6:07-CV 417 (E.D. Tex. Apr. 11, 2012).

Exhibit B-1, Excerpt from Deposition of Defense FY 2000/2001 Biennial Budget Estimates, (Feb. 1999).

Exhibit B-2, Collection of Reports and Presentations on DARPA Projects.

Exhibit B-3, Maryann Lawlor, Transient Partnerships Stretch Security Policy Management, Signal Magazine (Sep. 2001) <http://www.afcea.org/signal/articles/annviewer.asp?a=494&print=yes>.

Joel Snyder, Living in Your Own Private Idaho, Network World (Jan. 28, 1998) <http://www.networkworld.com/intranet/0126review.html>.

Time Greene, CEO's Chew the VPN Fat, CNN.com (Jun. 17, 1999), <http://www.cnn.com/TECH/computing/9906/17/vpnfat.ent.idg/index.html?iref=allsearch>.

Lee et al., "Uniform Resource Locators (URL)," Network Working Group, RFC 1738, Dec. 1994 (25 pages).

VPN 3000 Concentrator Series, User Guide; Release 2.5 Jul. 2000 (489 pages).

VPN 3000 Concentrator Series, Getting Started; Release 2.5 Jul. 2000 (122 pages).

Fratto, Altiga Concentrates on VPN Security (Hardware Review Evaluation), Network Computing, Mar. 22, 1999 (2 pages).

(56)

**References Cited**

## OTHER PUBLICATIONS

Response to RFP: Altiga, Network World Fusion, May 10, 1999 (7 pages).

Altiga Proves Multi-Vendor Interoperability for Seamless VPN Deployment; VPN Workshop Marks Significant Development in the VPN Market, Jul. 12, 1999 (2 pages).

Altiga VPN Concentrator Series (C50) Versus Nortel Networks Contivity Extranet Switch 4000 and 4500, VPN Tunneling competitive Evaluation, 1999 (6 pages).

VPN 3000 Client User Guide, Release 2.5, Jul. 2000 (94 pages).

Digital Certificates Design Specification for Release 2.0, May 17, 1999 (21 pages).

Altiga IPsec Client Architecture, Revision 1.0, Apr. 5, 1999 (34 pages).

Altiga IPsec Functional Specification, Revision 2.1, (17 pages).

Altiga Product Requirements, Revision 1.7, May 26, 1998 (17 pages).

Altiga Network Lists Feature Functional Specification, Revision 1.0, (7 pages).

Altiga Split Tunneling Functional/Design Specification, (15 pages).

Altiga Digital Certificate Support for IPsec Client V2.1 Functional Specification, Aug. 12, 1999 (24 pages).

Altiga IPsec LAN to LAN Tunnel Autodiscovery Functional Specification, (5 pages).

Altiga Split Tunneling Testplan, Revision 1.0, (8 pages).

Altiga VPN Concentrator Getting Started, Revision 1, Mar. 1999 (116 pages).

Altiga VPN Concentrator Getting Started, Version 2, Jun. 1999 (102 pages).

Altiga VPN Concentrator Getting Started, Version 3, Dec. 1999 (130 pages).

Altiga VPN Concentrator Getting Started, Version 4, Mar. 2000 (138 pages).

Altiga VPN Concentrator User Guide, Revision 1, Mar. 1999 (304 pages).

Altiga VPN Concentrator User Guide, Revision 1.1, Mar. 1999 (304 pages).

Altiga VPN Concentrator User Guide, Version 3, Jun. 1999 (478 pages).

Altiga VPN Concentrator User Guide, Version 4, Dec. 1999 (472 pages).

Altiga VPN Concentrator User Guide, Version 5, Mar. 2000 (606 pages).

Altiga VPN Client Installation and User Guide, Version 2, Jul. 1999 (92 pages).

Altiga VPN Concentrator VPN Client Installation and User Guide, Version 3, Dec. 1999 (113 pages).

Altiga VPN Concentrator VPN Client Installation and User Guide, Version 4, Mar. 2000 (118 pages).

Altiga Networks VPN Concentrator and VPN Client, as well as their Public Demonstrations and Testing, are also Described in Marketing Materials and Publications (4 pages).

Peter Alexander Invalidity Report.

Defendants' Second Supplemental Joint Invalidity Contentions.

Exhibit 118A, *Altiga VPN System vs. Claims of the '135 Patent*.

Exhibit 119A, *Altiga VPN System vs. Claims of the '151 Patent*.

Exhibit 120A, *Altiga VPN System vs. Claims of the '180 Patent*.

Exhibit 121A, *Altiga VPN System vs. Claims of the '211 Patent*.

Exhibit 122A, *Altiga VPN System vs. Claims of the '504 Patent*.

Exhibit 123A, *Altiga VPN System vs. Claims of the '759 Patent*.

Exhibit 12A, *SSL 3.0 vs. Claims of the '135 Patent*.

Exhibit 13A, *SSL 3.0 vs. Claims of the '504 Patent*.

Exhibit 14A, *SSL 3.0 vs. Claims of the '211 Patent*.

Exhibit 228A, *Understanding OSF DCE 1.1 for AIX and OS/2 (APP\_VX0556531-804) vs. Claims of the '135 Patent*.

Exhibit 229A, *Understanding OSF DCE 1.1 for AIX and OS/2 (APP\_VX0556531-804) vs. Claims of the '151 Patent*.

Exhibit 230A, *Understanding OSF DCE 1.1 for AIX and OS/2 (APP\_VX0556531-804) vs. Claims of the '180 Patent*.

Exhibit 231A, *Understanding OSF DCE 1.1 for AIX and OS/2 (APP\_VX0556531-804) vs. Claims of the '211 Patent*.

Exhibit 232A, *Understanding OSF DCE 1.1 for AIX and OS/2 (APP\_VX0556531-804) vs. Claims of the '504 Patent*.

Exhibit 233A, *Understanding OSF DCE 1.1 for AIX and OS/2 (APP\_VX0556531-804) vs. Claims of the '759 Patent*.

Exhibit 255, *Schulzrinne vs. Claims of the '135 Patent*.

Exhibit 256, *Schulzrinne vs. Claims of the '504 Patent*.

Exhibit 257, *Schulzrinne vs. Claims of the '211 Patent*.

Exhibit 258, *Schulzrinne vs. Claims of the '151 Patent*.

Exhibit 259, *Schulzrinne vs. Claims of the '180 Patent*.

Exhibit 260, *Schulzrinne vs. Claims of the '759 Patent*.

Exhibit 261, *SSL 3.0 vs. Claims of the '151 Patent*.

Exhibit 262, *SSL 3.0 vs. Claims of the '759 Patent*.

Exhibit 263, *Wang vs. Claims of the '135 Patent*.

*Wang vs. Claims of the '504 Patent*.

*Wang vs. Claims of the '211 Patent*.

Exhibit 1, Alexander CV.pdf.

Exhibit 2, Materials Considered by Peter Alexander.

Exhibit 3, Cross Reference Chart.

Exhibit 4, *RFC 2543 vs. Claims of the '135 Patent*.

Exhibit 5, *RFC 2543 vs. Claims of the '504 Patent*.

Exhibit 6, *RFC 2543 vs. Claims of the '211 Patent*.

Exhibit 7, *The Schulzrinne Presentation vs. Claims of the '135 Patent*.

Exhibit 8, *The Schulzrinne Presentation vs. Claims of the '504 Patent*.

Exhibit 9, *The Schulzrinne Presentation vs. Claims of the '211 Patent*.

Exhibit 10, *The Schulzrinne Presentation vs. Claims of the '151 Patent*.

Exhibit 11, *The Schulzrinne Presentation vs. Claims of the '180 Patent*.

Exhibit 12, *The Schulzrinne Presentation vs. Claims of the '759 Patent*.

Exhibit 13, *SSL 3.0 vs. Claims of the '135 Patent*.

Exhibit 14, *SSL 3.0 vs. Claims of the '504 Patent*.

Exhibit 15, *SSL 3.0 vs. Claims of the '211 Patent*.

Exhibit 16, *SSL 3.0 vs. Claims of the '151 Patent*.

Exhibit 17, *SSL 3.0 vs. Claims of the '759 Patent*.

Exhibit 18, *Kiuchi vs. Claims of the '135 Patent*.

Exhibit 19, *Kiuchi vs. Claims of the '504 Patent*.

Exhibit 20, *Kiuchi vs. Claims of the '211 Patent*.

Exhibit 21, *Kiuchi vs. Claims of the '151 Patent*.

Exhibit 22, *Kiuchi vs. Claims of the '180 Patent*.

Exhibit 23, *Kiuchi vs. Claims of the '759 Patent*.

Exhibit 24, U.S. Patent No. 6,119,234 (*hereinafter "Aziz"*) and *RFC 2401 vs. Claims of the '135 Patent*.

Exhibit 25, U.S. Patent No. 6,119,234 (*hereinafter "Aziz"*) and *RFC 2401 vs. Claims of the '504 Patent*.

Exhibit 26, U.S. Patent No. 6,119,234 (*hereinafter "Aziz"*) and *RFC 2401 vs. Claims of the '211 Patent*.

Exhibit 27, U.S. Patent No. 6,119,234 (*hereinafter "Aziz"*) and *RFC 2401 vs. Claims of the '151 Patent*.

Exhibit 28.

Exhibit 29, *The Altiga System vs. Claims of the '135 Patent*.

Exhibit 30, *The Altiga System vs. Claims of the '504 Patent*.

Exhibit 31, *The Altiga System vs. Claims of the '211 Patent*.

Exhibit 32, *The Altiga System vs. Claims of the '759 Patent*.

Exhibit 33, U.S. Patent No. 6,496,867 ("*Beser*") and *RFC 2401 vs. Claims of the '135 Patent*.

Exhibit 34, U.S. Patent No. 6,496,867 ("*Beser*") and *RFC 2401 vs. Claims of the '504 Patent*.

Exhibit 35, U.S. Patent No. 6,496,867 ("*Beser*") and *RFC 2401 vs. Claims of the '211 Patent*.

Exhibit 36, U.S. Patent No. 6,496,867 ("*Beser*") and *RFC 2401 vs. Claims of the '151 Patent*.

Exhibit 37, U.S. Patent No. 6,496,867 ("*Beser*") and *RFC 2401 vs. Claims of the '180 Patent*.

Exhibit 38, *Kent vs. Claims of the '759 Patent*.

Exhibit 39, *RFC 2538, Storing Certificates in the Domain Name System (DNS) vs. Claims of the '504 Patent*.

(56)

## References Cited

## OTHER PUBLICATIONS

- Exhibit 40, RFC 2538, *Storing Certificates in the Domain Name System (DNS) vs. Claims of the '211 Patent*.
- Exhibit 41, *Aziz ('646) vs. Claims of the '759 Patent*.
- Exhibit 42, *The PLX Firewall vs. Claims of the '759 Patent*.
- Exhibit A-1, *Kiuchi vs. Claims of the '135 Patent*.
- Exhibit B-1, *Kiuchi vs. Claims of the '211 Patent*.
- Exhibit C-1, *Kiuchi vs. Claims of the '504 Patent*.
- Exhibit D, Materials Considered.
- Exhibit E, Expert Report of Stuart G. Stubblebine, Ph.D.
- Exhibit F, Expert Report of Stuart G. Stubblebine, Ph.D.
- Exhibit G, Opening Expert Report of Dr. Stuart Stubblebine Regarding Invalidity of the '135, '211, and '504 Patents.
- U.S. Appl. No. 09/399,753, filed Sep. 22, 1998, Graig Miller et al.
- U.S. Appl. No. 60/134,547, filed May 17, 1999, Victory Sheymov.
- U.S. Appl. No. 60/151,563, filed Aug. 31, 1999, Bryan Whittles.
- Alan O. Frier et al., "The SSL Protocol Version 3.0", Nov. 18, 1996, printed from <http://www.netscape.com/eng/ss13/draft302.txt> on Feb. 4, 2002, 56 pages.
- August Bequai, "Balancing Legal Concerns Over Crime and Security in Cyberspace", *Computer & Security*, vol. 17, No. 4, 1998, pp. 293-298.
- D. B. Chapman et al., "Building Internet Firewalls", Nov. 1995, pp. 278-375.
- D. Clark, "US Calls for Private Domain-Name System", *Computer, IEEE Computer Society*, Aug. 1, 1998, pp. 22-25.
- Davila J et al, "Implementation of Virtual Private Networks at the Transport Layer", *Information Security, Second International Workshop, ISW'99. Proceedings (Lecture Springer-Verlag Berlin, Germany, [Online] 1999, pp. 85-102, XP002399276, ISBN 3-540-666.*
- Dolev, Shlomi and Ostrovsky, Rafil, "Efficient Anonymous Multicast and Reception" (Extended Abstract), 16 pages.
- Donald E. Eastlake, 3rd, "Domain Name System Security Extensions", *Internet Draft*, Apr. 1998, pp. 1-51.
- F. Halsall, "Data Communications, Computer Networks and Open Systems", Chapter 4, *Protocol Basics*, 1996, pp. 198-203.
- Glossary for the Linux FreeS/WAN project, printed from [http://liberty.freeswan.org/freeswan\\_trees/freeswan-1.3/doc/glossary.html](http://liberty.freeswan.org/freeswan_trees/freeswan-1.3/doc/glossary.html) on Feb. 21, 2002, 25 pages.
- J. Gilmore, "Swan: Securing the Internet against Wiretapping", printed from [http://liberty.freeswan.org/freeswan\\_trees/freeswan-1.3/doc/rationale.html](http://liberty.freeswan.org/freeswan_trees/freeswan-1.3/doc/rationale.html) on Feb. 21, 2002, 4 pages.
- James E. Bellare, "New Statement of Rules-Naming Internet Domains", *Internet Newsgroup*, Jul. 30, 1995, 1 page.
- Jim Jones et al., "Distributed Denial of Service Attacks: Defenses", *Global Integrity Corporation*, 2000, pp. 1-14.
- Laurie Wells (LANCASTERBIBELMAIL MSN COM); "Subject: Security Icon" *USENET Newsgroup*, Oct. 19, 1998, XP002200606, 1 page.
- Linux FreeS/WAN Index File, printed from [http://liberty.freeswan.org/freeswan\\_trees/freeswan-1.3/doc/](http://liberty.freeswan.org/freeswan_trees/freeswan-1.3/doc/) on Feb. 21, 2002, 3 Pages.
- P. Srisuresh et al., "DNS extensions to Network address Translators (DNS\_ALG)", *Internet Draft*, Jul. 1998, pp. 1-27.
- Reiter, Michael K. and Rubin, Aviel D. (AT&T Labs—Research), "Crowds: Anonymity for Web Transactions", pp. 1-23.
- RFC 2401 (dated Nov. 1998) *Security Architecture for the Internet Protocol (RTP)*.
- RFC 2543-SIP (dated Mar. 1999): *Session Initiation Protocol (SIP or SIPS)*.
- Rich Winkel, "CAQ: Networking With Spooks: The NET & The Control of Information", *Internet Newsgroup*, Jun. 21, 1997, 4 pages.
- Rubin, Aviel D., Geer, Daniel, and Ranum, Marcus J. (Wiley Computer Publishing), "Web Security Sourcebook", pp. 82-94.
- Search Report (dated Aug. 20, 2002), International Application No. PCT/US01/04340.
- Search Report (dated Aug. 23, 2002), International Application No. PCT/US01/13260.
- Search Report (dated Oct. 7, 2002), International Application No. PCT/US01/13261.
- Search Report, IPER (dated Nov. 13, 2002), International Application No. PCT/US01/04340.
- Search Report, IPER (dated Feb. 6, 2002), International Application No. PCT/US01/13261.
- Search Report, IPER (dated Jan. 14, 2003), International Application No. PCT/US01/13260.
- Shankar, A.U. "A verified sliding window protocol with variable flow control". *Proceedings of ACM SIGCOMM conference on Communications architectures & protocols*. pp. 84-91, ACM Press, NY, NY 1986.
- Shree Murthy et al., "Congestion-Oriented Shortest Multi-path Routing", *Proceedings of IEEE INFOCOM*, 1996, pp. 1028-1036.
- W. Stallings, "Cryptography and Network Security", 2nd, Edition, Chapter 13, *IP Security*, Jun. 8, 1998, pp. 399-440.
- Microsoft Corporation's Fourth Amended Invalidity Contentions dated Jan. 5, 2009, *VirnetX Inc. and Science Applications International Corp. v. Microsoft Corporation*.
- Appendix A of the Microsoft Corporation's Fourth Amended Invalidity Contentions dated Jan. 5, 2009.
- Concordance Table for the References Cited in Tables on pp. 6-15, 71-80 and 116-124 of the Microsoft Corporation's Fourth Amended Invalidity Contentions dated Jan. 5, 2009.
- I. P. Mockapetris, "DNS Encoding of Network Names and Other Types," *Network Working Group, RFC 1101 (Apr. 1989) RFC 1101, DNS SRV*.
- R. Atkinson, "An Internetwork Authentication Architecture," *Naval Research Laboratory, Center for High Assurance Computing Systems (Aug. 5, 1993)*. (Atkinson NRL, KX Records).
- Henning Schulzrinne, *Personal Mobility For Multimedia Services In The Internet*, *Proceedings of the Interactive Distributed Multimedia Systems and Services European Workshop at 143 (1996)*. (Schulzrinne 96).
- Microsoft Corp., *Microsoft Virtual Private Networking: Using Point-to-Point Tunneling Protocol for Low-Cost, Secure, Remote Access Across the Internet (1996)* (printed from 1998 PDC DVD-ROM). (Point to Point, Microsoft Prior Art VPN Technology).
- "Safe Surfing: How to Build a Secure World Wide Web Connection," *IBM Technical Support Organization*, (Mar. 1996). (Safe Surfing, Website Art).
- Goldschlag, et al., "Hiding Routing Information," *Workshop on Information Hiding, Cambridge, UK (May 1996)*. (Goldschlag II, Onion Routing).
- "IPSec Minutes From Montreal", *IPSEC Working Group Meeting Notes*, <http://www.sandleman.ca/ipsec/1996/08/msg00018.html> (Jun. 1996). (IPSec Minutes, FreeS/WAN).
- J. M. Galvin, "Public Key Distribution with Secure DNS," *Proceedings of the Sixth USENIX UNIX Security Symposium, San Jose, California, Jul. 1996*. (Galvin, DNSSEC).
- J. Gilmore, et al. "Re: Key Management, anyone? (DNS Keying)," *IPSec Working Group Mailing List Archives (Aug. 1996)*. (Gilmore DNS, FreeS/WAN).
- H. Orman, et al. "Re: 'Re: DNS? was Re: Key Management, anyone?'" *IETF IPsec Working Group Mailing List Archive (Aug. 1996-Sep. 1996)*. (Orman DNS, FreeS/WAN).
- Arnt Gullbrandsen & Paul Vixie, *ADNSRR for specifying the location of services (DNS SRV)*, *IETF RFC 2052 (Oct. 1996)*. (RFC 2052, DNS SRV).
- Freier, et al. "The SSL Protocol Version 3.0," *Transport Layer Security Working Group (Nov. 18, 1996)*. (SSL, Underlying Security Technology).
- M. Handley, H. Schulzrinne, E. Schooler, *Internet Engineering Task Force, Internet Draft, (Dec. 2, 1996)*. (RFC 2543 Internet Draft 1).
- M.G. Reed, et al. "Proxies for Anonymous Routing," *12th Annual Computer Security Applications Conference, San Diego, CA, Dec. 9-13, 1996*. (Reed, Onion Routing).
- Kenneth F. Alden & Edward P. Wobber, *The AltaVista Tunnel: Using the Internet to Extend Corporate Networks*, *Digital Technical Journal (1997)* (Alden, AltaVista).
- Automotive Industry Action Group, "ANX Release 1 Document Publication," *AIAG (1997)*. (AIAG, ANX).
- Automotive Industry Action Group, "ANX Release 1 Draft Document Publication," *AIAG Publications (1997)*. (AIAG Release, ANX).

(56)

## References Cited

## OTHER PUBLICATIONS

- Aventail Corp. "Aventail VPN Data Sheet," available at <http://www.archive.org/web/19970212013043/www.aventail.com/prod/vpndata.html> (1997). (Data Sheet, Aventail).
- Aventail Corp., "Directed VPN Vs. Tunnel," available at <http://web.archive.org/web/19970620030312/www.aventail.com/educate/directvpn.html> (1997). (Directed VPN, Aventail).
- Aventail Corp., "Managing Corporate Access to the Internet," Aventail AutoSOCKS White Paper available at <http://web.archive.org/199706200300312/www.aventail.com/educate/whitepaper/ipmw.html> (1997). (Corporate Access, Aventail).
- Aventail Corp., "VPN Server V2.0 Administration Guide," (1997). (VPN, Aventail).
- Goldschlag, et al. "Privacy on the Internet," Naval Research Laboratory, Center for High Assurance Computer Systems (1997). (Goldschlag I, Onion Routing).
- Microsoft Corp., *Installing Configuring and Using PPTP with Microsoft Clients and Servers* (1997). (Using PPTP, Microsoft Prior Art VPN Technology).
- Microsoft Corp., *IP Security for Microsoft Windows NT Server 5.0* (1997) (printed from 1998 PDC DVD-ROM). (IP Security, Microsoft Prior Art VPN Technology).
- Microsoft Corp., *Microsoft Windows NT Active Directory: An Introduction to the Next Generation Directory Services* (1997) (printed from 1998 PDC DVD-ROM). (Directory, Microsoft Prior Art VPN Technology).
- Microsoft Corp., *Routing and Remote Access Service for Windows NT Server New Opportunities Today and Looking Ahead* (1997) (printed from 1998 PDC DVD-ROM). (Routing, Microsoft Prior Art VPN Technology).
- Microsoft Corp., *Understanding Point-to-Point Tunneling Protocol PPTP* (1997) (printed from 1998 PDC DVD-ROM). (Understanding PPTP, Microsoft Prior Art VPN Technology).
- J. Mark Smith et al., *Protecting a Private Network: The AltaVista Firewall*, Digital Technical Journal (1997). (Smith, AltaVista).
- Naganand Doraswamy *Implementation of Virtual Private Networks (VPNs) with IP Security*, <draft-ietf-ipsec-vpn-00.txt> (Mar. 12, 1997). (Doraswamy).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Mar. 27, 1997). (RFC 2543 Internet Draft 2).
- Aventail Corp., "Aventail and Cybersafe to Provide Secure Authentication for Internet and Intranet Communication," Press Release, Apr. 3, 1997. (Secure Authentication, Aventail).
- D. Wagner, et al. "Analysis of the SSL 3.0 Protocol," (Apr. 15, 1997). (Analysis, Underlying Security Technologies).
- Automotive Industry Action Group, "ANXO Certification Authority Service and Directory Service Definition for ANX Release 1," AIAG Telecommunications Project Team and Bellcore (May 9, 1997). (AIAG Definition, ANX).
- Automotive Industry Action Group, "ANXO Certification Process and ANX Registration Process Definition for ANX Release 1," AIAG Telecommunications Project Team and Bellcore (May 9, 1997). (AIAG Certification, ANX).
- Aventail Corp., "Aventail Announces the First VPN Solution to Assure Interoperability Across Emerging Security Protocols," Jun. 2, 1997. (First VPN, Aventail).
- Syverson, et al. "Private Web Browsing," Naval Research Laboratory, Center for High Assurance Computer Systems (Jun. 2, 1997). (Syverson, Onion Routing).
- Bellcore, "Metrics, Criteria, and Measurement Technique Requirements for ANX Release 1," AIAG Telecommunications Project Team and Bellcore (Jun. 16, 1997). (AIAG Requirements, ANX).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Jul. 31, 1997). (RFC 2543 Internet Draft 3).
- R. Atkinson, "Key Exchange Delegation Record for the DNS," Network Working Group, RFC 2230 (Nov. 1997). (RFC 2230, KX Records).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Nov. 11, 1997). (RFC 2543 Internet Draft 4).
- 1998 Microsoft Professional Developers Conference DVD ("1998 PDC DVD-ROM") (including screenshots captured there from and produced as MSFTVX 00018827-00018832). (Conference, Microsoft Prior Art VPN Technology).
- Microsoft Corp., *Virtual Private Networking An Overview* (1998) (printed from 1998 PDC DVD-ROM) (Overview, Microsoft Prior Art VPN Technology).
- Microsoft Corp., *Windows NT 5.0 Beta Has Public Premiere at Seattle Mini-Camp Seminar attendees get first look at the performance and capabilities of Windows NT 5.0* (1998) (available at <http://www.microsoft.com/presspass/features/1998/10-19nt5.mspxpftue>). (NT Beta, Microsoft Prior Art VPN Technology).
- "What ports does SSL use" available at [stason.org/TULARC/security/ssl-talk/3-4-What-ports-does-ssl-use.html](http://stason.org/TULARC/security/ssl-talk/3-4-What-ports-does-ssl-use.html) (Ports, DNS SRV).
- Aventail Corp., "Aventail VPN V2.6 Includes Support for More Than Ten Authentication Methods Making Extranet VPN Development Secure and Simple," Press Release, Jan. 19, 1998. (VPN V2.6, Aventail).
- R. G. Moskowitz, "Network Address Translation Issues with IPsec," Internet Draft, Internet Engineering Task Force, Feb. 6, 1998. (Moskowitz).
- H. Schulzrinne, et al. "Internet Telephony Gateway Location," Proceedings of IEEE Infocom '98, The Conference on Computer Communications, vol. 2 (Mar. 29-Apr. 2, 1998). (Gateway, Schulzrinne).
- C. Huitema, 45 al. "Simple Gateway Control Protocol," Version 1.0 (May 5, 1998). (SGCP).
- DISA "Secret Internet Protocol Router Network," SIPRNET Program Management Office (D3113) DISN Networks, DISN Transmission Services (May 8, 1998). (DISA, SIPRNET).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (May 14, 1998). (RFC 2543 Internet Draft 5).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Jun. 17, 1998). (RFC 2543 Internet Draft 6).
- D. McDonald, et al. "PF\_KEY Key Management API, Version 2," Network Working Group, RFC 2367 (Jul. 1998). (RFC 2367).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Jul. 16, 1998). (RFC 2543 Internet Draft 7).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Aug. 7, 1998). (RFC 2543 Internet Draft 8).
- Microsoft Corp., *Company Focuses on Quality and Customer Feedback* (Aug. 18, 1998). (Focus, Microsoft Prior Art VPN Technology).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Sep. 18, 1998). (RFC 2543 Internet Draft 9).
- Atkinson, et al. "Security Architecture for the Internet Protocol," Network Working Group, RFC 2401 (Nov. 1998). (RFC 2401, Underlying Security Technologies).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Nov. 12, 1998). (RFC 2543 Internet Draft 10).
- Donald Eastlake, *Domain Name System Security Extensions*, IETF DNS Security Working Group (Dec. 1998). (DNSSEC-7).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Dec. 15, 1998). (RFC 2543 Internet Draft 11).
- Aventail Corp., "Aventail Connect 3.1/2.6 Administrator's Guide," (1999). (Aventail Administrator 3.1, Aventail).
- Aventail Corp., "Aventail Connect 3.1/2.6 User's Guide," (1999). (Aventail User 3.1, Aventail).
- Aventail Corp., "Aventail ExtraWeb Server v3.2 Administrator's Guide," (1999). (Aventail ExtraWeb 3.2, Aventail).
- Kaufman et al. "Implementing IPsec," (Copyright 1999). (Implementing IPSEC, VPN References).
- Network Solutions, Inc. "Enabling SSL," NSI Registry (1999). (Enabling SSL, Underlying Security Technologies).
- Check Point Software Technologies Ltd. (1999) (Check Point, Checkpoint FW).
- Arnt Gullbrandsen & Paul Vixie, *ADNSRR for specifying the location of services (DNS SRV)*, <draft-ietf-dnsind-frc2052bis-02.txt> (Jan. 1999). (Gullbrandsen 99, DNS SRV).
- C. Scott, et al. *Virtual Private Networks*, O'Reilly and Associates, Inc., 2nd ed. (Jan. 1999). (Scott VPNs).
- M. Handley, H. Schulzrinne, E. Schooler, Internet Engineering Task Force, Internet Draft, (Jan. 15, 1999). (RFC 2543 Internet Draft 12).

(56)

## References Cited

## OTHER PUBLICATIONS

- Goldschlag, et al., "Onion Routing for Anonymous and Private Internet Connections," Naval Research Laboratory, Center for High Assurance Computer Systems (Jan. 28, 1999). (Goldschlag III, Onion Routing).
- H. Schulzrinne, "Internet Telephony: architecture and protocols—an IETF perspective," *Computer Networks*, vol. 31, No. 3 (Feb. 1999). (Telephony, Schulzrinne).
- M. Handley, et al. "SIP: Session Initiation Protocol," Network Working Group, RFC 2543 and Internet Drafts (Dec. 1996-Mar. 1999). (Handley, RFC 2543).
- FreeS/WAN Project, *Linux FreeS/WAN Compatibility Guide* (Mar. 4, 1999). (FreeS/WAN Compatibility Guide, FreeS/WAN).
- Telcordia Technologies, "ANX Release 1 Document Corrections," AIAG (May 11, 1999). (Telcordia, ANX).
- Ken Hornstein & Jeffrey Altman, *Distributing Kerberos KDC and Realm Information with DNS* <draft-eitf-cat-krb-dns-locate-oo.txt> (Jun. 21, 1999). (Hornstein, DNS SRV).
- Bhattacharya, et al., "An LDAP Schema for Configuration and Administration of IPsec Based Virtual Private Networks (VPNs)," IETF Internet Draft (Oct. 1999). (Bhattacharya LDAP VPN).
- B. Patel, et al. "DHCP Configuration of IPSEC Tunnel Mode," IPSEC Working Group, Internet Draft 02 (Oct. 15, 1999). (Patel).
- Goncalves, et al. *Check Point FireWall-1 Administration Guide*, McGraw-Hill Companies (2000). (Goncalves, Checkpoint FW).
- "Building a Microsoft VPN: A Comprehensive Collection of Microsoft Resources," FirstVPN, (Jan. 2000). (FirstVPN Microsoft).
- Gulbrandsen, Vixie, & Esibov, *A DNS RR for specifying the location of services (DNS SRV)*, IETF RFC 2782 (Feb. 2000). (RFC 2782, DNS SRV).
- MITRE Organization, "Technical Description," Collaborative Operations in Joint Expeditionary Force Experiment (JEFX) 99 (Feb. 2000). (MITRE, SIPRNET).
- H. Schulzrinne, et al. "Application-Layer Mobility Using SIP," *Mobile Computing and Communications Review*, vol. 4, No. 3. pp. 47-57 (Jul. 2000). (Application, SIP).
- Kindred et al., "Dynamic VPN Communities: Implementation and Experience," DARPA Information Survivability Conference and Exposition II (Jun. 2001). (DARPA, VPN Systems).
- ANX 101: Basic ANX Service Outline. (Outline, ANX).
- ANX 201: Advanced ANX Service. (Advanced, ANX).
- Appendix A: Certificate Profile for ANX IPsec Certificates. (Appendix, ANX).
- Assured Digital Products. (Assured Digital).
- Aventail Corp., "Aventail AutoSOCKS the Client Key to Network Security," Aventail Corporation White Paper. (Network Security, Aventail).
- Cindy Moran, "DISN Data Networks: Secret Internet Protocol Router Network (SIPRNet)." (Moran, SIPRNET).
- Data Fellows F-Secure VPN+ (F-Secure VPN+).
- Interim Operational Systems Doctrine for the Remote Access Security Program (RASP) Secret Dial-In Solution. (RASP, SIPRNET).
- Onion Routing*, "Investigation of Route Selection Algorithms," available at <http://www.onion-router.net/Archives/Route/index.html>. (Route Selection, Onion Routing).
- Secure Computing, "Bullet-Proofing an Army Net," Washington Technology. (Secure, SIPRNET).
- SPARTA "Dynamic Virtual Private Network." (Sparta, VPN Systems).
- Standard Operation Procedure for Using the 1910 Secure Modems. (Standard, SIPRNET).
- Publicly available emails relating to FreeS/WAN (MSFTVX00018833-MSFTVX00019206). (FreeS/WAN emails, FreeS/WAN).
- Kaufman et al., "Implementing IPsec," (Copyright 1999) (Implementing IPsec).
- Network Associates *Gauntlet Firewall For Unix User's Guide Version 5.0* (1999). (Gauntlet User's Guide—Unix, Firewall Products).
- Network Associates *Gauntlet Firewall For Windows NT Getting Started Guide Version 5.0* (1999) (Gauntlet Getting Started Guide—NT, Firewall Products).
- Network Associates *Gauntlet Firewall For Unix Getting Started Guide Version 5.0* (1999) (Gauntlet Unix Getting Started Guide, Firewall Products).
- Network Associates *Release Notes Gauntlet Firewall for Unix 5.0* (Mar. 19, 1999) (Gauntlet Unix Release Notes, Firewall Products).
- Network Associates *Gauntlet Firewall For Windows NT Administrator's Guide Version 5.0* (1999) (Gauntlet NT Administrator's Guide, Firewall Products).
- Trusted Information Systems, Inc. *Gauntlet Internet Firewall Firewall-to-Firewall Encryption Guide Version 3.1* (1996) (Gauntlet Firewall-to-Firewall, Firewall Products).
- Network Associates *Gauntlet Firewall Global Virtual Private Network User's Guide for Windows NT Version 5.0* (1999) (Gauntlet NT GVPN, GVPN).
- Network Associates *Gauntlet Firewall For UNIX Global Virtual Private Network User's Guide Version 5.0* (1999) (Gauntlet Unix GVPN, GVPN).
- Dan Sterne *Dynamic Virtual Private Networks* (May 23, 2000) (Sterne DVPN, DVPN).
- Darrell Kindred *Dynamic Virtual Private Networks (DVPN)* (Dec. 21, 1999) (Kindred DVPN, DVPN).
- Dan Sterne et al. *TIS Dynamic Security Perimeter Research Project Demonstration* (Mar. 9, 1998) (Dynamic Security Perimeter, DVPN).
- Darrell Kindred *Dynamic Virtual Private Networks Capability Description* (Jan. 5, 2000) (Kindred DVPN Capability, DVPN) 11. Oct. 7, and 28 1997 email from Domenic J. Turchi Jr. (SPARTA00001712-1714, 1808-1811) (Turchi DVPN email, DVPN).
- James Just & Dan Sterne *Security Quickstart Task Update* (Feb. 5, 1997) (Security Quickstart, DVPN).
- Virtual Private Network Demonstration dated Mar. 21, 1998 (SPARTA00001844-54) (DVPN Demonstration, DVPN).
- GTE Internetworking & BBN Technologies DARPA *Information Assurance Program Integrated Feasibility Demonstration (IFD) 1.1 Plan* (Mar. 10, 1998) (IFD 1.1, DVPN).
- Microsoft Corp. Windows NT Server Product Documentation: Administration Guide—Connection Point Services, available at <http://www.microsoft.com/technet/archive/winntas/proddocs/inetconctservice/cpsops.mspx> (Connection Point Services).
- Microsoft Corp. Windows NT Server Product Documentation: Administration Kit Guide—Connection Manager, available at <http://www.microsoft.com/technet/archive/winntas/proddocs/inetconctservice/cmak.mspx> (Connection Manager).
- Microsoft Corp. Autodial Heuristics, available at <http://support.microsoft.com/kb/164249> (Autodial Heuristics).
- Microsoft Corp., Cariplo: Distributed Component Object Model, (1996) available at [http://msdn2.microsoft.com/en-us/library/ms809332\(printer\).aspx](http://msdn2.microsoft.com/en-us/library/ms809332(printer).aspx) (Cariplo I).
- Marc Levy, COM Internet Services (Apr. 23, 1999), available at [http://msdn2.microsoft.com/en-us/library/ms809302\(printer\).aspx](http://msdn2.microsoft.com/en-us/library/ms809302(printer).aspx) (Levy).
- Markus Horstmann and Mary Kirtland, DCOM Architecture (Jul. 23, 1997), available at [http://msdn2.microsoft.com/en-us/library/ms809311\(printer\).aspx](http://msdn2.microsoft.com/en-us/library/ms809311(printer).aspx) (Horstmann).
- Microsoft Corp., DCOM: A Business Overview (Apr. 1997), available at [http://msdn2.microsoft.com/en-us/library/ms809320\(printer\).aspx](http://msdn2.microsoft.com/en-us/library/ms809320(printer).aspx) (DCOM Business Overview I).
- Microsoft Corp., DCOM Technical Overview (Nov. 1996), available at [http://msdn2.microsoft.com/en-us/library/ms809340\(printer\).aspx](http://msdn2.microsoft.com/en-us/library/ms809340(printer).aspx) (DCOM Technical Overview I).
- Microsoft Corp., DCOM Architecture White Paper (1998) available in PDC DVD-ROM (DCOM Architecture).
- Microsoft Corp, DCOM—The Distributed Component Object Model, A Business Overview White Paper (Microsoft 1997) available in PDC DVD-ROM (DCOM Business Overview II).
- Microsoft Corp., DCOM—Cariplo Home Banking Over the Internet White Paper Microsoft 1996) available in PDC DVD-ROM (Cariplo II).

(56)

## References Cited

## OTHER PUBLICATIONS

Microsoft Corp., DCOM Solutions in Action White Paper (Microsoft 1996) available in PDC DVD-ROM (DCOM Solutions in Action).  
Microsoft Corp., DCOM Technical Overview White Paper (Microsoft 1996) available in PDC DVD-ROM (DCOM Technical Overview II).

125. Scott Suhy & Glenn Wood, DNS and Microsoft Windows NT 4.0, (1996) available at [http://msdn2.microsoft.com/en-us/library/ms810277\(printer\).aspx](http://msdn2.microsoft.com/en-us/library/ms810277(printer).aspx) (Suhy).

126. Aaron Skonnard, *Essential Winlnet* 313-423 (Addison Wesley Longman 1998) (Essential Winlnet).

Microsoft Corp. Installing, Configuring, and Using PPTP with Microsoft Clients and Servers, (1998) available at [http://msdn2.microsoft.com/enus/library/ms811078\(printer\).aspx](http://msdn2.microsoft.com/enus/library/ms811078(printer).aspx) (Using PPTP).  
Microsoft Corp., Internet Connection Services for MS RAS, Standard Edition, <http://www.microsoft.com/technet/Archive/winntas/proddocs/inetconctservice/bcgstart.mspix> (Internet Connection Services I).

Microsoft Corp., Internet Connection Services for RAS, Commercial Edition, available at <http://www.microsoft.com/technet/archive/winntas/proddocs/inetconctservice/bcgstrtc.mspix> (Internet Connection Services II).

Microsoft Corp., Internet Explorer 5 Corporate Deployment Guide—Appendix B: Enabling Connections with the Connection Manager Administration Kit, available at <http://www.microsoft.com/technet/prodtechnol/ie/deploy/deploy5/appendb.mspix> (IE5 Corporate Development).

Mark Minasi, *Mastering Windows NT Server 4* 1359-1442 (6th ed., Jan. 15, 1999) (Mastering Windows NT Server).

*Hands on, Self-Paced Training for Supporting Version 4.0* 371-473 (Microsoft Press 1998) (Hands On).

Microsoft Corp., MS Point-to-Point Tunneling Protocol (Windows NT 4.0), available at <http://www.microsoft.com/technet/archive/winntas/maintain/feasability/pptpwp3.mspix> (MS PPTP).

Kenneth Gregg, et al., *Microsoft Windows NT Server Administrator's Bible* 173-206, 883-911, 974-1076 (IDG Books Worldwide 1999) (Gregg).

Microsoft Corp., Remote Access (Windows), available at [http://msdn2.microsoft.com/enus/library/bb545687\(VS.85.printer\).aspx](http://msdn2.microsoft.com/enus/library/bb545687(VS.85.printer).aspx) (Remote Access).

Microsoft Corp., Understanding PPTP (Windows NT 4.0), available at <http://www.microsoft.com/technet/archive/winntas/plan/pptpdst.mspix> (Understanding PPTP NT 4).

Microsoft Corp., Windows NT 4.0: Virtual Private Networking, available at <http://www.microsoft.com/technet/archive/winntas/deploy/confeat/vpntwk.mspix> (NT4 VPN).

Anthony Northrup, *NT Network Plumbing: Routers, Proxies, and Web Services* 299-399 (IDG Books Worldwide 1998) (Network Plumbing).

Microsoft Corp., Chapter 1—Introduction to Windows NT Routing with Routing and Remote Access Service, available at <http://www.microsoft.com/technet/archive/winntas/proddocs/rras40/rrasch01.mspix> (Intro to RRAS).

Microsoft Corp., Windows NT Server Product Documentation: Chapter 5—Planning for Large-Scale Configurations, available at <http://www.microsoft.com/technet/archive/winntas/proddocs/rras40/rrasch05.mspix> (Large-Scale Configurations).

F-Secure, *F-Secure NameSurfer* (May 1999) (from FSECURE 00000003) (NameSurfer 3).

F-Secure, *F-Secure VPN Administrator's Guide* (May 1999) (from FSECURE 00000003) (F-Secure VPN 3).

F-Secure, *F-Secure SSH User's & Administrator's Guide* (May 1999) (from FSECURE 00000003) (SSH Guide 3).

F-Secure, *F-Secure SSH2.0 for Windows NT and 95* (May 1999) (from FSECURE 00000003) (SSH 2.0 Guide 3).

F-Secure, *F-Secure VPN+ Administrator's Guide* (May 1999) (from FSECURE 00000003) (VPN+ Guide 3).

F-Secure, *F-Secure VPN+ 4.1* (1999) (from FSECURE 00000006) (VPN+ 4.1 Guide 6).

F-Secure, *F-Secure SSH* (1996) (from FSECURE 00000006) (F-Secure SSH 6).

F-Secure, *F-Secure SSH 2.0 for Windows NT and 95* (1998) (from FSECURE 00000006) (F-Secure SSH 2.0 Guide 6).

F-Secure, *F-Secure SSH User's & Administrator's Guide* (Sep. 1998) (from FSECURE 00000009) (SSH Guide 9).

F-Secure, *F-Secure SSH 2.0 for Windows NT and 95* (Sep. 1998) (from FSECURE 00000009) (F-Secure SSH 2.0 Guide 9).

F-Secure, *F-Secure VPN+* (Sep. 1998) (from FSECURE 00000009) (VPN+ Guide 9).

F-Secure, *F-Secure Management Tools, Administrator's Guide* (1999) (from FSECURE 00000003) (F-Secure Management Tools).

F-Secure, *F-Secure Desktop, User's Guide* (1997) (from FSECURE 00000009) (FSecure Desktop User's Guide).

SafeNet, Inc., *VPN Policy Manager* (Jan. 2000) (VPN Policy Manager).

F-Secure, *F-Secure VPN+ for Windows NT 4.0* (1998) (from FSECURE 00000009) (FSecure VPN+).

IRE, Inc., *SafeNet/Security Center Technical Reference Addendum* (Jun. 22, 1999) (Safenet Addendum).

IRE, Inc., *System Description for VPN Policy Manager and SafeNet/SoftPK* (Mar. 30, 2000) (VPN Policy Manager System Description).

IRE, Inc., *About SafeNet / VPN Policy Manager* (1999) (About Safenet VPN Policy Manager).

Trusted Information Systems, Inc., *Gauntlet Internet Firewall, Firewall Product Functional Summary* Jul. 22, 1996) (Gauntlet Functional Summary).

Trusted Information Systems, Inc., *Running the Gauntlet Internet Firewall, An Administrator's Guide to Gauntlet Version 3.0* (May 31, 1995) (Running the Gauntlet Internet Firewall).

Ted Harwood, *Windows NT Terminal Server and Citrix Metaframe* (New Riders 1999) (Windows NT Harwood) 79.

Todd W. Mathers and Shawn P. Genoway, *Windows NT Thing Client Solutions: Implementing Terminal Server and Citrix MetaFrame* (Macmillan Technical Publishing 1999) (Windows NT Mathers).

Bernard Aboba et al., *Securing L2TP using IPSEC* (Feb. 2, 1999).  
156. *Finding Your Way Through the VPN Maze* (1999) ("PGP").

Linux FreeS/WAN Overview (1999) (Linux FreeS/WAN Overview).

TimeStep, *The Business Case for Secure VPNs* (1998) ("TimeStep").

WatchGuard Technologies, Inc., *WatchGuard LiveSecurity for MSS Powerpoint* (Feb. 14, 2000).

WatchGuard Technologies, Inc., *MSS Version 2.5, Add-On for WatchGuard SOHO Release Notes* (Jul. 21, 2000).

WatchGuard Technologies, Inc., *MSS Firewall Specifications* (1999).

WatchGuard Technologies, Inc., *Request for Information, Security Services* (2000).

WatchGuard Technologies, Inc., *Protecting the Internet Distributed Enterprise, White Paper* (Feb. 2000).

Air Force Research Laboratory, *Statement of Work for Information Assurance System Architecture and Integration, PR No. N-8-6106* (Contract No. F30602-98-C-0012) (Jan. 29, 1998).

Technologies, Inc., *WatchGuard Firebox System Powerpoint* (2000).

GTE Internetworking & BBN Technologies DARPA Information Assurance Program Integrated Feasibility Demonstration 1FD 1.2 Report, Rev. 1.0 (Sep. 21, 1998).

BBN Information Assurance Contract, *TIS Labs Monthly Status Report* (Mar. 16-Apr. 30, 1998).

DARPA, *Dynamic Virtual Private Network (VPN) Powerpoint*.

GTE Internetworking, *Contractor's Program Progress Report* (Mar. 16-Apr. 30, 1998).

Darrell Kindred, *Dynamic Virtual Private Networks (DVPN) Countermeasure Characterization* (Jan. 30, 2001).

*Virtual Private Networking Countermeasure Characterization* (Mar. 30, 2000).

*Virtual Private Network Demonstration* (Mar. 21, 1998).

Information Assurance/NAI Labs, *Dynamic Virtual Private Networks (VPNs) and Integrated Security Management* (2000).

Information Assurance/NAI Labs, *Create/Add DVPN Enclave* (2000).

NAI Labs, *IFE 3.1 Integration Demo* (2000).

Information Assurance, *Science Fair Agenda* (2000).

Darrell Kindred et al., *Proposed Threads for IFE 3.1* (Jan. 13, 2000).

*IFE 3.1 Technology Dependencies* (2000).



(56)

## References Cited

## OTHER PUBLICATIONS

- IFE 3.1 Topology* (Feb. 9, 2000).
- Information Assurance, *Information Assurance Integration: IFE 3.1, Hypothesis & Thread Development* Jan. 10-11, 2000).
- Information Assurance/NAI Labs, *Dynamic Virtual Private Networks Presentation* (2000).
- Information Assurance/NAI Labs, *Dynamic Virtual Private Networks Presentation v.2* (2000).
- Information Assurance/NAI Labs, *Dynamic Virtual Private Networks Presentation v.3* (2000).
- T. Braun et al., *Virtual Private Network Architecture*, Charging and Accounting Technology for the Internet (Aug. 1, 1999) (VPNA).
- Network Associates Products—*PGP Total Network Security Suite, Dynamic Virtual Private Networks* (1999).
- Microsoft Corporation, *Microsoft Proxy Server 2.0* (1997) (Proxy Server 2.0, Microsoft Prior Art VPN Technology).
- David Johnson et. al., *A Guide to Microsoft Proxy Server 2.0* (1999) (Johnson, Microsoft Prior Art VPN Technology).
- Microsoft Corporation, *Setting Server Parameters* (1997 (copied from Proxy Server 2.0 CD labeled MSFTVX00157288) (Setting Server Parameters, Microsoft Prior Art VPN Technology).
- Kevin Schuler, *Microsoft Proxy Server 2* (1998) (Schuler, Microsoft Prior Art VPN Technology).
- Erik Rozell et. al., *MCSE Proxy Server 2 Study Guide* (1998) (Rozell, Microsoft Prior 15 Art VPN Technology).
- M. Shane Stigler & Mark A Linsenhardt, *IIS 4 and Proxy Server 2* (1999) (Stigler, Microsoft Prior ART VPN Technology).
- David G. Schaer, *MCSE Test Success: Proxy Server 2* (1998) (Schaer, Microsoft Prior Art VPN Technology).
- John Savill, *The Windows NT and Windows 2000 Answer Book* (1999) (Savill, Microsoft Prior ART VPN Technology).
- File History for U.S. Appl. No. 09/653,201, Applicant(s): Whittle Bryan, et al., Filing Date Aug. 31, 2000.
- AutoSOCKS v2. 1*, Datasheet, <http://web.archive.org/web/19970212013409/www.aventail.com/prod/autoskds.html>.
- Ran Atkinson, *Use of DNS to Distribute Keys*, Sep. 7, 1993, [http://ops.ietf.org/lists/namedroppers/namedroppers\\_1\\_99x/msg00945.html](http://ops.ietf.org/lists/namedroppers/namedroppers_1_99x/msg00945.html).
- FirstVPN Enterprise Networks, Overview.
- Chapter 1: Introduction to Firewall Technology, Administration Guide; Dec. 19, 2007, [http://www.books24x7.com/book/id\\_762/viewer\\_.asp?bookid=762&chunked=41065062](http://www.books24x7.com/book/id_762/viewer_.asp?bookid=762&chunked=41065062).
- The TLS Protocol Version 1.0; Jan. 1999; p. 65 of 71.
- Elizabeth D. Zwicky, et al., *Building Internet Firewalls*, 2nd Ed. Virtual Private Networks—Assured Digital Incorporated—ADI 4500; <http://web.archive.org/web/19990224050035/www.assured-digital.com/products/prodvpn/adia4500.htm>.
- Accessware—The Third Wave in Network Security, Conclave from Internet Dynamics; <http://web.archive.org/web/11980210013830/interdyn.com/Accessware.html>.
- Extended System Press Release, Sep. 2, 1997; *Extended VPN Uses The Internet to Create Virtual Private Networks*, [www.extendedsystems.com](http://www.extendedsystems.com).
- Socks Version 5; Executive Summary; <http://web.archive.org/web/199970620031945/www.aventail.com/educate/whitepaper/sockswp.html>.
- Internet Dynamics First to Ship Integrated Security Solutions for Enterprise Intranets and Extranets; Sep. 15, 1997; <http://web.archive.org/web/19980210014150/interdyn.com>.
- Emails from various individuals to Linux IPsec re: DNS-LDAP Splicing.
- Fasbender, A., et al., Variable and Scalable Security: Protection of Location Information in Mobile IP, IEEE VTS, 46th, 1996, 5 pp.
- David Kosiur, “Building and Managing Virtual Private Networks” (1998).
- Request for Inter Partes Reexamination of Patent No. 6,502,135, dated Nov. 25, 2009.
- Request for Inter Partes Reexamination of Patent No. 7,188,180, dated Nov. 25, 2009.
- Yuan Dong Feng, “A novel scheme combining interleaving technique with cipher in Rayleigh fading channels,” Proceedings of the International Conference on Communication technology, 2:S47-02-1-S47-02-4 (1998).
- Davies and Price, edited by Tadahiro Uezono, “Network Security”, Japan, Nikkei McGraw-Hill, Dec. 5, 1958, First Edition, first copy, p. 102-108.
- Davies et al., “An Introduction to Data Security in Teleprocessing and Electronic Funds Transfer,” Security for Computer Networks, Second Edition, pp. 98-101 (1989).
- Baumgartner et al., “Differentiated Services: A New Approach for Quality of Service in the Internet,” International Conference on High Performance Networking, 255-273 (1998).
- Chapman et al., “Domain Name System (DNS),” 278-296 (1995).
- Davila et al., “Implementation of Virtual Private Networks at the Transport Layer,” M. Mambo, Y. Zheng (Eds), Information Security (Second International) Workshop, ISW’ 99. Lecture Notes in Computer Science (LNCS), vol. 1729; 85-102 (1999).
- De Raadt et al., “Cryptography in OpenBSD,” 10 pages (1999).
- Eastlake, “Domain Name System Security Extensions,” Internet Citation, Retrieved from the Internet: URL: <ftp://ftp.inet.no/pub/ietf/internet-drafts/draft-ietf-dnssec-secext2-05.txt> (1998).
- Gunter et al., “An Architecture for Managing QoS-Enabled VRNs Over the Internet,” Proceedings 24th Conference on Local Computer Networks. LCN’ 99 IEEE Comput. Soc Los Alamitos, CA, pp. 122-131 (1999).
- Shimizu, “Special Feature: Mastering the Internet with Windows 2000”, Internet Magazine, 63:296-307 (2000).
- Stallings, “Cryptography and Network Security,” Principals and Practice, 2nd Edition, pp. 399-440 (1999).
- Takata, “U.S. Vendors Take Serious Action to Act Against Crackers—A Tracking Tool and a Highly Safe DNS Software are Released”, Nikkei Communications, 257:87(1997).
- Wells, Email (Lancaster1be@mail.msn.com), Subject: “Security Icon,” (1998).
- Microsoft Corporation’s Fifth Amended Invalidity Contentions dated Sep. 18, 2009, *VirnetX Inc. and Science Applications International Corp. v. Microsoft Corporation* and invalidity claim charts for U.S. Patent Nos. 7,188,180 and 6,839,759.
- The IPSEC Protocol as described in Atkinson, et al., “Security Architecture for the Internet Protocol,” Network Working Group, RFC 2401 (Nov. 1998) (“RFC 2401”); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- S. Kent and R. Atkinson, “IP Authentication Header,” RFC 2402 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- C. Madson and R. Glenn, “The Use of HMAC-MD5-96 within ESP and AH,” RFC 2403 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- C. Madson and R. Glenn, “The Use HMAC-SHA-1-96 within ESP and AH,” RFC 2404 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- C. Madson and N. Doraswamy, “The ESP DES-CBC Cipher Algorithm With Explicit IV”, RFC 2405 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- S. Kent and R. Atkinson, “IP Encapsulating Security Payload (ESP),” RFC 2406 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- Derrell Piper, “The Internet IP Security Domain of Interpretation for ISAKMP,” RFC 2407 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- Douglas Maughan, et al, “Internet Security Association and Key Management Protocol (ISAKMP),” RFC 2408 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).

(56)

## References Cited

## OTHER PUBLICATIONS

- D. Harkins and D. Carrell, "The Internet Key Exchange (IKE)," RFC 2409 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- R. Glenn and S. Kent, "The NULL Encryption Algorithm and Its Use With IPsec." RFC 2410 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- R. Thayer, et al., "IP Security Document Roadmap," RFC 2411 (Nov. 1998); [http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu\\_eng.html](http://web.archive.org/web/19991007070353/http://www.imib.med.tu-dresden.de/imib/Internet/Literatur/ipsec-docu_eng.html).
- Hilarie K. Orman, "The OAKLEY Key Determination Protocol," RFC 2412 (Nov. 1998) in combination with J.M. Galvin, "Public Key Distribution with Secure DNS," Proceedings of the Sixth USENIX UNIX Security Symposium, San Jose California (Jul. 1996) ("Galvin").
- DNS-related correspondence dated Sep. 7, 1993 to Sep. 20, 1993. (Pre KX, KX Records).
- Aventail Corp., "AutoSOCKS v. 2.1 Datasheet," available at <http://www.archive.org/web/19970212013409/www.aventail.com/prod/autosk2ds.html> (1997) (AutoSOCKS, Aventail).
- Aventail Corp., "Socks Version 5," Aventail Whitepaper, available at [http://web.archive.org/web/19970620030312/www.aventail.com/educate/whitepaper/soc\\_kswp.html](http://web.archive.org/web/19970620030312/www.aventail.com/educate/whitepaper/soc_kswp.html) (1997). (Socks, Aventail).
- Goncalves, et al. *Check Point Fire Wall -1 Administration Guide*, McGraw-Hill Companies (2000). (Goncalves, Checkpoint FW).
- F-Secure, *F-Secure Evaluation Kit* (May 1999) (FSECURE 00000003) (Evaluation Kit 3).
- F-Secure, *F-Secure Evaluation Kit* (Sep. 1998) (FSECURE 00000009) (Evaluation Kit 9).
- IRE, Inc., *SafeNet/Soft-PK Version 4* (Mar. 28, 2000) (Soft-PK Version 4).
- IRE/SafeNet Inc., *VPN Technologies Overview* (Mar. 28, 2000) (Safenet VPN Overview).
- IRE, Inc., *SafeNet/VPN Policy Manager Quick Start Guide Version 1* (1999) (SafeNet VPN Policy Manager).
- PCT International Search Report for related PCT Application No. PCT/US01/13261, 8 pages.
- PCT International Search Report for related PCT Application No. PCT/US99/25323, 3 pages.
- PCT International Search Report for related PCT Application No. PCT/US99/25325, 3 pages.
- Deposition Transcript for Gary Tomlinson dated Feb. 27, 2009.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 8, 2010, 8:45 AM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 8, 2010, 1:30 PM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 9, 2010, 9:00 AM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 9, 2010, 1:30 PM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 10, 2010, 9:00 AM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 10, 2010, 1:00 PM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 11, 2010, 9:00 AM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 11, 2010, 1:30 PM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 12, 2010, 9:00 AM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 12, 2010, 1:15 PM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 15, 2010, 9:00 AM.
- Trial Transcript, *VirnetX vs. Microsoft Corporation* dated Mar. 15, 2010, 12:35 PM.
- European Search Report dated Jan. 24, 2011 from corresponding European Application No. 10011949.4.
- European Search Report dated Mar. 17, 2011 from corresponding European Application No. 10184502.2.
- Hollenbeck et al., Registry Registrar Protocol (RRP) Version 1.1.0; Internet Engineering Task Force, 34 pages. (1999).
- Tannenbaum, "Computer Networks," pp. 202-219 (1996).
- Defendants' Preliminary Joint Invalidation Contentions dated Jul. 1, 2011.
- Appendix B: DNS References to Defendants' Preliminary Joint Invalidation Contentions dated Jul. 1, 2011.
- Appendix A to Defendants' Preliminary Joint Invalidation Contentions dated Jul. 1, 2011.
- Exhibit 1, IETF RFC 2065: *Domain Name System Security Extensions*; Published Jan. 1997' vs. *Claims of the '211 Patent*.
- Exhibit 2, IETF RFC 2065: *Domain Name System Security Extensions*; Published Jan. 1997' vs. *Claims of the '504 Patent*.
- Exhibit 3, *RFC 2543 vs. Claims of the '135 Patent*.
- Exhibit 4, *RFC 2543 vs. Claims of the '211 Patent*.
- Exhibit 6, *SIP Draft v.2 vs. Claims of the '135 Patent*.
- Exhibit 7, *SIP Draft v.2 vs. Claims of the '211 Patent*.
- Exhibit 8, *SIP Draft v.2 vs. Claims of the '504 Patent*.
- Exhibit 9, *H.323 vs. Claims of the '135 Patent*.
- Exhibit 10, *H.323 vs. Claims of the '211 Patent*.
- Exhibit 11, *H.323 vs. Claims of the '504 Patent*.
- Exhibit 12, *SSL 3.0 vs. Claims of the '135 Patent*.
- Exhibit 13, *SSL 3.0 vs. Claims of the '211 Patent*.
- Exhibit 15, *RFC 2487 vs. Claims of the '135 Patent*.
- Exhibit 16, *RFC 2487 vs. Claims of the '211 Patent*.
- Exhibit 17, *RFC 2487 vs. Claims of the '504 Patent*.
- Exhibit 18, *RFC 2595 vs. Claims of the '135 Patent*.
- Exhibit 19, *RFC 2595 vs. Claims of the '211 Patent*.
- Exhibit 20, *RFC 2595 vs. Claims of the '504 Patent*.
- Exhibit 21, *iPass vs. Claims of the '135 Patent*.
- Exhibit 22, *iPASS vs. Claims of the '211 Patent*.
- Exhibit 23, *iPASS vs. Claims of the '504 Patent*.
- Exhibit 24, "US '034" vs. *Claims of the '135 Patent*.
- Exhibit 25, US Patent No. 6,453,034 ("US '034") vs. *Claims of the '211 Patent*.
- Exhibit 26, US Patent No. 6,453,034 ("US '034") vs. *Claims of the '504 Patent*.
- Exhibit 27, *US '287 vs. Claims of the '135 Patent*.
- Exhibit 28, *US '287 vs. Claims of the '211 Patent*.
- Exhibit 29, *US '287 vs. Claims of the '504 Patent*.
- Exhibit 30, *Overview of Access VPNs vs. Claims of the '135 Patent*.
- Exhibit 31, *Overview of Access VPNs vs. Claims of the '211 Patent*.
- Exhibit 32, *Overview of Access VPNs vs. Claims of the '504 Patent*.
- Exhibit 34, *RFC 1928 vs. Claims of the '135 Patent*.
- Exhibit 35, *RFC 1928 vs. Claims of the '211 Patent*.
- Exhibit 36, *RFC 1928 vs. Claims of the '504 Patent*.
- Exhibit 37, *RFC 2661 vs. Claims of the '135 Patent*.
- Exhibit 38, *RFC 2661 vs. Claims of the '211 Patent*.
- Exhibit 39, *RFC 2661 vs. Claims of the '504 Patent*.
- Exhibit 40, *SecureConnect vs. Claims of the '135 Patent*.
- Exhibit 41, *SecureConnect vs. Claims of the '211 Patent*.
- Exhibit 42, *SecureConnect vs. Claims of the '504 Patent*.
- Exhibit 43, *SFS-HTTP vs. Claims of the '135 Patent*.
- Exhibit 44, *SFS-HTTP vs. Claims of the '211 Patent*.
- Exhibit 45, *SFS-HTTP vs. Claims of the '504 Patent*.
- Exhibit 46, *US '883 vs. Claims of the '135 Patent*.
- Exhibit 47, *US '883 vs. Claims of the '211 Patent*.
- Exhibit 48, *US '883 vs. Claims of the '504 Patent*.
- Exhibit 49, *US '132 vs. Claims of the '135 Patent*.
- Exhibit 50, *US '132 vs. Claims of the '211 Patent*.
- Exhibit 51, *US '132 vs. Claims of the '504 Patent*.
- Exhibit 52, *US '213 vs. Claims of the '135 Patent*.
- Exhibit 53, *US '213 vs. Claims of the '211 Patent*.
- Exhibit 54, *US '213 vs. Claims of the '504 Patent*.
- Exhibit 55, *B&M VPNs vs. Claims of the '135 Patent*.
- Exhibit 56, *B&M VPNs vs. Claims of the '211 Patent*.
- Exhibit 57, *B&M VPNs vs. Claims of the '504 Patent*.
- Exhibit 58, *BorderManager vs. Claims of the '135 Patent*.
- Exhibit 59, *BorderManager vs. Claims of the '211 Patent*.

(56)

## References Cited

## OTHER PUBLICATIONS

- Exhibit 60, *BorderManager* vs. *Claims of the '504 Patent*.  
 Exhibit 61, *Prestige 128 Plus* vs. *Claims of the '135 Patent*.  
 Exhibit 62, *Prestige 128 Plus* vs. *Claims of the '211 Patent*.  
 Exhibit 63, *Prestige 128 Plus* vs. *Claims of the '504 Patent*.  
 Exhibit 64, *RFC 2401* vs. *Claims of the '135 Patent*.  
 Exhibit 65, *RFC 2401* vs. *Claims of the '211 Patent*.  
 Exhibit 66, *RFC 2401* vs. *Claims of the '504 Patent*.  
 Exhibit 67, *RFC 2486* vs. *Claims of the '135 Patent*.  
 Exhibit 68, *RFC 2486* vs. *Claims of the '211 Patent*.  
 Exhibit 69, *RFC 2486* vs. *Claims of the '504 Patent*.  
 Exhibit 70, *Understanding IPSec* vs. *Claims of the '135 Patent*.  
 Exhibit 71, *Understanding IPSec* vs. *Claims of the '211 Patent*.  
 Exhibit 72, *Understanding IPSec* vs. *Claims of the '504 Patent*.  
 Exhibit 73, *US '820* vs. *Claims of the '135 Patent*.  
 Exhibit 74, *US '820* vs. *Claims of the '211 Patent*.  
 Exhibit 75, *US '820* vs. *Claims of the '504 Patent*.  
 Exhibit 76, *US '019* vs. *Claims of the '211 Patent*.  
 Exhibit 77, *US '019* vs. *Claims of the '504 Patent*.  
 Exhibit 78, *US '049* vs. *Claims of the '135 Patent*.  
 Exhibit 79, *US '049* vs. *Claims of the '211 Patent*.  
 Exhibit 80, *US '049* vs. *Claims of the '504 Patent*.  
 Exhibit 81, *US '748* vs. *Claims of the '135 Patent*.  
 Exhibit 82, *US '261* vs. *Claims of the '135 Patent*.  
 Exhibit 83, *US '261* vs. *Claims of the '211 Patent*.  
 Exhibit 84, *US '261* vs. *Claims of the '504 Patent*.  
 Exhibit 85, *US '900* vs. *Claims of the '135 Patent*.  
 Exhibit 86, *US '900* vs. *Claims of the '211 Patent*.  
 Exhibit 87, *US '900* vs. *Claims of the '504 Patent*.  
 Exhibit 88, *US '671* vs. *Claims of the '135 Patent*.  
 Exhibit 89, *US '671* vs. *Claims of the '211 Patent*.  
 Exhibit 90, *US '671* vs. *Claims of the '504 Patent*.  
 Exhibit 91, *JP '704* vs. *Claims of the '135 Patent*.  
 Exhibit 92, *JP '704* vs. *Claims of the '211 Patent*.  
 Exhibit 93, *JP '704* vs. *Claims of the '504 Patent*.  
 Exhibit 94, *GB '841* vs. *Claims of the '135 Patent*.  
 Exhibit 95, *GB '841* vs. *Claims of the '211 Patent*.  
 Exhibit 96, *GB '841* vs. *Claims of the '504 Patent*.  
 Exhibit 97, *US '318* vs. *Claims of the '135 Patent*.  
 Exhibit 98, *US '318* vs. *Claims of the '211 Patent*.  
 Exhibit 99, *US '318* vs. *Claims of the '504 Patent*.  
 Exhibit 100, *VPNA/LAN* vs. *Claims of the '135 Patent*.  
 Exhibit 101, *Nikkei* vs. *Claims of the '135 Patent*.  
 Exhibit 102, *Nikkei* vs. *Claims of the '211 Patent*.  
 Exhibit 103, *Nikkei* vs. *Claims of the '504 Patent*.  
 Exhibit 104, *Special Anthology* vs. *Claims of the '135 Patent*.  
 Exhibit 105, *Omron* vs. *Claims of the '135 Patent*.  
 Exhibit 106, *Gauntlet System* vs. *Claims of the '135 Patent*.  
 Exhibit 107, *Gauntlet System* vs. *Claims of the '151 Patent*.  
 Exhibit 108, *Gauntlet System* vs. *Claims of the '180 Patent*.  
 Exhibit 109, *Gauntlet System* vs. *Claims of the '211 Patent*.  
 Exhibit 110, *Gauntlet System* vs. *Claims of the '504 Patent*.  
 Exhibit 111, *Gauntlet System* vs. *Claims of the '759 Patent*.  
 Exhibit 112, *IntraPort System* vs. *Claims of the '135 Patent*.  
 Exhibit 113, *IntraPort System* vs. *Claims of the '151 Patent*.  
 Exhibit 114, *IntraPort System* vs. *Claims of the '180 Patent*.  
 Exhibit 115, *IntraPort System* vs. *Claims of the '211 Patent*.  
 Exhibit 116, *IntraPort System* vs. *Claims of the '504 Patent*.  
 Exhibit 117, *IntraPort System* vs. *Claims of the '759 Patent*.  
 Exhibit 118, *Altiga VPN System* vs. *Claims of the '135 Patent*.  
 Exhibit 119, *Altiga VPN System* vs. *Claims of the '151 Patent*.  
 Exhibit 120, *Altiga VPN System* vs. *Claims of the '180 Patent*.  
 Exhibit 121, *Altiga VPN System* vs. *Claims of the '211 Patent*.  
 Exhibit 122, *Altiga VPN System* vs. *Claims of the '504 Patent*.  
 Exhibit 123, *Altiga VPN System* vs. *Claims of the '759 Patent*.  
 Exhibit 124, *Kiuchi* vs. *Claims of the '135 Patent*.  
 Exhibit 125, *Kiuchi* vs. *Claims of the '151 Patent*.  
 Exhibit 126, *Kiuchi* vs. *Claims of the '180 Patent*.  
 Exhibit 127, *Kiuchi* vs. *Claims of the '211 Patent*.  
 Exhibit 128, *Kiuchi* vs. *Claims of the '504 Patent*.  
 Exhibit 129, *Kiuchi* vs. *Claims of the '759 Patent*.  
 Exhibit 130, *Overview of Access VPNs and Tunneling Technologies* (“Overview”) vs. *Claims of the '135 Patent*.  
 Exhibit 131, *Overview of Access VPNs and Tunneling Technologies* (“Overview”) vs. *Claims of the '151 Patent*.  
 Exhibit 132, *Overview of Access VPNs and Tunneling Technologies* (“Overview”) vs. *Claims of the '180 Patent*.  
 Exhibit 133, *Overview of Access VPNs and Tunneling Technologies* (“Overview”) vs. *Claims of the '211 Patent*.  
 Exhibit 134, *Overview of Access VPNs and Tunneling Technologies* (“Overview”) vs. *Claims of the '504 Patent*.  
 Exhibit 135, *Overview* vs. *Claims of the '759 Patent*.  
 Exhibit 136, *RFC 2401* vs. *Claims of the '759 Patent*.  
 Exhibit 137, *Schulzrinne* vs. *Claims of the '135 Patent*.  
 Exhibit 138, *Schulzrinne* vs. *Claims of the '151 Patent*.  
 Exhibit 139, *Schulzrinne* vs. *Claims of the '180 Patent*.  
 Exhibit 140, *Schulzrinne* vs. *Claims of the '211 Patent*.  
 Exhibit 141, *Schulzrinne* vs. *Claims of the '504 Patent*.  
 Exhibit 142, *Schulzrinne* vs. *Claims of the '759 Patent*.  
 Exhibit 143, *Solana* vs. *Claims of the '135 Patent*.  
 Exhibit 144, *Solana* vs. *Claims of the '151 Patent*.  
 Exhibit 145, *Solana* vs. *Claims of the '180 Patent*.  
 Exhibit 146, *Solana* vs. *Claims of the '211 Patent*.  
 Exhibit 147, *Solana* vs. *Claims of the '504 Patent*.  
 Exhibit 148, *Solana* vs. *Claims of the '759 Patent*.  
 Exhibit 149, *Atkinson* vs. *Claims of the '135 Patent*.  
 Exhibit 150, *Atkinson* vs. *Claims of the '151 Patent*.  
 Exhibit 151, *Atkinson* vs. *Claims of the '180 Patent*.  
 Exhibit 152, *Atkinson* vs. *Claims of the '211 Patent*.  
 Exhibit 153, *Atkinson* vs. *Claims of the '504 Patent*.  
 Exhibit 154, *Atkinson* vs. *Claims of the '759 Patent*.  
 Exhibit 155, *Marino* vs. *Claims of the '135 Patent*.  
 Exhibit 156, *Marino* vs. *Claims of the '151 Patent*.  
 Exhibit 157, *Marino* vs. *Claims of the '180 Patent*.  
 Exhibit 158, *Marino* vs. *Claims of the '211 Patent*.  
 Exhibit 159, *Marino* vs. *Claims of the '504 Patent*.  
 Exhibit 160, *Marino* vs. *Claims of the '759 Patent*.  
 Exhibit 161, *Aziz ('646)* vs. *Claims of the '759 Patent*.  
 Exhibit 162, *Wesinger* vs. *Claims of the '135 Patent*.  
 Exhibit 163, *Wesinger* vs. *Claims of the '151 Patent*.  
 Exhibit 164, *Wesinger* vs. *Claims of the '180 Patent*.  
 Exhibit 165, *Wesinger* vs. *Claims of the '211 Patent*.  
 Exhibit 166, *Wesinger* vs. *Claims of the '504 Patent*.  
 Exhibit 167, *Wesinger* vs. *Claims of the '759 Patent*.  
 Exhibit 168, *Aziz ('234)* vs. *Claims of the '135 Patent*.  
 Exhibit 169, *Aziz (234)* vs. *Claims of the '151 Patent*.  
 Exhibit 170, *Aziz ('234)* vs. *Claims of the '180 Patent*.  
 Exhibit 171, *Aziz ('234)* vs. *Claims of the '211 Patent*.  
 Exhibit 172, *Aziz (234)* vs. *Claims of the '504 Patent*.  
 Exhibit 173, *Aziz ('234)* vs. *Claims of the '759 Patent*.  
 Exhibit 174, *Schneider* vs. *Claims of the '759 Patent*.  
 Exhibit 175, *Valencia* vs. *Claims of the '135 Patent*.  
 Exhibit 176, *Valencia* vs. *Claims of the '151 Patent*.  
 Exhibit 177, *Valencia* vs. *Claims of the '180 Patent*.  
 Exhibit 178, *Valencia* vs. *Claims of the '211 Patent*.  
 Exhibit 179, *Valencia* vs. *Claims of the '504 Patent*.  
 Exhibit 180, *RFC 2401* in Combination with U.S. Patent No. 6,496,867 vs. *Claims of the '180 Patent*.  
 Exhibit 181, *Davison* vs. *Claims of the '135 Patent*.  
 Exhibit 182, *Davison* vs. *Claims of the '151 Patent*.  
 Exhibit 183, *Davison* vs. *Claims of the '180 Patent*.  
 Exhibit 184, *Davison* vs. *Claims of the '211 Patent*.  
 Exhibit 185, *Davison* vs. *Claims of the '504 Patent*.  
 Exhibit 186, *Davison* vs. *Claims of the '759 Patent*.  
 Exhibit 187, *AutoSOCKS v2.1* vs. *Claims of the '135 Patent*.  
 Exhibit 188, *AutoSOCKS v2.1* vs. *Claims of the '151 Patent*.  
 Exhibit 189, *AutoSOCKS v2.1 Administrator's Guide* vs. *Claims of the '180 Patent*.  
 Exhibit 190, *AutoSOCKS* vs. *Claims of the '759 Patent*.  
 Exhibit 191, *Aventail Connect 3.01/2.51* vs. *Claims of the '135 Patent*.  
 Exhibit 192, *Aventail Connect v3.01/2.51* vs. *Claims of the '151 Patent*.

(56)

## References Cited

## OTHER PUBLICATIONS

Exhibit 193, *Aventail Connect 3.01/2.51 vs. Claims of the '180 Patent.*

Exhibit 194, *Aventail Connect 3.01/2.5 vs. Claims of the '759 Patent.*

Exhibit 195, *Aventail Connect 3.1/2.6 Administrator's Guide vs. Claims of the '135 Patent.*

Exhibit 196, *Aventail Connect 3.1/2.6 Administrator's Guide vs. Claims of the '151 Patent.*

Exhibit 197, *Aventail Connect 3.1/2.6 vs. Claims of the '180 Patent.*

Exhibit 198, *Aventail Connect 3.1/2.6 vs. Claims of the '759 Patent.*

Exhibit 199, *BinGO! User's User's Guide/Extended Features Reference vs. Claims of the '151 Patent.*

Exhibit 200, *BinGO! User's User's Guide/Extended Features Reference vs. Claims of the '135 Patent.*

Exhibit 201, *BinGO! vs. Claims of the '180 Patent.*

Exhibit 202, *BinGO! vs. Claims of the '759 Patent.*

Exhibit 203, *Broadband Forum Technical Report TR-025 (Issue 1.0/5.0) vs. Claims of the '135 Patent.*

Exhibit 204, *Domain Name System (DNS) Security vs. Claims of the '211 Patent.*

Exhibit 205, *Domain Name System (DNS) Security vs. Claims of the '504 Patent.*

Exhibit 206, *RFC 2230, Key Exchange Delegation Record for the DNS vs. Claims of the '211 Patent.*

Exhibit 207, *RFC 2230, Key Exchange Delegation Record for the DNS vs. Claims of the '504 Patent.*

Exhibit 208, *RFC 2538, Storing Certificates in the Domain Name System (DNS) vs. Claims of the '211 Patent.*

Exhibit 209, *RFC 2538, Storing Certificates in the Domain Name System (DNS) vs. Claims of the '504 Patent.*

Exhibit 210, *IETF RFC 2065: Domain Name System Security Extensions; Published Jan. 1997 vs. Claims of the '504 Patent.*

Exhibit 211, *IETF RFC 2065: Domain Name System Security Extensions; Published Jan. 1997 vs. Claims of the '211 Patent.*

Exhibit 212, *RFC 2486, RFC 2661, RFC 2401, and Internet-Draft, "Secure Remote Access with L2TP" vs. Claims of the '135 Patent.*

Exhibit 213, *U.S. Patent No. 7,100,195 in Combination with RFC 2401 and U.S. Patent No. 6,496,867 vs. Claims of the '135 Patent.*

Exhibit 214, *U.S. Patent No. 7,100,195 in Combination with RFC 2401 and U.S. Patent No. 6,496,867 vs. Claims of the '151 Patent.*

Exhibit 215, *U.S. Patent No. 6,643,701 vs. Claims of the '135 Patent.*

Exhibit 216, *U.S. Patent No. 6,643,701 vs. Claims of the '151 Patent.*

Exhibit 217, *U.S. Patent No. 6,496,867 in Combination with RFC 2401 vs. Claims of the '151 Patent.*

Exhibit 218, *U.S. Patent No. 6,496,867 in Combination with RFC 2401 vs. Claims of the '135 Patent.*

Exhibit 219, *U.S. Patent No. 6,496,867 vs. Claims of the '211 Patent.*

Exhibit 220, *U.S. Patent No. 6,496,867 vs. Claims of the '504 Patent.*

Exhibit 221, *RFC 2486, RFC 2661, RFC 2401, and Internet-Draft, "Secure Remote Access with L2TP" vs. Claims of the '151 Patent.*

Exhibit 222, *U.S. Patent No. 6,557,037 vs. Claims of the '211 Patent.*

Exhibit 223, *U.S. Patent No. 6,557,037 vs. Claims of the '504 Patent.*

Exhibit 224, *RFC 2230, Key Exchange Delegation Record for the DNS vs. Claims of the '135 Patent.*

Exhibit 225, *RFC 2230, Key Exchange Delegation Record for the DNS vs. Claims of the '151 Patent.*

Exhibit Cisco-1, *Cisco's Prior Art Systems vs. Claims of the '135 Patent.*

Exhibit Cisco-2, *Cisco's Prior Art Systems vs. Claims of the '151 Patent.*

Exhibit Cisco-3, *Cisco's Prior Art Systems vs. Claims of the '180 Patent.*

Exhibit Cisco-4, *Cisco's Prior Art Systems vs. Claims of the '211 Patent.*

Exhibit Cisco-5, *Cisco's Prior Art Systems vs. Claims of the '504 Patent.*

Exhibit Cisco-6, *Cisco's Prior Art Systems vs. Claims of the '759 Patent.*

Exhibit Cisco-7, *Cisco's Prior Art PIX System vs. Claims of the '759 Patent.*

Exhibit A: U.S. Patent No. 6,502,135.

Exhibit A: U.S. Patent No. 7,490,151.

Exhibit B: Certificate of Service to Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 6,502,135).

Exhibit B: Certificate of Service to Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 7,490,151).

Exhibit B-1: File History of U.S. Patent 6,502,135.

Exhibit B-2: Reexamination Record No. 95/001,269.

Exhibit C1: Claim Chart—Aventail Connect v3.1 (Patent No. 6,502,135).

Exhibit C2: Claim Chart Aventail Connect V3.01 (Patent No. 6,502,135).

Exhibit C-1: U.S. Patent No. 7,010,604.

Exhibit C2: Claim Chart Aventail Autosocks (Patent No. 7,490,151).

Exhibit C1: Claim Chart Aventail Connect v3.01 (Patent No. 7,490,151).

Exhibit C-2: U.S. Appl. No. 60/106,261.

Exhibit C3: Claim Chart Aventail AutoSOCKS (Patent No. 6,502,135).

Exhibit C3: Claim Chart BinGO (Patent No. 7,490,151).

Exhibit C-3: U.S. Appl. No. 60/137,704.

Exhibit C4: Claim Chart Wang (Patent No. 6,502,135).

Exhibit C4: Claim Chart Beser (Patent No. 7,490,151).

Exhibit C5: Claim Chart Beser (Patent No. 6,502,135).

Exhibit C5: Claim Chart Wang (Patent No. 7,490,151).

Exhibit C6: Claim Chart BinGO (Patent No. 6,502,135).

Exhibit D: Memorandum Opinion in *VimetX v. Microsoft*.

Exhibit D-1: Takahiro Kiuchi and Shigekoto Kaihara, "C-HTTP—The Development of a Secure, Closed HPPT-Based Network on the Internet," Published in the Proceedings of SNDSS 1996.

Exhibit D-10: D.E. Denning and G.M. Sacco, "Time-stamps in Key Distribution Protocols," *Communications of the ACM*, vol. 24, N.8, pp. 533-536. Aug. 1981.

Exhibit D-11: C.I. Dalton and J.F. Griffin, "Applying Military Grade Security to the Internet," Proceedings of the 8th Joint European Networking Conference (JENC 8), (May 12-15, 1997).

Exhibit D-12: Steven M. Bellovin and Michael Merritt, "Encrypted Key Exchange: Password-Based protocols Secure against Dictionary Attacks," 1992 IEEE Symposium on Security and Privacy (1992).

Exhibit D-2: U.S. Pat. No. 5,898,830.

Exhibit D-3: Eduardo Solana and Jürgen Harms, "Flexible Internet Secure Transactions Based on Collaborative Domains," *Security Protocols Workshop 1997*, pp. 37-51.

Exhibit D-4: U.S. Pat. No. 6,119,234.

Exhibit D-5: Jeff Sedayao, "Mosaic Will Kill My Network!"—Studying Network Traffic Patterns of Mosaic Use," in *Electron. Proc. 2nd World Wide Web Conf. '94: Mosaic and the Web*, Chicago, IL, Oct. 1994.

Exhibit D-6: M. Luby Juels and R. Ostrovsky, "Security of Blind Digital Signatures," *Crypto '97, LNCS 1294*, pp. 150-164, Springer-Verlag, Berlin, 1997.

Exhibit D-8: David M. Martin, "A Framework for Local Anonymity in the Internet," Technical Report. Boston University, Boston, MA, USA (Feb. 21, 1998).

Exhibit D-9: U.S. Pat. No. 7,764,231.

Exhibit E-1: Claim Charts Applying Kiuchi and Other References to Claims of the '135 Patent.

Exhibit E1: Declaration of Chris Hopen (Patent No. 6,502,135).

Exhibit E1: Declaration of Chris Hopen (Patent No. 7,490,151).

Exhibit E-2: Claim Charts Applying Wesinger and Other References to Claims of the '135 Patent.

Exhibit E2: Declaration of Michael Fratto (Patent No. 6,502,135).

Exhibit E2: Declaration of Michael Fratto (Patent No. 7,490,151).

Exhibit E-3: Claim Charts Applying Solana and Other References to Claims of the '135 Patent.

Exhibit E3: Declaration of James Chester (Patent No. 6,502,135).

Exhibit E3: Declaration of James Chester (Patent No. 7,490,151).

Exhibit E-4: Claim Charts Applying Aziz and Other References to Claims of the '135 Patent.

Exhibit X1: *Aventail Connect Administrator's Guide v3.1/v2.6.*, pp. 1-20 (1996-1999).

Exhibit X10: U.S. Patent No. 4,885,778.

Exhibit X11: U.S. Patent No. 6,615,357.

(56)

## References Cited

## OTHER PUBLICATIONS

- Exhibit X2: Aventail Connect Administrator's Guide v3.01/v2.51., pp. 1-116 (1996-1999).
- Exhibit X3: Aventail AutoSOCKS Administration & User's Guide v2.1., pp. 1-70 (1996-1999).
- Exhibit X4: Reed et al., "Proxies for Anonymous Routine," 12th Annual Computer Security Applications Conference, San Diego, CA, Dec. 9-13, pp. 1-10 (1996).
- Exhibit X5 Wang, The Broadband Forum Technical Report, "TR-025—Core Network Architecture Recommendations for Access to Legacy Data Networks over ADSL," Issue 1.0; pp. 1-24, v1.0 (1999).
- Exhibit X6: U.S. Patent No. 6,496,867.
- Exhibit X7: BinGO! User's Guide Incorporating by Reference BinGO! Extended Feature Reference.
- Exhibit X7: Kent et al., "Security Architecture for the Internet Protocol," Network Working Group Request for Comments (RFC) 2401, pp. 1-70 (1998).
- Exhibit X8: U.S. Patent No. 6,182,141.
- Exhibit X9: BinGO! User's Guide v1.6 (1999).
- Exhibit Y1: Aventail Extranet Server 3.0 Administrator's Guide.
- Exhibit Y10: Hanks, S., et al., RFC1701, "Generic Routing Encapsulation (GRE)," 1994, Is Accessible at <http://www.ietf.org/rfc/rfc1701.txt>.
- Exhibit Y10: Socolofsky, T. et al., RFC 1180, "A TCP/IP Tutorial," Jan. 1991.
- Exhibit Y11: Simpson, W., editor, RFC 1661, "The Point-to-Point Protocol (PPP)," Jul. 1994.
- Exhibit Y11: Simpson, W., RFC1994, "PPP Challenge Handshake Authentication Protocol (CHAP)," 1996, <http://www.ietf.org/rfc/rfc1994.txt>.
- Exhibit Y12: Meyer, G., RFC 1968, "The PPP Encryption Control Protocol (ECP)," Jun. 1996.
- Exhibit Y12: Perkins, D., RFC1171, "The Point-To-Point Protocol for the Transmission of Multi-Protocol Datagrams over Point-To-Point Links," 1990, Is Accessible at <http://www.ietf.org/rfc/rfc1171.txt>.
- Exhibit Y13: Kummert, H., RFC 2420, "The PPP Triple-DES Encryption Protocol (3DESE)," Sep. 1998.
- Exhibit Y14: Townsley, W.M., et al., RFC 2661, "Layer Two Tunneling Protocol 'L2TP'," Aug. 1999.
- Exhibit Y15: Pall, G.S., RFC 2118, "Microsoft Point-To-Point Encryption (MPPE) Protocol," Mar. 1997.
- Exhibit Y16: Gross, G., et al., RFC 2364, "PPP Over AAL5," Jul. 1998.
- Exhibit Y17: Srisuresh, P., RFC 2663, "IP Network Address Translator (NAT) Terminology and Considerations," Aug. 1999.
- Exhibit Y18: Heinanen, J., RFC 1483, "Multiprotocol Encapsulation over ATM Adaptation Layer 5," Jul. 1993.
- Exhibit Y2: Goldschlag et al., "Hiding Routing Information" (1996).
- Exhibit Y3: U.S. Patent No. 5,950,519.
- Exhibit Y4: Ferguson, P. and Huston, G., "What Is a VPN", The Internet Protocol Journal, vol. 1., No. 1 (Jun. 1998 ("Ferguson")).
- Exhibit Y5: Mockapetris, P., RFC 1034, "Domain Names—Concepts and Facilities," Nov. 1987 ("RFC1034").
- Exhibit Y6: Mockapetris, P., RFC 1035, "Domain Names—Implementation and Specification," Nov. 1987 ("RFC1035").
- Exhibit Y8: Fielding, R., et al., RFC 2068, "Hypertext Transfer Protocol—HTTP/1.1," Jan. 1997
- Exhibit Y8: Woodburn, R.A., et al., RFC1241, "A Scheme for an Internet Encapsulation Protocol: Version 1," 1991.
- Exhibit Y9: Leech, M., et al., RFC 1928, "Socks Protocol Version 5," Mar. 1996.
- Exhibit Y9: Simpson, W., RFC1853, "IP in IP Tunneling," 1995, Is Accessible at <http://www.ietf.org/rfc/rfc1853.txt>.
- Form PTO/SB/42, Listing Each Patent and Printed Publication Relied Upon to Provide a Substantial New Question of Patentability (Patent No. 6,502,135).
- Form PTO/SB/42, Listing Each Patent and Printed Publication Relied Upon to Provide a Substantial New Question of Patentability (Patent No. 7,490,151).
- Request for Inter Partes Reexamination (Patent No. 6,502,135).
- Request for Inter Partes Reexamination Transmittal Form (PTO/SB/58) (Patent No. 6,502,135).
- Request for Inter Partes Reexamination Transmittal Form (PTO/SB/58) (Patent No. 7,490,151).
- Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 6,502,135).
- Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 7,490,151).
- Transmittal Letter (Patent No. 6,502,135).
- Transmittal Letter (Patent No. 7,490,151).
- Joint Claim Construction and Prehearing Statement.
- Exhibit A: Agreed Upon Terms; P.R. 4-3 Joint Claims Construction and Prehearing Statement.
- Exhibit B: Disputed Claim Terms; P.R. 4-3 Joint Claim Construction and Prehearing Statement.
- Exhibit C; VirnetX's Proposed Construction of Claim Terms and Supporting Evidence.
- Exhibit D; Defendants' Intrinsic and Extrinsic Support; P.R. 4-3 Joint Claim Construction and Prehearing Statement.
- File History of U.S. Patent 6,839,759.
- Exhibit B-4; *VirnetX, Inc. v. Microsoft Corp.*, Case No. 6:07-cv-80, Microsoft's Motion for Partial Summary Judgment of Invalidity of U.S. Patent No. 6,839,759 (E.D. Tex. Dec. 18, 2009).
- Exhibit D-2; Kent et al., "Security Architecture for the Internet Protocol," Internet Engineering Task Force, Internet Draft, (Feb. 1998).
- Exhibit D-3; Aziz et al., U.S. Patent 5,548,646 to Aziz et al., "System for Signatureless Transmission and Reception of Data Packets Between Computer Networks," Filed Sep. 15, 1994 and issued Aug. 20, 1996.
- Exhibit D-4; Yinger; U.S. Patent 5,960,204 to Yinger et al., System and Method for Installing Applications on a Computer on an as needed basis, Filed on Oct. 28, 1996 and Issued Sep. 28, 1999.
- Exhibit D-8; Barlow; U.S. Patent 5,204,961 to Barlow, "Computer Network Operating with Multilevel Hierarchical Security with Selectable Common Trust Realms and Corresponding Security Protocols," Filed on Jun. 25, 1990 and Issued Apr. 20, 1993.
- Exhibit D-12; RFC 1122, Braden, "Requirements for Internet Hosts—Communication Layers," RFC 1122 (Oct. 1989).
- Exhibit D-13; RFC 791; Information Sciences Institute, "Internet Protocol," DARPA Internet Program Specification RFC 791 (Sep. 1981).
- Exhibit D-14; Caronni et al., "SKIP—Securing the Internet," 5th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE '96) (Jun. 19-21, 1996).
- Exhibit D-15; Maughan et al., "Internet Security Association and Key Management Protocol (ISAKMP)," IPSEC Work Group Draft (Jul. 26, 1997).
- Exhibit E-1; Claim Charts Applying Kiuchi as a Primary Reference to the '759 Patent.
- Exhibit E-2; Claim Charts Applying Kent as a Primary Reference to the '759 Patent.
- Exhibit E-3; Claim Charts Applying Aziz as a Primary Reference to the '759 Patent.
- Exhibit E-4; Claim Charts Applying Kent in view of Caronni as a Primary Combination of References to the '759 Patent.
- Exhibit D-5; Edwards et al., "High Security Web Servers and Gateways," Computer Networks and ISDN System 29, pp. 927-938 (Sep. 1997).
- Exhibit D-10; Lee et al., "Hypertext Transfer Protocol—HTTP/1.0," RFC 1945 (May 1996).
- Exhibit E-3; Claim Charts Applying Blum to Claims of the '151 Patent.
- Exhibit B-1, File History of U.S. Patent 7,490,151.
- Exhibit E-1, Claim Charts Applying Kiuchi, and Kiuchi and Martin to Claims of the '151 Patent.
- Exhibit E-2, Claim Charts Applying Wesinger, and Wesinger and Martin to Claims of the '151 Patent.
- Exhibit E-4, Claim Charts Applying Aziz and Edwards, and Aziz, Edwards, and Martin to Claims of the '151 Patent.
- Exhibit E-6, Claim Charts Applying Wesinger and Edwards, and Wesinger, Edwards, and Martin to Claims of the '151 Patent.

(56)

## References Cited

## OTHER PUBLICATIONS

*VirnetX Inc., V. Mitel Networks Corp.*; Defendants' Joint Invalidity Contentions.

Exhibit 42, SecureConnect vs. Claims of the '504 Patent.

Exhibit 46, US 883 vs. Claims of the '135 Patent.

Exhibit 49, *Chuah vs. Claims of the '135 Patent.*

Exhibit 50, *Chuah vs. Claims of the '211 Patent.*

Exhibit 51, *Chuah vs. Claims of the '504 Patent.*

Exhibit 52, U.S. '648 vs. Claims of the '135 Patent.

Exhibit 53, U.S. '648 vs. Claims of the '211 Patent.

Exhibit 65, RFC 2401 vs. Claims of the '211 Patent.

Exhibit 67, US '072 vs. Claims of the '135 Patent.

Exhibit 69, RFC 2486' vs. Claims of the '504 Patent.

Exhibit 70 Understanding IPSecl vs. Claims of the '135 Patent.

Exhibit 100, VPN/VLAN vs. Claims of the '135 Patent.

Exhibit 106-A, *Gauntlet System vs. Claims of the '135 Patent.*

Exhibit 109-A, *Gauntlet System vs. Claims of the '211 Patent.*

Exhibit 110-A, *Gauntlet System vs. Claims of the '504 Patent.*

Exhibit 137, *Schulzrinne vs. Claims of the '135 (Final) Patent.*

Exhibit 168, *Aziz vs. Claims of the '135 Patent.*

Exhibit 171, U.S. '234 vs. Claims of the '211 Patent.

Exhibit 172, *Aziz vs. Claims of the '504 Patent.*

Exhibit 200, BinGO! User's Guide/Extended Features Reference vs. Claims of the '135 Patent.

Exhibit 212, RFC 2486, RFC 2661, RFC 2401 and Internet-Draft, Secure Remote Access with L2TP' vs. Claims of the '135 Patent.

Exhibit 218, U.S. Patent No. 6,496,867 in combination with RFC 2401' vs. Claims of the '135 Patent.

Exhibit 228, U.S. 588 vs. Claims of the '211 Patent (Final).

Exhibit 229, U.S. 588 vs. Claims of the '504 Patent (Final).

Exhibit 230, *Microsoft VPN vs. Claims of the '135 Patent (Final).*

Exhibit 231, *Microsoft VPN vs. Claims of the '211 Patent (Final).*

Exhibit XX, *Microsoft VPN vs. Claims of the '504 Patent.*

Exhibit Cisco-1, *Cisco's Prior Art System vs. Claims of the '135 Patent.*

Exhibit Cisco-4, *Cisco's Prior Art System vs. Claims of the '211 Patent.*

Exhibit Cisco-5, *Cisco's Prior Art System vs. Claims of the '504 Patent.*

Exhibit 225, US '037 vs. Claims of the '135 Patent.

Exhibit 226, *ITU-T Standardization Activities vs. Claims of the '135 Patent.*

Exhibit 227, US '393 vs. Claims of the '135 Patent.

Exhibit 233, *The Miller Application vs. Claim 13 of the '135 Patent.*

Exhibit 234, Aventail Connect 3.1/2.6 Administrator's Guide ("*Aventail Connect*") vs. Claims of the '504 Patent.

Exhibit 235, *Microsoft VPN vs. Claims of the '504 Patent.*

Exhibit 1, IETF RFC 2065: Domain Name System Security Extensions; published Jan. 1997 vs. Claims of the '211 Patent.

Exhibit 2, IETF RFC 2065: Domain Name System Security Extensions; published Jan. 1997 vs. Claims of the '504 Patent.

Exhibit 24, U.S. Patent No. 6,453,034 ('034 Patent) vs. Claims of the 135 Patent.

Exhibit 25, U.S. Patent No. 6,453,034 ('034 Patent) vs. Claims of the 211 Patent.

Exhibit 26, U.S. Patent No. 6,453,034 ('034 Patent) vs. Claims of the 504 Patent.

Exhibit 27, U.S. Patent No. 6,223,287 ("287 Patent") vs. Claims of the 135 Patent.

Exhibit 28, U.S. Patent No. 6,223,287 ("287 Patent") vs. Claims of the 211 Patent.

Exhibit 29, U.S. Patent No. 6,223,287 ("287 Patent") vs. Claims of the 504 Patent.

Exhibit 106, Gauntlet System and Gauntlet References vs. Claims of the '135 Patent.

Exhibit 109, Gauntlet System and Gauntlet References vs. Claims of the '211 Patent.

Exhibit 110, Gauntlet System vs. Claims of the '504 Patent.

Exhibit 165, *Wesingervs. Claims of the '211 Patent.*

Exhibit 191, Aventail Connect 3.01/2.51 ("*Aventail Connect*") vs. Claims of the '135 Patent.

Exhibit 195, Aventail Connect 3.1/2.6 Administrator's Guide ("*Aventail Connect*") vs. Claims of the '135 Patent.

Exhibit 205, *Domain Name System (DNS) Security ("DNS Security")* vs. Claims of the '504 Patent.

Exhibit 210, *Lendenmann vs. Claims of the '211 Patent.*

Exhibit 211, *Lendenmann vs. Claims of the '504 Patent.*

Exhibit 215, *Aziz vs. Claims of the '135 Patent.*

Cisco '180, Efiling Acknowledgment.

Exhibit A, U.S. Patent 7,188,180.

Exhibit 81, File History of U.S. Patent 7,188,180.

Exhibit B2, File History of U.S. Appl. No. 09/588,209.

Exhibit 83, File History of Reexamination Control No. 95/001,270, Reexamination of U.S. 7,188,180 requested by Microsoft Corp.

Exhibit D1, "Lendenmann": Rolf Lendenman, Understanding OSF DCE 1.1 for AIX and OS/2, IBM International Technical Support Organization (Oct. 1995).

Exhibit D5, "Schneier": Bruce Schneier, Applied Cryptography (1996).

Exhibit D6, RFC 793; Information Sciences Institute, "Transmission Control Protocol," DARPA Internet Program Specification RFC 793 (Sep. 1981).

Exhibit D7, "Schimpf"; Brian C. Schimpf, "Securing Web Access with DCE," Presented at Network and Distributed System Security (Feb. 10-11, 1997).

Exhibit D8 "Rosenberry"; Ward Rosenberry, David Kenney, and Gerry Fisher, Understanding DCE (1993).

Exhibit D9, Masys; Daniel R. Masys & Dixie B. Baker, "Protecting Clinical Data on Web Client Computers: the PCASSO Approach," Proceedings of the AMIA '98 Annual Symposium, Orlando, Florida (Nov. 7-11, 1998).

Exhibit E1, Claim Charts Applying Lendenmann as a Primary Reference to the '180 Patent.

Exhibit E2, Claim Charts Applying Kiuchi as a Primary Reference to the '180 Patent.

Exhibit E3, Claim Charts Applying Solana as a Primary Reference to the '180 Patent.

Exhibit E4, Claim Charts Applying Schimpf and Rosenberry as a Primary Reference to the '180 Patent.

Request for Inter Partes Reexamination of Patent No. 7,188,180.

Request for Inter Partes Reexamination Transmittal Form No. 7,188,180.

Exhibit A; U.S. Patent 7,921,211 with Terminal Disclaimer.

Exhibit B, Certificate of Service to Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 7,921,211).

Exhibit C1, Claim Chart—USP 7,921,211 Relative to Solana, Alone and in Conjunction with RFC 920, Reed and Beser.

Exhibit C2, Claim Chart—USP 7,921,211 Relative to Solana in view of RFC 2504 and Further in conjunction with RFC 920, Reed, and Beser.

Exhibit C3, Claim Chart—USP 7,921,211 Relative to Provino, Alone and in Conjunction with RFC 920, Reed, and Beser).

Exhibit C4, Claim Chart—USP 7,921,211 Relative to Provino in view of RFC 2230 and Further in Conjunction with RFC 920, Reed and Beser.

Exhibit C5, Claim Chart—USP 7,921,211 Relative to Provino in view of RFC 2504 and in Further Conjunction with RFC 920, Reed and Beser.

Exhibit C6, Claim Chart—USP 7,921,211 Relative to Beser, Alone and in Conjunction with RFC 920, RFC 2401, and Reed.

Exhibit C7, Claim Chart—USP 7,921,211 Relative to RFC 2230, Alone and in Conjunction with RFC 920, RFC 2401, Reed, and Beser.

Exhibit C8, Claim Chart—USP 7,921,211 Relative to RFC 2538, Alone and in Conjunction with RFC 920, RFC 2401, Reed, Beser, and RFC 2065.

Exhibit D1, Asserted Claim and Infringement Contentions by Plaintiff *VirnetX, Inc. In VirnetX, Inc. v. Cisco Systems, Inc., Apple Inc., Aastra Technologies Ltd, NEC Corporation, NEC Corporation of America and Aastra USA, Inc.*, Civ. Act 6:2010cv00417 (E.D. Tex).  
Exhibit D2, Asserted Claims and Infringement Contentions by Plaintiff *VirnetX, Inc.* against *Apple* based on 7,921,211 Patent.

(56)

## References Cited

## OTHER PUBLICATIONS

- Exhibit X1, Solana, E. et al. "Flexible Internet Secure Transactions Based on Collaborative Domains".
- Exhibit X2, U.S. Patent 6,557,037.
- Exhibit X4, Atkinson, R., IETF RFC 2230, "Key Exchange Delegation Record for the DNS" (Nov. 1997).
- Exhibit X6, Kent, et al., IETF RFC 2401, "Security Architecture for the Internet Protocol" (Nov. 1998) Is Accessible at: <http://www.ietf.org/rfc/rfc2401.txt>.
- Exhibit X7, Eastlake, D. et al., IETF RFC 2065, "Domain Name System Security Extensions" (Jan. 1997) Is Accessible at: <http://www.ietf.org/rfc/rfc2065.txt>.
- Exhibit X9, Guttman, E. et al., IETF RFC 2504, "Users' Security Handbook" (Feb. 1999) Is Accessible At: <http://www.ietf.org/rfc/rfc2504.txt>.
- Exhibit Y3, Braden, R., RFC 1123, "Requirements for Internet Hosts—Application and Support," Oct. 1989 ("RFC1123").
- Exhibit Y4, Atkinson, R., RFC 1825, Security Architecture for the Internet Protocol (Aug. 1995) Is Accessible At: <http://www.ietf.org/rfc/rfc1825.txt>.
- Exhibit Y5, Housley, R. et al., RFC 2459, "Internet X.509 Public Key Infrastructure Certificate and CRL Profile" (Jan. 1999) Is accessible At: <http://www.ietf.org/rfc/rfc2459.txt>.
- Exhibit A, U.S. Patent 7,418,504.
- Exhibit B, Certificate of Service to Request for Inter Partes Reexamination Under 35 U.S.C. § 311 (Patent No. 7,418,504).
- Exhibit C1, Claim Chart—USP 7,418,504 Relative to Solana, Alone and in Conjunction with RFC 920, Reed, and Beser.
- Exhibit C2, Claim Chart—USP 7,418,504 Relative to Solana in view of RFC 2504 and Further in Conjunction with RFC 920, Reed, and Beser.
- Exhibit C3, Claim Chart—USP 7,418,504 Relative to Provino, Alone and in Conjunction with RFC 920, Reed, and Beser.
- Exhibit C4, Claim Chart—USP 7,418,504 Relative to Provino in View of RFC 2230 and Further in Conjunction with RFC 920, Reed and Beser.
- Exhibit C5, Claim Chart—USP 7,418,504 Relative to Provino in View of RFC 2504 and in Further Conjunction with RFC 920, Reed, and Beser.
- Exhibit C6, Claim Chart—USP 7,418,504 Relative to Beser, Alone and in Conjunction with RFC 920, RFC 2401, and Reed.
- Exhibit C7, Claim Chart—USP 7,418,504 Relative to RFC 2230, Alone and in Conjunction with RFC 920, RFC 2401, Reed, and Beser.
- Exhibit C8, Claim Chart—USP 7,418,504 Relative to RFC 2538, Alone and in Conjunction with RFC 920, RFC 2401, Reed, Beser, and RFC 2065.
- Exhibit D1, Asserted Claims and Infringement Contentions by Plaintiff VlnetX Inc. In *VlnetX, Inc. v. Cisco Systems, Inc., Applce, Inc, Aastra Technologies Ltd., NEC Corporation, NEC Corporation of America and Aastra USA, Inc.*, Civ. Act. 6:2010cv00417 (E.D. Tex).
- Exhibit D2, Asserted Claims and Infringement Contentions by Plaintiff VlnetX Inc. against Apple Inc. Based on the 7,418,504.
- Exhibit X5, Eastlake, D., et al., IETF RFC 2538, "Storing Certificates in the Domain Name System (DNS)" (Mar. 1999).
- Exhibit X6, Kent, S. IETF RFC 2401, Security Architecture for the Internet Protocol, (Nov. 1998) <http://www.ietf.org/rfc/rfc2401.txt>.
- Exhibit X8 Postel, J. et al., IETF RFC 920, "Domain Requirements" (Oct. 1984) Is Accessible at <http://www.ietf.org/rfc/rfc920.txt>.
- Exhibit X10, Reed, M. et al. "Proxies for Anonymous Routing," 12th Annual Computer Security Applications Conference, San Diego, CA, Dec. 9-13, 1996.
- Request for Inter Partes Reexamination Transmittal form. Transmittal Letter.
- Request for Inter Partes Reexamination Under 35 U.S.C. § 311.
- Exhibit D-7, "Thomas": Brian Thomas, Recipe for E-Commerce, IEEE Internet Computing, (Nov.-Dec. 1997).
- Exhibit D-9, "Kent II": Stephen Kent & Randall Atkinson, "IP Encapsulating Security Payload (ESP)," Internet Engineering Task Force, Internet Draft (Feb. 1998).
- Exhibit C1, Claim Chart—USP 7,921,211 Relative to Solana, Alone and in Conjunction with RFC 920, Reed and Beser (Came from Inval. Cisco dtd Nov. 18, 2011).
- Exhibit C3, Claim Chart—USP 7,921,211 Relative to Provino, Alone and in Conjunction with RFC 920, Reed, and Beser.
- Exhibit C7, Claim Chart—USP 7,921,211 Relative to RFC 2230, Alone and in Conjunction with RFC 920, Reed, and Beser.
- 211 Request for Inter Partes Reexamination.
- Exhibit C1, Claim Chart—USP 7,418,504 Relative to Solana, Alone and in Conjunction with RFC 920, Reed and Beser.
- Exhibit C5, Claim Chart—USP 7,418,504 Relative to Provino in View of RFC 2504 and in Further Conjunction with RFC 920, Reed and Beser.
- Exhibit C6, USP 7,418,504 Relative to Beser, Alone and in Conjunction with RFC 920, RFC 2401, and Reed.
- 504 Request for Inter Partes Reexamination.
- Defendants' Supplemental Joint Invalidity Contentions.
- Exhibit 226, *Securing Web Access with DCE vs. Claims of the '135 Patent*.
- Exhibit 227, *Securing Web Access with DCE vs. Claims of the '151 Patent*.
- Exhibit 228, Understanding OSF DCE 1.1 for AIX and OS/2 vs. Claims of the '135 Patent.
- Exhibit 229, Understanding OSF DCE 1.1 for AIX and OS/2' vs. Claims of the '151 Patent.
- Exhibit 230, Understanding OSF DCE 1.1 for AIX and OS/2 vs. Claims of the '180 Patent.
- Exhibit 231, Understanding OSF DCE 1.1 for AIX and OS/2 vs. Claims of the '211 Patent.
- Exhibit 232, Understanding OSF DCE 1.1 for AIX and OS/2 vs. Claims of the '504 Patent.
- Exhibit 233, Understanding OSF DCE 1.1 for AIX and OS/2 vs. Claims of the '759 Patent.
- Exhibit 234, U.S. '648 vs. Claims of the '135 Patent.
- Exhibit 235, U.S. '648 vs. Claims of the '211 Patent.
- Exhibit 236, U.S. '648 vs. Claims of the '504 Patent.
- Exhibit 237, U.S. '648 vs. Claims of the '135 Patent.
- Exhibit 238, *Gauntlet System vs. Claims of the '211 Patent*.
- Exhibit 239, *Gauntlet System vs. Claims of the '504 Patent*.
- Exhibit 240, *Gauntlet System vs. Claims of the '135 Patent*.
- Exhibit 241, U.S. '588 vs. Claims of the '211 Patent.
- Exhibit 242, U.S. '588 vs. Claims of the '504 Patent.
- Exhibit 243, *Microsoft VPN vs. Claims of the '135 Patent*.
- Exhibit 244, *Microsoft VPN vs. Claims of the '211 Patent*.
- Exhibit 245, *Microsoft VPN vs. Claims of the '504 Patent*.
- Exhibit 246, ITU-T Standardization Activities vs. Claims of the '135 Patent.
- Exhibit 247, U.S. '393 vs. Claims of the '135 Patent.
- Exhibit 248, The Miller Application vs. Claim 13 of the '135 Patent.
- Exhibit 249, *Gauntlet System vs. Claims of the '151 Patent*.
- Exhibit 250, ITU-T Standardization Activities vs. Claims of the '151 Patent.
- Exhibit 251, U.S. Patent No. 5,940,393 vs. Claims of the '151 Patent.
- Exhibit 252, *Microsoft VPN vs. Claims of the '151 Patent*.
- Exhibit 253, U.S. Patent No. 6,324,648 vs. Claims of the '151 Patent.
- Exhibit 254, U.S. Patent No. 6,857,072 vs. Claims of the '151 Patent.
- Exhibit A, Aventail Press Release, May 2, 1997.
- Exhibit B, InfoWorld, "Aventail Delivers Highly Secure, Flexible VPN Solution," InfoWorld, p. 64D, (1997).
- Exhibit C, Aventail AutoSOCKS v2.1 Administrator's Guide.
- Exhibit D, Aventail Press Release, Oct. 12, 1998.
- Exhibit G, Aventail Press Release, May 26, 1999.
- Exhibit H, Aventail Press Release, Aug. 9, 1999.
- Exhibit J, Aventail ExtraNet Center 3.1: Security with Solid Management, Network Computing, Jun. 28, 1999.
- Petition in Opposition to Patent Owner's Petition to Vacate Inter Partes ReExamination Determination on Certain Prior Art.
- Exhibit B, Certificate of Service to Request for Inter Partes Reexamination Under U.S.C. § 311.
- Exhibit C1, Claim Chart Aventail Connect v3.1.
- Exhibit C2, Claim Chart Aventail Connect v3.01.
- Exhibit C3, Claim Chart Aventail AutoSOCKS.
- Exhibit C4, Claim Chart Wang.

(56)

## References Cited

## OTHER PUBLICATIONS

- Exhibit C5, Claim Chart Beser.  
 Exhibit C6, Claim Chart BINGO.  
 Exhibit X6, U.S. Patent 6,496,867.  
 Exhibit X10, U.S. Patent 4,885,778.  
 Exhibit X11, U.S. Patent 6,615,357.  
 Exhibit Y3 U.S. Patent 5,950,519.  
 Exhibit D, v3.1 Administrator's Guide.  
 Exhibit E-1, Claim Charts Applying Kiuchi to Various Claims of the '135 Patent.  
 Exhibit E-2, Claim Charts Applying Wesinger to Various Claims of the '135 Patent.  
 Exhibit E-3, Claim Charts Applying Solana to Various Claims of the '135 Patent.  
 Exhibit E-4, Claim Charts Applying Aziz to Various Claims of the '135 Patent.  
 Request for Inter Partes Reexamination.  
 Exhibit C1, Claim Chart Aventura Connect v3.01.  
 Exhibit C2, Claim Chart Aventura AutoSOCKS.  
 Exhibit C3, Claim Chart BINGO.  
 Exhibit C4, Claim Chart Beser.  
 Exhibit C5, Claim Chart Wang.  
 Exhibit B, Certificate of Service to Request for Inter Partes Reexamination Under 35 U.S.C. § 311.  
 Exhibit E-3, Claim Charts Applying Blum to Claims of the '151 Patent.  
 Exhibit E-5, Claim Charts Applying Kiuchi and Edwards, and Kiuchi, Edwards, and Martin to Claims of the '151 Patent.  
 Exhibit A, U.S. Patent 6,839,759.  
 Exhibit C-1, U.S. Patent 6,502,135.  
 Exhibit E-1, Claim Charts Applying Kiuchi, as Primary Reference to the '759 Patent.  
 Exhibit E-2, Claim Charts Applying Kent as a Primary Reference to the '759 Patent.  
 Exhibit E-3, Claim Charts Applying Aziz as a Primary Reference to the '759 Patent.  
 Exhibit E-4, Claim Charts Applying Kent in View of Caronni as a Primary Combination of References to the '759 Patent.  
 Certificate of Service to Request for Inter Partes Reexamination Under 35 U.S.C. § 311.  
 Exhibit B1 File History of U.S. Patent 7,418,504.  
 Exhibit B2, File History of U.S. Appl. No. 09/558,210.  
 Exhibit D-10, Gaspoz et al., "VPN on DCE: From Reference Configuration to Implementation," Bringing Telecommunication Services to the People—IS&N '95, Third International Conference on Intelligence in Broadband Services and Networks, Oct. 1995 Proceedings, Lecture Notes in Computer Science, vol. 998 (Springer, 1995).  
 Exhibit D-11, U.S. Patent No. 6,269,099.  
 Exhibit D-11, U.S. Patent No. 6,560,634.  
 Exhibit D-13, Pallen, "The World Wide Web," British Medical Journal, vol. 311 at 1554 (Dec. 1995).  
 Exhibit D-14, Rivest et al., "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," Communications of the ACM, 21:120-126 (Feb. 1978).  
 Exhibit D-15, U.S. Patent No. 4,952,930.  
 Exhibit D-17, Pfaffenberger, Netscape Navigator 3.0: Surfing the Web and Exploring the Internet, Academic Press (1996).  
 Exhibit D-18, Gittler et al., "The DCE Security Service," Hewlett-Packard Journal, pp. 41-48 (Dec. 1995).  
 Exhibit D-6, U.S. Patent No. 5,689,641.  
 Exhibit D-9, Lawton, "New Top-Level Domains Promise Descriptive Names," Sunworld Online, 1996.  
 Exhibit E-1, Catalog Listing by IBM for RS/6000 Redbooks Collection which includes a Link to the *Lendenmann* reference. The link to the *Lendenmann* reference was archived at archive.org on Dec. 7, 1998 and retrieved by the Wayback Machine.  
 Exhibit E-10, An Archived Version of the Lawton reference archived at archive.org on Feb. 19, 1999 and retrieved by the Wayback Machine.  
 Exhibit E-11, Abstracts of the Proceedings of the Symposium on Network and Distributed System Security, 1996, Archived at archive.org on Apr. 10, 1997, and retrieved by the Wayback Machine.  
 Exhibit E-12, 1996 Symposium on Network and Distributed System Security, Website Archived by archive.org (Apr. 10, 1997), Retrieved by the Wayback Machine at <http://web.archive.org/web/19970410114853/http://computer.org/cspress/catalog/proc9.htm>.  
 Exhibit E-13, Search Results for ISBN 0-12-553153-2 (Pfaffenberger) from [www.isbnsearch.org](http://www.isbnsearch.org).  
 Exhibit F-1, Claim Charts applying Lendenmann as a Primary Reference to the '504 Patent.  
 Exhibit F-2, Claim Charts applying Aziz as a Primary Reference to the '504 Patent.  
 Exhibit F-3, Claim Charts applying Kiuchi and Pfaffenberger as Primary References to the '504 Patent.  
 Exhibit E-2, First Page of U.S. Patent No. 5,913,217 published Jun. 15, 1999 and citing a portion of the Lendenmann reference as a prior art reference.  
 Exhibit E-3, Request for Comments 2026, "The Internet Standards Process—Revision 3," Oct. 1996.  
 Exhibit E-4, First Page of U.S. 5,463,735, published Oct. 31, 1995 and citing RFC 793 as a prior art Reference.  
 Exhibit E-5, Catalog listing from Boston University Digital Common Website, listing the Martin reference with an issue date of Feb. 21, 1998.  
 Exhibit E-6, Technical Reports Archive Listing from Boston University Computer Science Department which includes a link to the Martin paper. The link to the Martin paper was archived at archive.org on Jan. 22, 1998 and Retrieved by the Wayback Machine.  
 Exhibit E-7, Boston University Computer Science Department Technical Reports Instructions, available at: <http://www.cs.bu.edu/techreports/INSTRUCTIONS>.  
 Exhibit E-8, U. Möller, "Implementation eines Anonymisierungsverfahrens für WWW-Zugriffe," Diplomarbeit, Universität Hamburg (Jul. 16, 1999), citing to Martin at p. 77.  
 Exhibit E-9, First page of U.S. 5,737,423, published Apr. 7, 1998 and citing Schneider as Prior Art Reference.  
 Request for Inter Partes ReExamination; U.S. Patent 7,418,504.  
 Request for Inter Partes ReExamination Transmittal Form; U.S. Patent 7,418,504.  
 Request for Inter Partes Reexamination Transmittal 7,418,504.  
 Exhibit C7, Claim Chart—USP 7,921,211 relative to RFC 2230, alone and in conjunction with RFC 2401, Reed, and Beser.  
 Exhibit C6 USP 7,418 504 relative to Beser, alone and in conjunction with RFC 920, RFC 2401, and Reed.  
 Exhibit 229, Understanding OSF DCE 1.1 for AIX and OS/2 vs. Claims of the '151 Patent.  
 Exhibit 237, U.S. '072 vs. Claims of the '135 Patent.  
 Exhibit 248, *The Miller Application vs. Claim of the '135 Patent*.  
 Exhibit 253, U.S. Patent No. 6,324,648 1 vs. Claims of the '151 Patent.  
 Petition in Opposition to Patent Owner's Petition to Vacate *Inter Partes* Reexamination.  
 Exhibit B1, File History of U.S. Patent 7,921,211.  
 Exhibit B2, File History of U.S. Appl. No. 10/714,849.  
 Exhibit B4, *VirnetX, Inc. v. Microsoft Corp.*, Case No. 6:07-cv-80, Memorandum Opinion on Claim Construction (E.D. Tex. Jul. 30, 2009).  
 Exhibit D15, U.S. Patent 4,952,930.  
 Exhibit F1, Claim Charts Applying Lendenmann as a Primary Reference to the '211 Patent.  
 Exhibit F2, Claim Charts Applying Aziz as a Primary Reference to the '211 Patent.  
 Exhibit F3, Claim Charts Applying Kiuchi and Pfaffenberger as Primary References to the '211 Patent.  
 Exhibit 2, Letter and attachment from Ramzi Khazen, Counsel for VirnetX, to Dmitriy Kheyfits, Counsel for Cisco Systems (Jun. 23, 2011).  
 Exhibit P, Malkin, "Dial-In Virtual Private Networks Using Layer 3 Tunneling".  
 Exhibit Q, Ortiz, "Virtual Private Networks: Leveraging the Internet".



(56)

**References Cited**

## OTHER PUBLICATIONS

- Exhibit R, Keromytis, "Creating Efficient Fail-Stop Cryptographic Protocols".  
 Transcript of Markman Hearing Dated Jan. 5, 2012.  
 Declaration of John P. J. Kelly, Ph.D.  
 Defendants' Responsive Claim Construction Brief; Exhibits A-P and 1-7.  
 Joint Claim Construction and Prehearing Statement Dated Nov. 8, 2011.  
 Exhibit A: Agreed Upon Terms Dated Nov. 8, 2011.  
 Exhibit B: Disputed Claim Terms Dated Nov. 8, 2011.  
 Exhibit C: VimetX's Proposed Construction of Claim Terms and Supporting Evidence Dated Nov. 8, 2011.  
 Exhibit D: Defendant's Intrinsic and Extrinsic Support Dated Nov. 8, 2011.  
 Declaration of Austin Curry in Support of VimetX Inc.'s Opening Claim Construction Brief.  
 Declaration of Mark T. Jones Opening Claims Construction Brief.  
 VimetX Opening Claim Construction Brief.  
 VimetX Reply Claim Construction Brief.  
 European Search Report from corresponding EP Application No. 11005789.  
 European Search Report from corresponding EP Application No. 11005792.  
 Eastlake, "Domain Name System Security Extensions," Network Working Group, RFC: 2535 pp. 2-11 (Mar. 1999).  
 Press Release; VimetX and Aastra Sign a Patent License Agreement, 4 pages, May 2012, Printed from Website: <http://virmetx.com/virmetx-and-aastra-sign-a-patent-license-agreement/>.  
 Press Release; VimetX and Mitel Networks Corporation Sign a Patent License Agreement, 5 pages, Jul. 2012, Printed from Website: <http://virmetx.com.virmetx-and-mitel-networks-corporation-sign-a-patent-license-agreement/>.  
 Press Release; VimetX and NEC Corporation and NEC Corporation of America Sign a Patent License Agreement, 5 pages, Aug. 2012, Printed from Website: <http://virmetx.com/virmetx-and-nec-corporation-and-nec-corporation-of-america-sign-a-patent-license-agreement/>.  
 Supplemental Declaration of Angeles D. Keromytis, Ph.D from Control No. 95001789 pp. 1-18, dated Dec. 20, 2012.  
 Supplemental Declaration of Angelos D. Keromytis, Ph.D from Control No. 95001851 pp. 1-13, dated Dec. 30, 2012.  
 Supplemental Declaration of Angelos D. Keromytis, Ph.D from Control No. 95001788 pp. 1-18, dated Dec. 18, 2012.  
 Supplemental Declaration of Angelos D. Keromytis, Ph.D from Control No. 95001856 pp. 1-13, dated Dec. 30, 2012.  
*VirmetX vs Apple* Transcript of Trial. Afternoon Session, 12:05 p.m., dated Nov. 5, 2012.  
 Certified Copy dated Sep. 18, 2012 of U.S. Patent No. 6,502,135, 73 pages.  
 Certified Copy dated Dec. 30, 2009 of Assignment for U.S. Appl. No. 95/047,83 12 pages.  
 Certified Copy dated Mar. 11, 2008 of U.S. Appl. No. 09/504,783, 1500 pages.  
 Certified Copy dated Mar. 30, 2011 of U.S. Patent No. 7,418,504, 74 pages.  
 Certified Copy dated Oct. 17, 2012 of Assignment for U.S. Appl. No. 10/714,849, 10 pages.  
 Certified Copy dated Apr. 4, 2011 of U.S. Appl. No. 10/714,849, 1170 pages.  
 Certified Copy dated Mar. 30, 2011 of U.S. Patent No. 7,490,151, 63 pages.  
 Certified Copy dated Oct. 17, 2012 of Assignment for U.S. Appl. No. 10/259,494, 19 pages.  
 Certified Copy dated Apr. 4, 2011 of U.S. Appl. No. 10/259,454, 1359 pages.  
 Certified Copy dated Apr. 12, 2011 of U.S. Patent No. 7,921,211, 78 pages.  
 Certified Copy dated Oct. 17, 2012 of Assignment for U.S. Appl. No. 11/840,560, 12 pages.  
 Certified Copy dated Apr. 20, 2011 of U.S. Appl. No. 11/840,560, 3 pages.  
 iPhone User Guide for iPhone OS 3.1 Software, 217 pages, 2009.  
 iPhone User Guide for iOS 4.2 and 4.3 Software, 274 pages, 2011.  
 iPhone User Guide for iPhone and iPhone 3G, 154 pages, 2008.  
 iPhone User Guide for iOS 5.0 Software, 163 pages, 2011.  
 iPad User Guide for iOS 5.0 Software, 141 pages, 2011.  
 iPad User Guide for iOS 4.2 Software, 181 pages, 2010.  
 iPad User Guide for iOS 4.3 Software, 198 pages, 2011.  
 iPad User Guide, 154 pages, 2010.  
 iPod Touch User Guide for iOS 5.0 Software, 143 pages, 2011.  
 iPod Touch User Guide, 122 pages, 2008.  
 iPod Touch User Guide for iPhone OS 3.0 Software, 153 pages, 2009.  
 iPod Touch User Guide for iPhone OS 3.1 Software, 169 pages, 2009.  
 iPod Touch User Guide for iOS 4.3 Software, 230 pages, 2011.  
 iPod Touch Features Guide, 98 pages, 2008.  
 VPN Server Configuration for iOS; Networking & Internet Enterprise Deployment, 12 pages, 2011.  
 iPhone Configuration Utility User Guide, 26 pages, 2010.  
 iPhone Configuration Utility; Networking & Internet: Enterprise Deployment, 26 pages, 2011.  
 iPhone Configuration Utility; Networking>Internet & Web, 24 pages, 2010.  
 iOS Configuration Profile Reference; Networking & Internet; Enterprise Deployment, 24 pages, 2011.  
 iPhone OS Enterprise Deployment Guide; Second Edition, for Version 3.1 or Later, 91 pages, 2009.  
 iPhone OS; Enterprise Deployment Guide; Second Edition, for Version 3.2 or Later, 90 pages, 2010.  
 CFHost Reference; Developer, 20 pages, 2008.  
 CFNetwork Programming Guide; Developer, 60 pages, 2011.  
 CFStream Socket Additions; Developer, 22 pages, 2010.  
 Mac OS X Developer Library; CFHostSample.c, 1 page.  
 Mac OS X Developer Library; CFHostSample, 1 page, 2004.  
 Mac OS X Developer Library; Document Revision History, 1 page, 2004.  
 Apple Push Notification Service, Distribution Service, Version 1.0, 6 pages, 2009.  
 iOS Human Interface Guidelines; Developer, 184 pages, 2012.  
 Networking & Internet Starting Point, 3 pages, 2011.  
 Server Admin. 10.5 Help; Viewing a VPN Overview, 1 page.  
 iOS: Supported Protocols for VPN, 2 pages, 2010.  
 iPhone in Business Virtual Private Networks (VPN), 3 pages, 2010.  
 iPhone and iPad in Business Deployment Scenarios, 26 pages, 2011.  
 Deploying iPhone and iPad Virtual Private Networks, 3 pages, 2011.  
 Deploying iPhone and iPad; Security Overview, 6 pages, 2011.  
 iPad in Business; "Ready for Work," 2012, 5 pages.  
 iOS: Using FaceTime, 2 pages, 2011, Printed from website <http://support.apple.com/kb/HT4317>.  
 MobileMe: "Secure Chat" is Unavailable in OS X Lion, 2 pages, 2012, Printed from Website: <http://support.apple.com/kb/T33902>.  
 iPhone 4 and Pod Touch (4th Generation): Using FaceTime, 2 pages, 2010, Printed from Website: <http://support.apple.com/kb/HT4319>.  
 iPhone; "Picking Up Where Amazing Left Off," 11 pages, 2012, Printed from Website: <http://www.apple.com/iphone/features/facetime>.  
 FaceTime for Mac; "Say Hello to FaceTime for Mac," 4 pages, 2012, Printed from Website: <http://www.apple.com/mac/facetime>.  
 iPad; "Your New Favorite Way to do Just About Everything," 8 pages, 2012, Printed from Website" <http://www.apple.com/ipad/built-in-apps/>.  
 iPod Touch; FaceTime, "Oh I see what you're saying," 2 pages.  
 Apple Press Info; Apple Presents iPhone 4, Printed from Website: <http://www.apple.com/pr/library/apple-presents-iphone>.  
 iPod Touch; FaceTime, "Oh I See What You're Saying," 3 pages, 2012, Printed from Website: <http://www.apple.com/iPodtouch/built-in-apps/facetime.htm>.  
 IOS 4, The World's Most Advanced Mobile Operating System, 5 pages, Printed from Website: <http://www.apple.com/iphone/ios4>.  
 Apple Press Info; Apple Reinvents the Phone with iPhone, 3 pages, 2007, Printed from Website: <http://www.apple.com/pr/library/2007/01/09Apple-reinvents-the-phone>.

(56)

## References Cited

## OTHER PUBLICATIONS

- Apple Press Info; Apple Announces the New iPhone 3Gs—The Fastest, Most Powerful iPhone Yet, 3 pages, 2009, Printed from the Website: <http://www.apple.com/pr/library/2009/06/08Apple-Announces-the-new-iphone3GS>.
- Apple Press Info; Apple Launches iPhone 4S, ios 5 & iCloud, iPhone 4S Features Dual-Core A5 Chip, All New Camera, full 1080p HD Video Recording & Introduces Siri, 2011, 2 pages, Printed from website: <http://www.apple.com/pr/library/2011/10/04Apple-Launches-iPhone-4S-ios-5-iCloud.html>.
- Apple Press Info; Apple Introduces New iPod Touch, Features Retina Display, A4 Chip, FaceTime Video Calling, HD Video Recording & Game Center, 2 pages, 2010, Printed from Website <http://www.apple.com/pr/library/2010/09/01Apple-Introduces-New-iPod-touch.html>.
- Apple Press Info; Apple Launches iPad, Magical & Revolutionary Device at an Unbelievable Price, 2 pages, 2010, Printed from Website: <http://www.apple.com/pr/library/2010/01/27Apple-Launches-iPad.html>.
- Apple Press Info; Apple Launces New iPad, New iPad Features Retina Display, A5X Chip, 5 Megapixel iSight Camera & Ultrafast 4G LTE, 2012, 3 pages, Printed from the Website: <http://www.apple.com/pr/library/2012/03/07Apple-Launches-New-iPad.html>.
- FaceTime; “Phone Calls Like You’ve Never Seen Before,” 3 pages.
- Apple Press Info; Apple Brings FaceTime to the Mac, 1 pages, Printed from Website <https://www.apple.com/pr/library/2010/10/20Apple-Brings-FaceTime-to-the-Mac.html>.
- iPad at Work; “Mobile Meetings Made Easy,” 4 pages, 2011.
- iPad—Technical Specifications, 49 pages, Printed from Website: <http://support.apple.com/kb/sp58C>.
- Stirling Design, 8 pages, 2008.
- Quick Guide: SSL VPN Technical Primer, 11 pages, 2010.
- Silva, “Secure iPhone Access to Corporate Web Applications,” Technical Brief, 10 pages.
- Defendant Apple Inc.’s Third Supplemental Responses to VirnetX Inc.’s First Request for Admission to Apple Inc. dated, Apr. 13, 2012, 207 pages.
- Apple Support Communities, 4 pages, Printed from Website <http://discussions.apple.com/thread/486096?start=0&tstart=0>.
- VirnetX—Products; License and Service Offerings, 1 page.
- VirnetX Contact Information, 4 pages, 2011.
- VirnetX Launches Secure Domain Name Initiative; 4G/LTE Security, 1 page, 2010.
- VirnetX Gabriel Connection; Enabling Safe Network Neighborhoods, 2 pages, 2012.
- Baughner et al., “The Secure Real-Time Transport Protocol (SRTP),” Network Working Group, RFC:3711, 39 pages, 2004.
- Jennings et al., “Resource Location and Discovery (Reload) Draft-Bryan-P2PSIP-Reload-04,” Internet-Draft, Dec. 12, 2008, pp. 1-127.
- Barnes et al., “Verification Involving PSTN Reachability: Requirements and Architecture Overview,” Internet Draft, 27 pages, 2012.
- April Inc. Form 10-K (Annual Report) filed Dec. 1, 2005 for the Period Ending Sep. 24, 2005, Edgar Online, 1400 pages, 2011.
- Phone, Facetime; “Be in Two Places at Once,” 3 pages, Printed from the Website <http://www.apple.com/ios/facetime>.
- Apple Press Info; Apple Presents iPhone 4, All-New Design with FaceTime Video Calling, Retina. Display, 5 Megapixel Camera & HD Video Recording, 3 pages, 2010.
- NYSE AMEX:VHC; Cowen and Co. 39th Annual Technology Media & Telecom Conference, 36 pages, Jun. 2, 2011.
- Pindyck et al., “Market Power: Monopoly and Monopsony,” Microeconomics, Sixth Edition, pp. 370-371.
- Press Release; IpCapital Group Completes VirnetX IP Licensing Evaluation, 3 pages.
- Microsoft Real-Time Communications: Protocols and Technologies, Microsoft TechNet, 22 pages, 2010.
- Filing Receipt dated Sep. 23, 2011 for U.S. Appl. No. 13/223,259.
- Email Communications Regarding Apple Product Innovations, 6 pages, 2010.
- Mathy et al., “Traversal Using Relays Around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN),” Internet Engineering Task Force (IETF), RFC; 5766, 67 pages; 2010.
- Egevang et al., “The IP Network Address Translator (NAT),” Network Working Group, RFC; 1631, 10 pages, 1994.
- Srisuresh et al., “IP Network Address Translator (NAT) Terminology and Considerations,” Network Working Group, RFC:2663, 30 pages, 1999.
- Sisalem, et al., “Introduction to Cryptographic Mechanisms,” SIP Security, 356 pages, 2009.
- Curriculum Vitae, Mark T Jones, 9 pages.
- Curriculum Vitae, Roy Weinstein; 5 pages.
- How to Configure IPSec Tunneling in Windows 2000; 8 pages.
- iPhone, FaceTime; “Be in Two Places at Once,” 3 pages, Printed from Website: <http://www.apple.com/ios/facetime/>.
- iPhone, “It Does Everything Better,” 6 pages, Printed from Website: <http://www.apple.com/iphone/built-in-apps>.
- My Apple ID, “What’s an Apple ID,” 1 pages, Printed from Website: <https://appleid.apple.com/cgi-bin/webobjects/myappleid.woa>.
- Rosenberg et al., “Session Initiation Protocol (SIP): Locating SIP Servers,” Network Working Group, RFC: 3263, 17 pages, 2002.
- Certified Copy dated Sep. 21, 2012 of Reexamination Certificate No. 6,502,135 issued Jun. 6, 2011, 11 pages.
- Certified Copy dated Sep. 20, 2012 of U.S. Appl. No. 95/001,269.
- Chatterjee et al., “Bargaining Under Incomplete Information,” Operations Research, 31:835-851, 1983.
- Nash, “The Bargaining Problem,” *Econometrica*, 18:155-162, 1950.
- Nash, “Two-Person Cooperative Games,” *Econometrica*, 21:128-140, 1953.
- Choi et al., “An Analytical Solution to Reasonable Royalty Rate Calculations,” *IDEA: The Journal of Law and Technology*, 13 pages, 2001.
- The Prize in Economics 1994—Press Release dated Oct. 11, 1994, 4 pages, Printed from Website: [http://www.nobelprize.org/nobel\\_prizes/economics/laureates/1994/press.html](http://www.nobelprize.org/nobel_prizes/economics/laureates/1994/press.html).
- Putnam et al., “Bargaining and the Construction of Economically Consistent Hypothetical License Negotiations,” *The Licensing Journal*, pp. 8-15, 2004.
- Scherling et al., “Rational Reasonable Royalty Damages: A Return to the Roots,” *Landslide*, vol. 4, 4 pages, 2011.
- Jaroszc et al., “Application of Game Theory to Intellectual Property Royalty Negotiations,” Chapter 17, pp. 241-265.
- Goldscheider, Licensing Best Practices; Strategic, Territorial, and Technology Issues, 2 pages, 2006.
- iPhone Configuration Utility, 19 pages, 2012.
- VPN Server Configuration for iOS Devices, 6 pages, 2012.
- Samuelson et al., *Economics*, Fourteenth Edition, pp. 258-259, 1992.
- Stigler et al., *The Theory of Price*, Forth Edition, pp. 215-216, 1987.
- Truett et al., “Joint Profit Maximization, Negotiation, and the Determinacy of Price in Bilateral Monopoly,” *Journal of Economic Education*, pp. 260-270.
- Binmore et al., “Noncooperative Models of Bargaining,” *The Handbook of Game Theory*, 1:(7)181-225, 1992.
- Spindler et al., “Endogenous Bargaining Power in Bilateral Monopoly and Bilateral Exchange,” *Canadian Journal of Economics- Revue Canadienne D Economie*, pp. 464-474, 1974.
- Myerson, “Game Theory; Analysis or Conflict,” *Harvard University Press*, pp. 375-392.
- Binmore, “The Nash Bargaining Solution in Economic Modelling,” *The Rand Journal of Economics*, 17:176-188, 1996.
- Rubinstein et al., “On the Interpretation of the Nash Bargaining Solution and its Extension to Non-Expected Utility Preference,” *Econometrica*, 60:1171-1186, 1992.
- Greenleaf et al., “Guarantees in Auctions: The Auction House as Negotiator and Managerial Decision Maker,” *Management Science*, 39:1130-1145, 1993.
- Chan, “Trade Negotiations in a Nash Bargaining Model,” *Journal of International Economics*, 25:253-363, 1987.
- Lemley et al., “Patent Holdup and Royalty Stacking,” *Texas Law Review*, 85:1991-2049.
- Cauley, “Winning the Patent Damages Case; A Litigator’s Guide to Economic Models and Other Damage Strategies,” *Oxford Press*, pp. 29-30, 2044.

(56)

**References Cited**

## OTHER PUBLICATIONS

Degnan et al., "A Survey of Licensed Royalties," Les Nouvelles, pp. 91, 93, 94, 1997.

Kahn, "The Review of Economics and Statistics," pp. 157-164, 1993.

Microsoft Company Information; Including Stocks and Financial Information, 83 pages.

Apple Press Info: Apple Updates MacBook Pro with Next Generation Processors, Graphics & Thunderbolt I/O Technology, 3 pages, Printed from Website: <http://www.apple.com/pr/library/2011/02/24Apple-Updates-MacBook-Pro-with-Next-Generation-Processors-Graphics-Thunderbolt-I-O-Technology.html>.

Apple Press Info: Apple to Ship Mac OS X Snow Leopard on Aug. 28, 2 pages, Printed from the Website: <http://www.apple.com/pr/library/2009/08/24/apple-to-ship-mac-os-x>.

iPad, Facetime; "Once Again, iPad gets the World Talking," 3 pages, Printed from the Website: <http://www.apple.com/ipad/built-in-apps/facetime/html>.

Apple iOS: Setting up VPN, 2 pages, Printed from Website: <http://support.apple.com/kb/HT1424>.

Apple iPhone User Guide for iOS 5.1 Software, 179 pages, 2012.

Apple, Communicating with HTTP Servers, CFNetworking Programming Guide, 6 pages, 2011, Printed from the Website: <https://developer.apple.com/library/ios/documentation/networking/conceptual/CFNetwork/CFHT>.

VirnetX, Gabriel Connection Technology™ White Paper, 7 pages, 2012.

VirnetX, Technology, 4 pages, 2012.

Certified Copy dated Jan. 15, 2008 of U.S. Patent No. 6,502,135, 64 pages.

Inter Partes Reexamination Certificate date Jun. 7, 2011 for Patent No. 6,502,135.

Proceedings of The Symposium on Network and Distributed System Security, 7 pages, Feb. 22-23, 1996.

In-Q-Tel; Corporate Overview, 2 pages, 2004.

Davies, Supervisor of Translation: Tadahiro Uezono, Security for Computer Networks, Japan, Nikkei-McGraw-Hill Inc., First Edition, First Copy, p. 126-129 (Dec. 5, 1985)—(English Version and Japanese Version Submitted).

Comer, "Translated by Jun Murai and Hiroyuki Kusumoto, "Internetworking with TCP/IP vol. 1: Principles, Protocols, and Architecture, Third Edition, Japan Kyoritsu Shuppan Co., Ltd., First Edition, First Copy, p. 161-193 (Aug. 10, 1997) (English Version and Japanese Version Submitted).

Lynch et al., Supervisor of Translation: Jun Murai, "Internet System Handbook," Japan Impress Co. Ltd. First Edition p. 152-157 and p. 345-351 (Aug. 11, 1996) (English Version and Japanese Version Submitted).

Office Action dated Dec. 27, 2012 from Corresponding Canadian Patent Application No. 2723504.

Office Action dated Dec. 5, 2012 from Corresponding Japanese Patent Application No. 2011-081417.

Office Action dated Dec. 13, 2012 from Corresponding Japanese Patent Application No. 2011-085052.

Office Action dated Dec. 13, 2012 from Corresponding Japanese Patent Application No. 2011-083415.

\* cited by examiner




FIG. 1

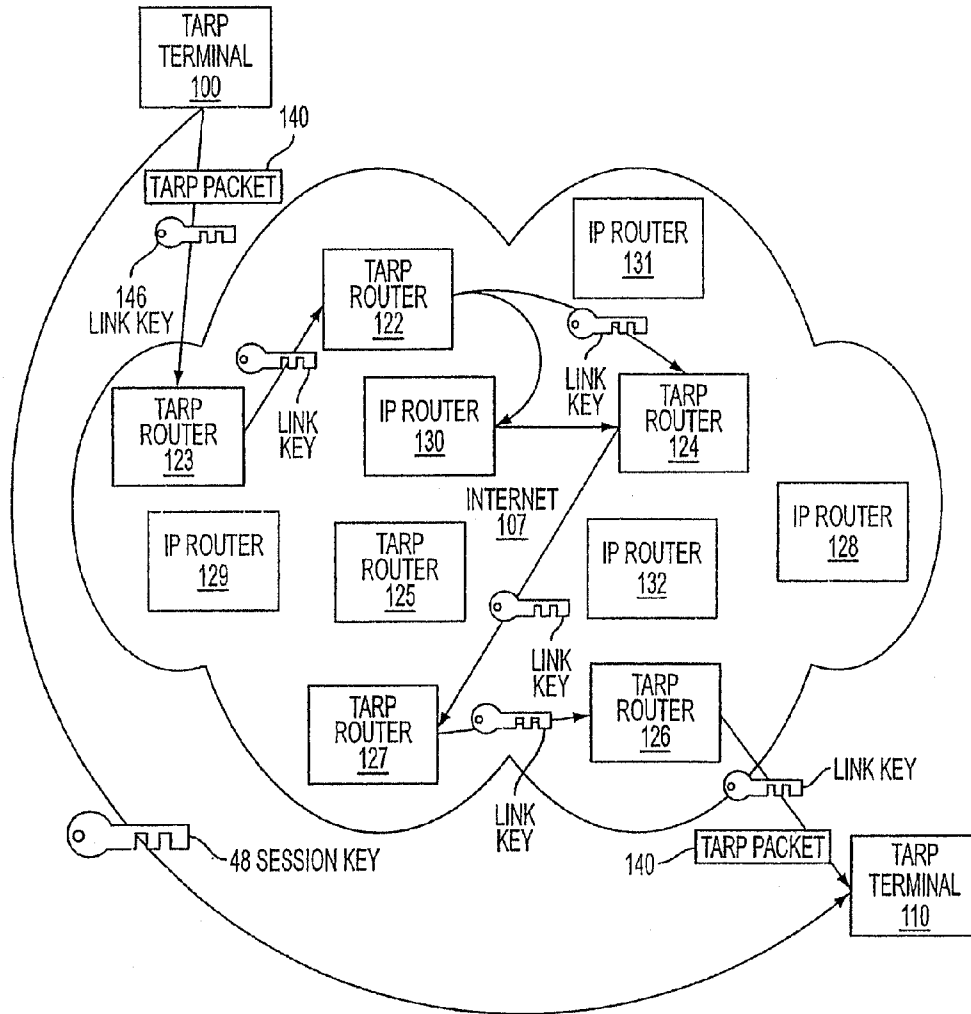


FIG. 2

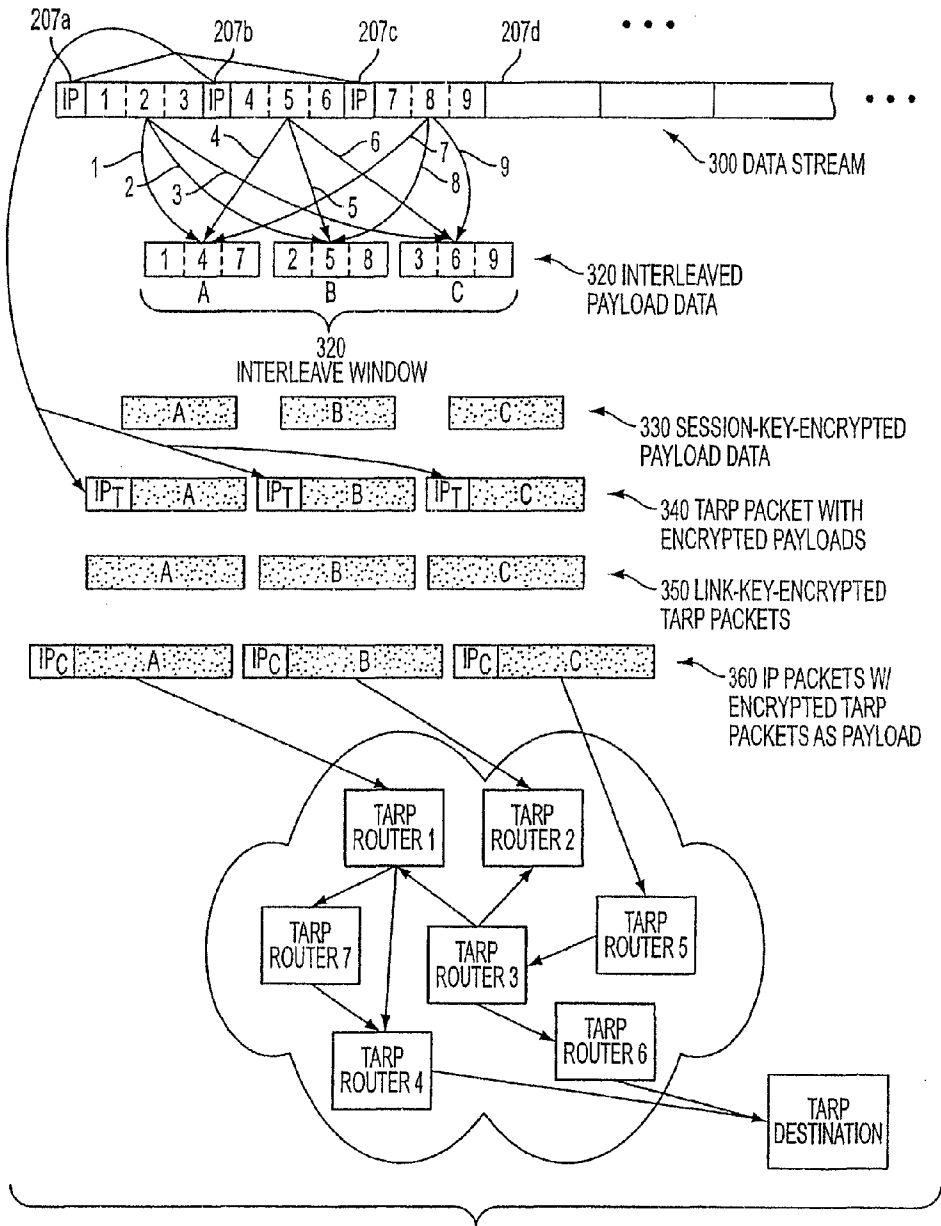


FIG. 3A




FIG. 3B




FIG. 4



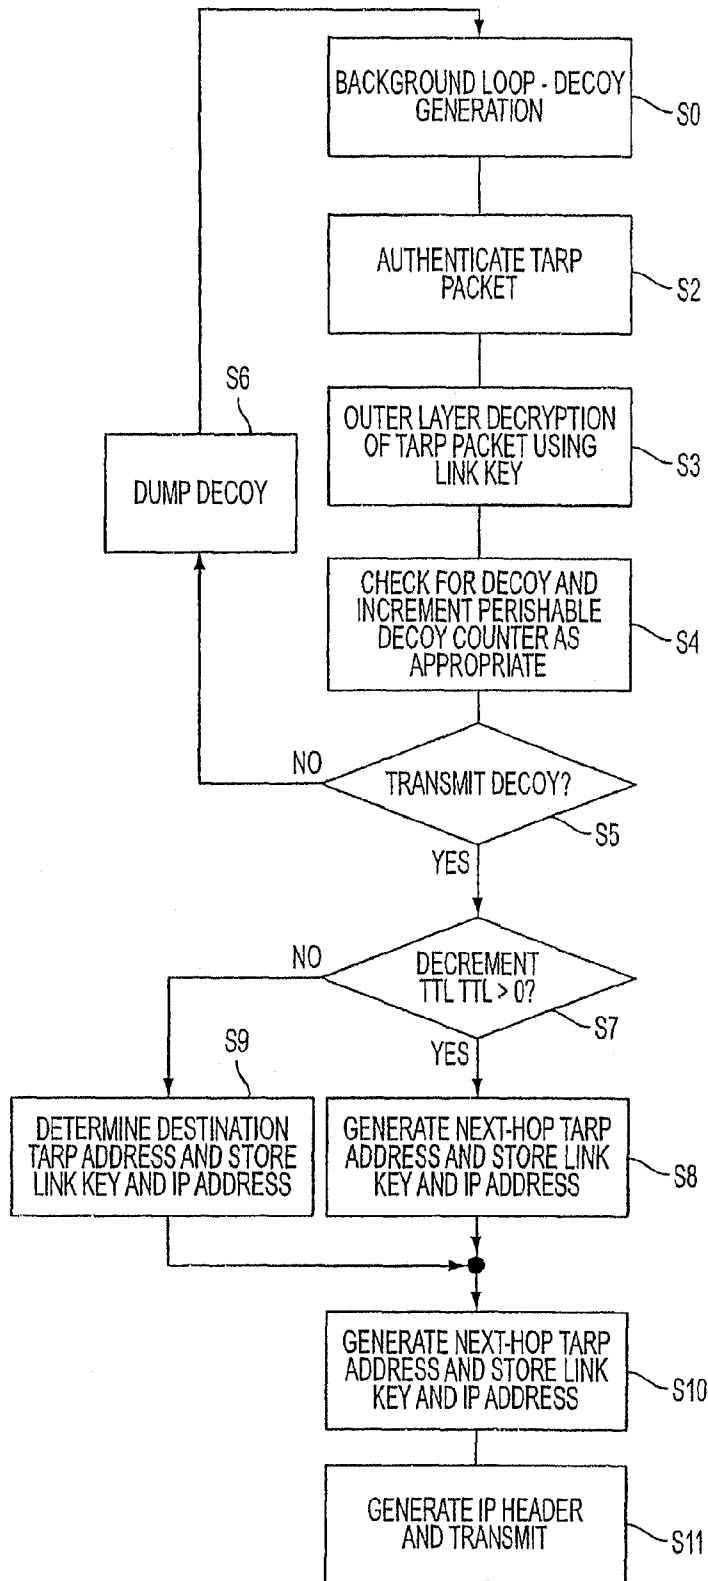


FIG. 5




FIG. 6

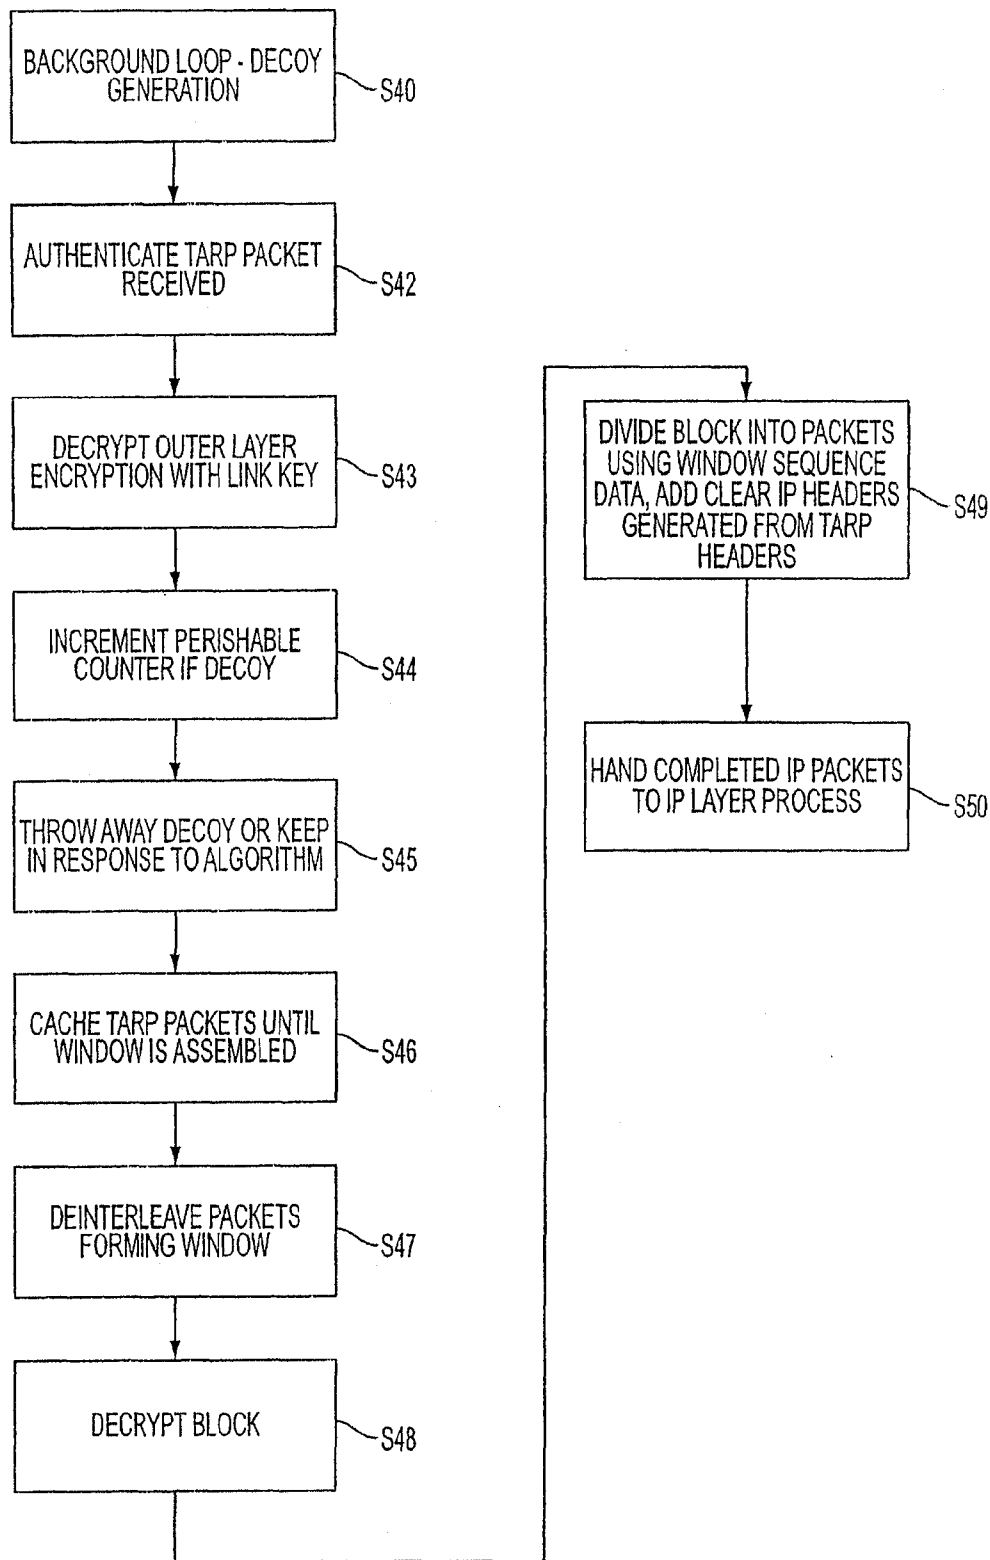


FIG. 7




FIG. 8

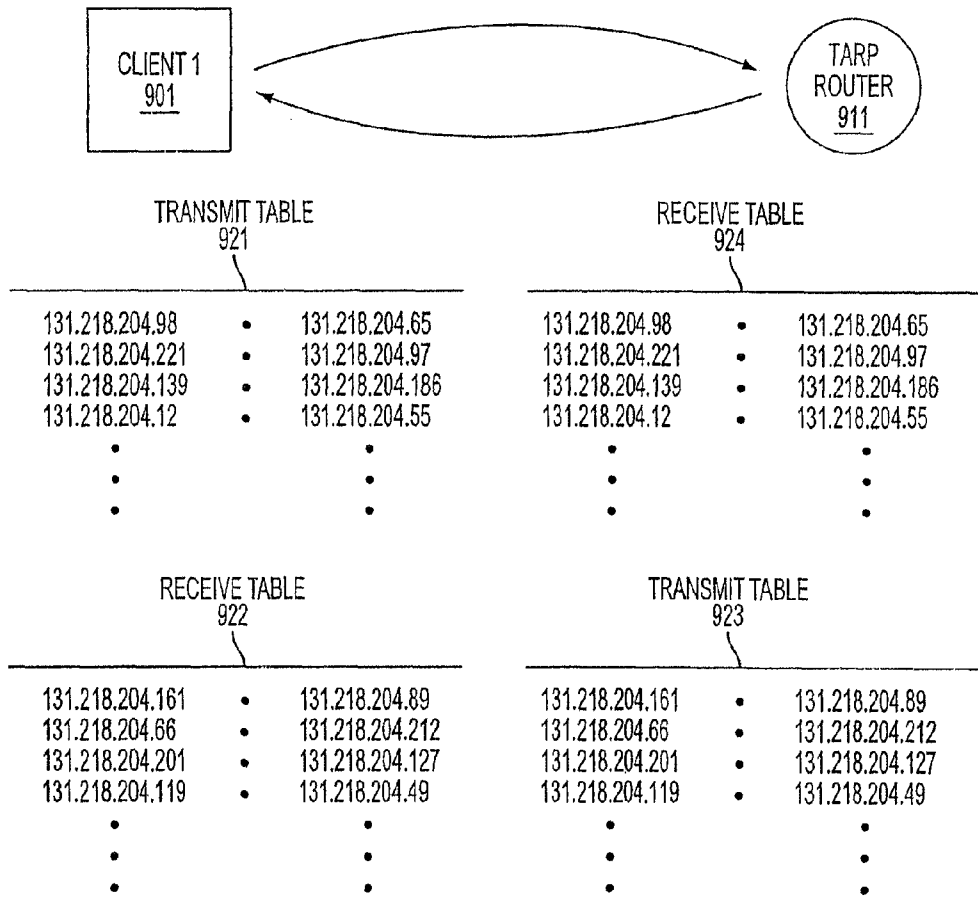


FIG. 9




FIG. 10




FIG. 11




FIG. 12A



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

FIG. 12B




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17




FIG. 18


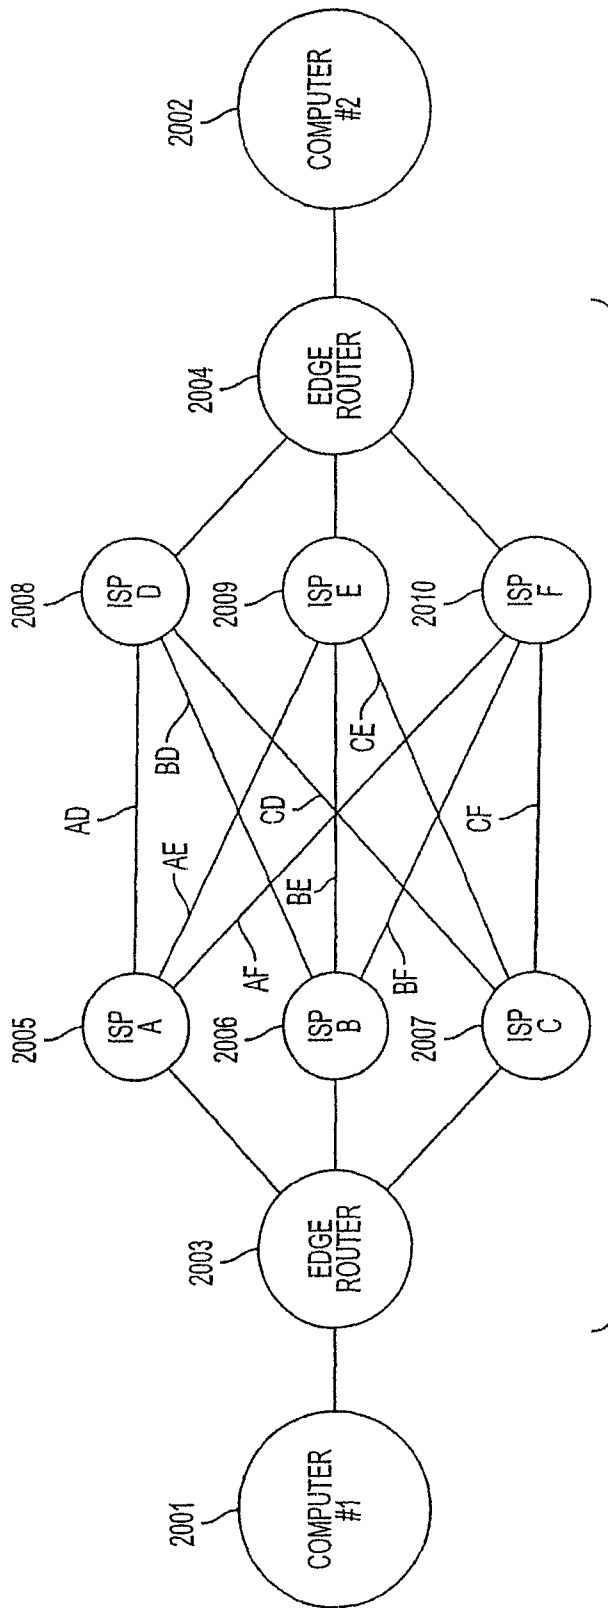


FIG. 19





2011  
FIG. 20




FIG. 21

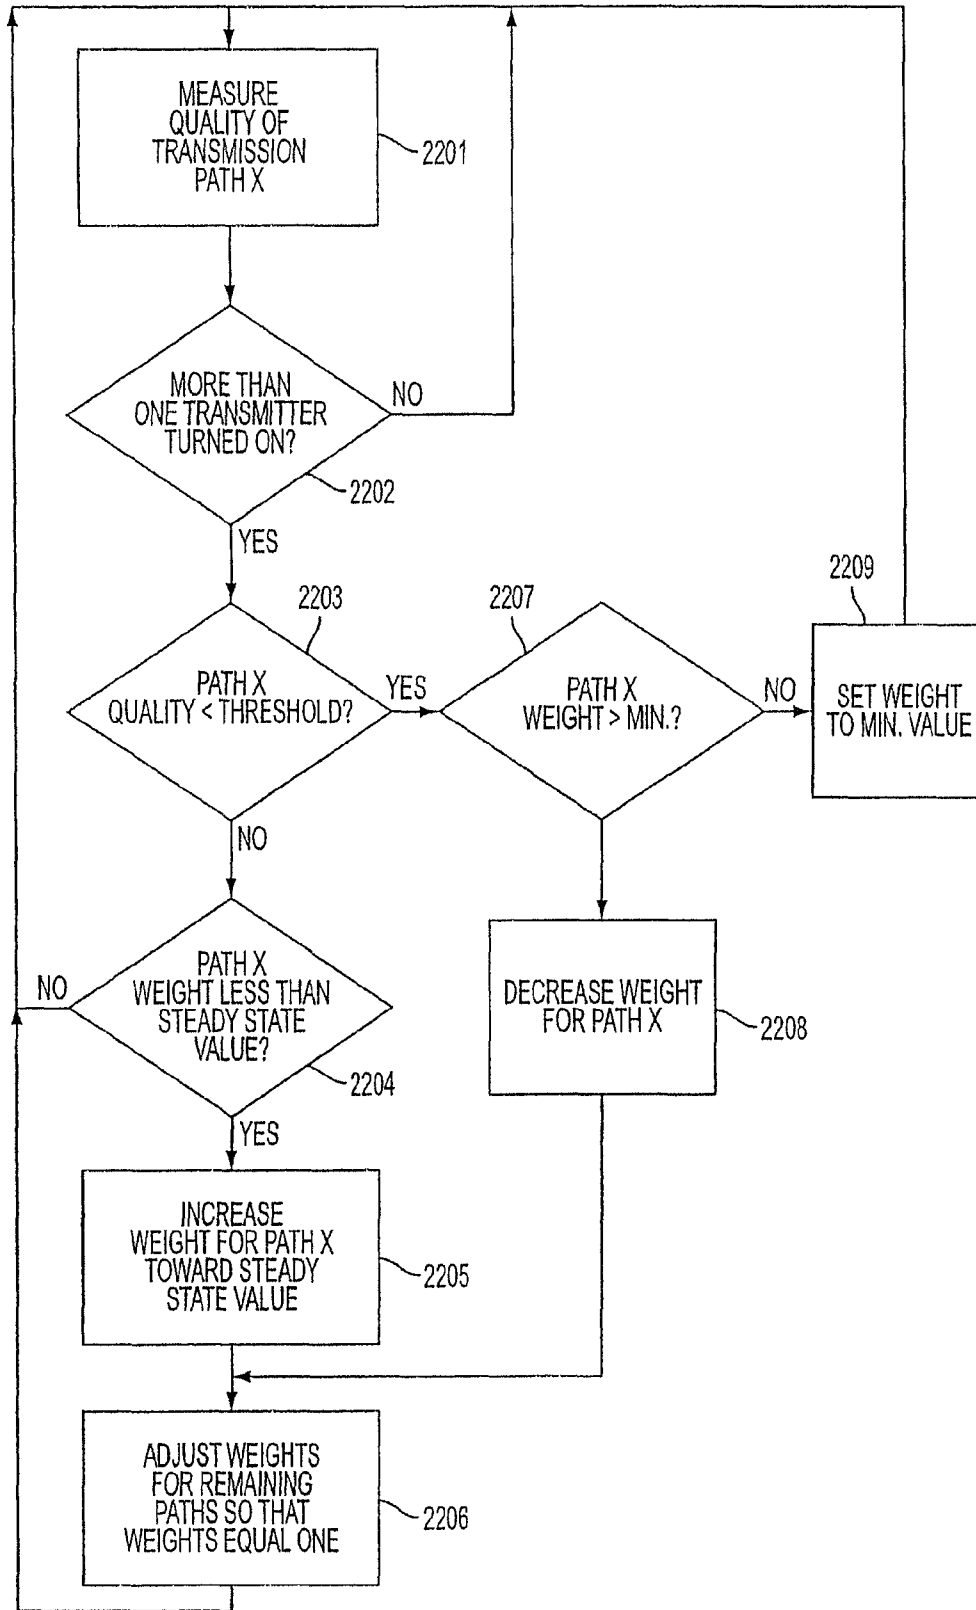


FIG. 22A

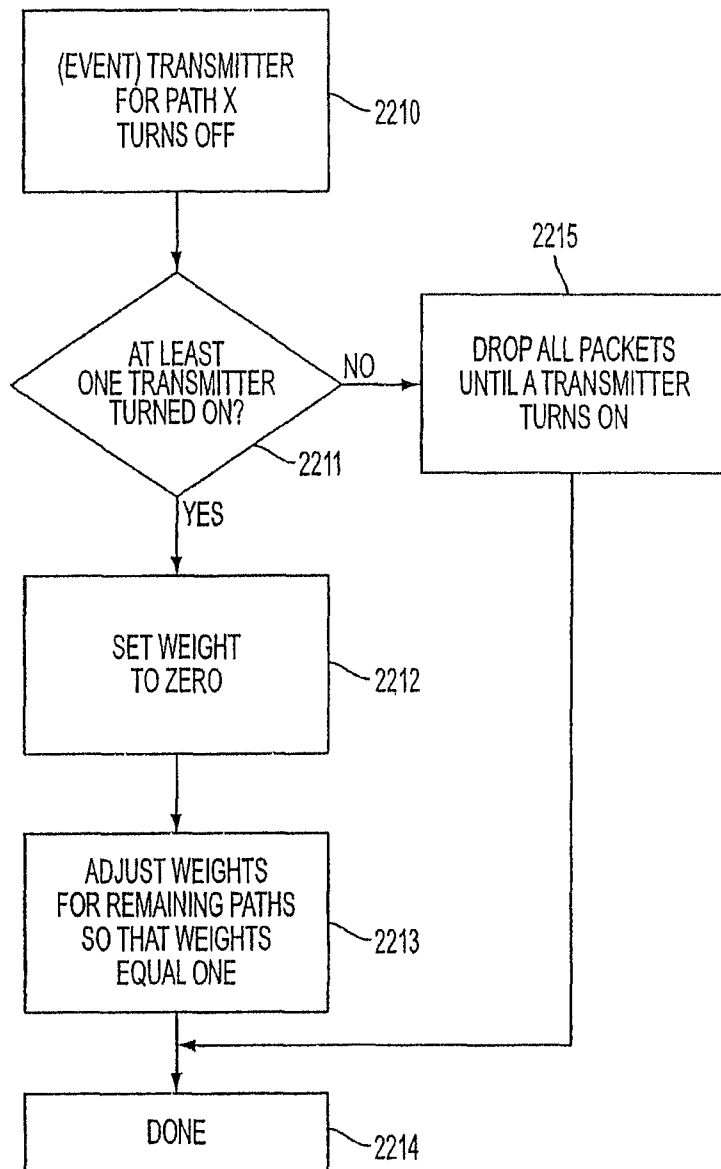


FIG. 22B




FIG. 23




FIG. 24




FIG. 25  
(PRIOR ART)




FIG. 26



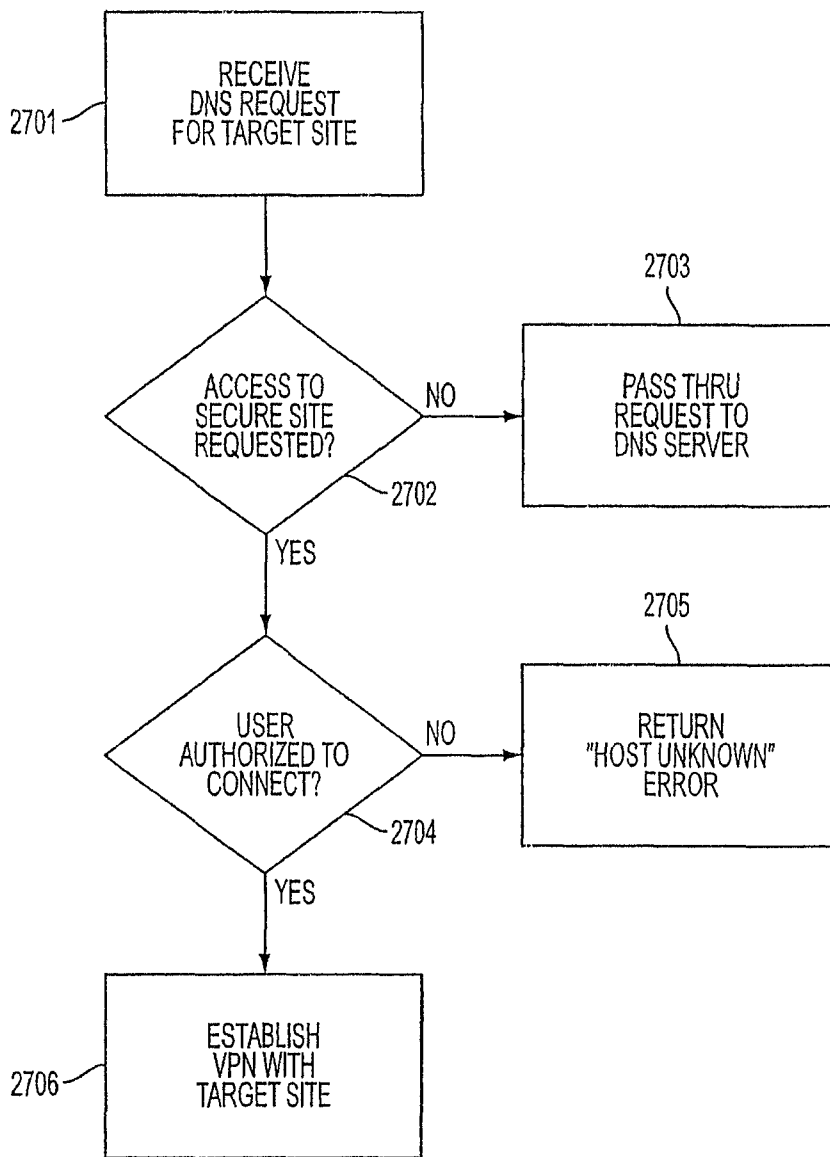


FIG. 27




FIG. 28




FIG. 29




FIG. 30




FIG. 31




FIG. 32




FIG. 33

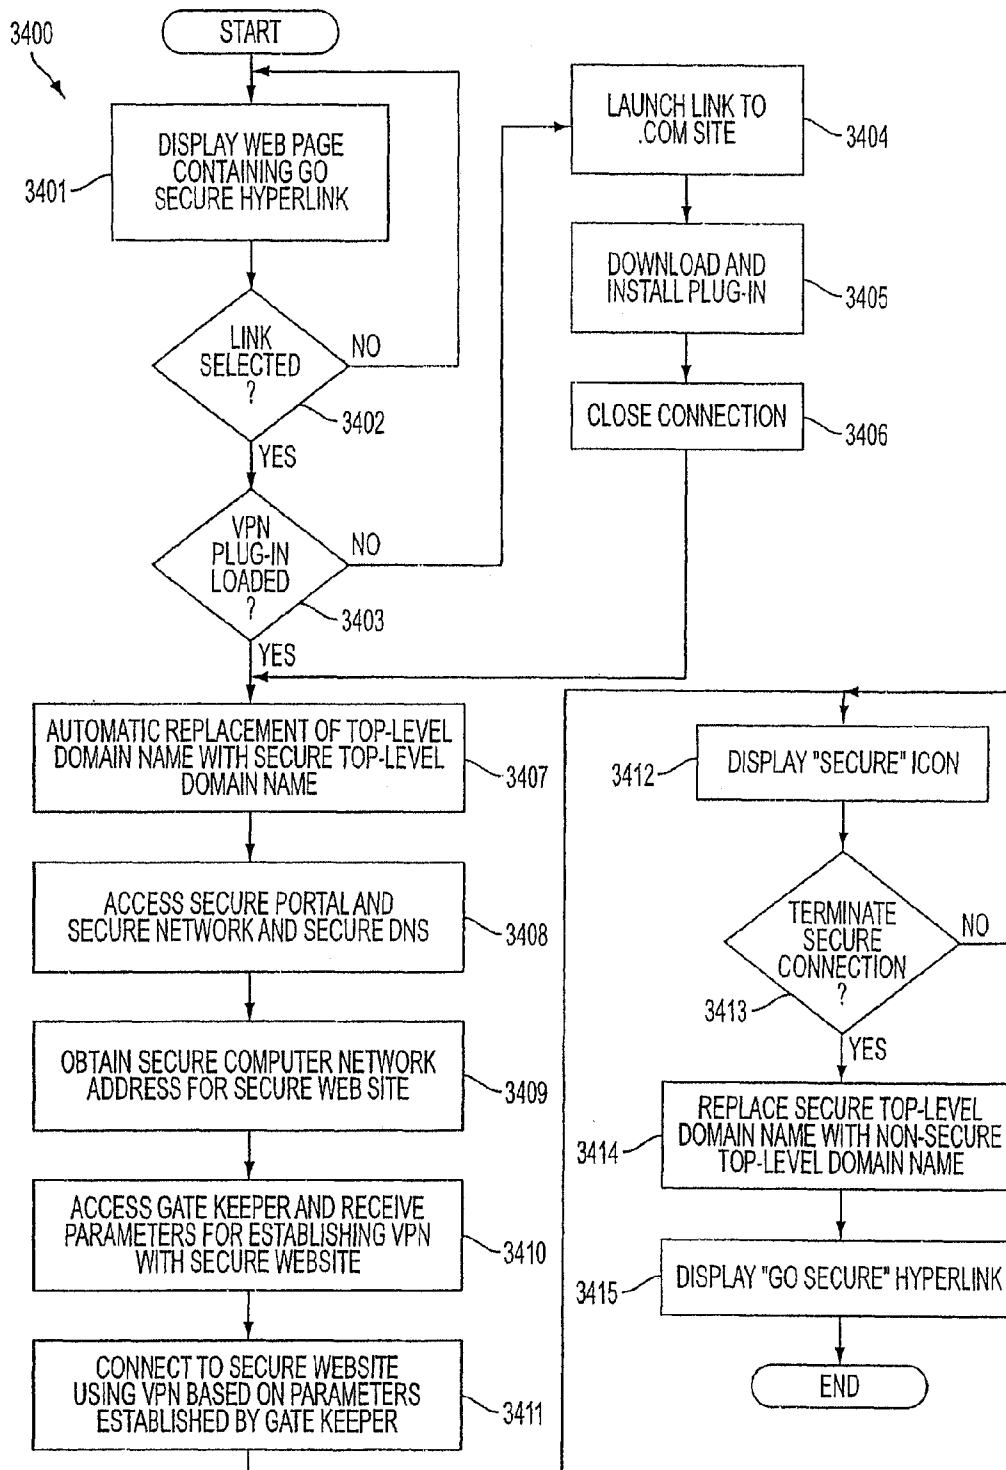


FIG. 34



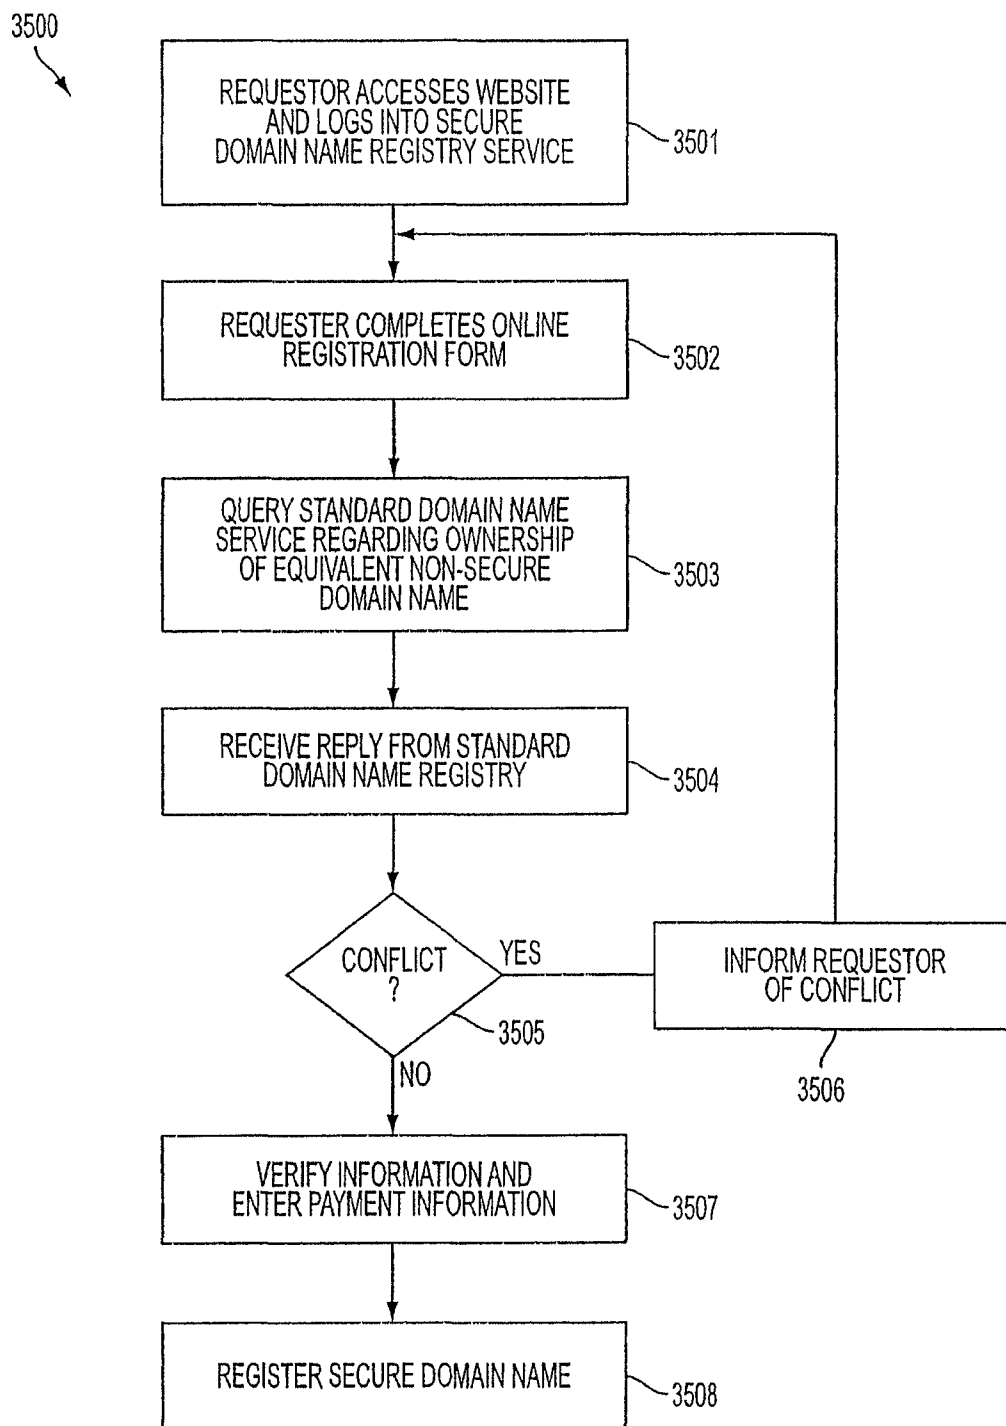


FIG. 35




FIG. 36




FIG. 37

**SYSTEM AND METHOD EMPLOYING AN  
AGILE NETWORK PROTOCOL FOR SECURE  
COMMUNICATIONS USING SECURE  
DOMAIN NAMES**

CROSS-REFERENCE TO RELATED  
APPLICATIONS

This application claims priority from and is a continuation of a co-pending U.S. Application No. 13/049,552, filed Mar. 16, 2011, which is a continuation of U.S. application Ser. No. 11/840,560, filed Aug. 17, 2007, now U.S. Pat. No. 7,921,211, which is a continuation of U.S. application Ser. No. 10/714,849, filed Nov. 18, 2003, now U.S. Pat. No. 7,418,504, which is a continuation of U.S. application Ser. No. 09/558,210, filed Apr. 26, 2000, now abandoned, which is a continuation-in-part of U.S. application Ser. No. 09/504,783, filed on Feb. 15, 2000, now U.S. Pat. No. 6,502,135, issued Dec. 31, 2002, which claims priority from and is a continuation-in-part patent application of previously-filed U.S. application Ser. No. 09/429,643, filed on Oct. 29, 1999, now U.S. Pat. No. 7,010,604, issued Mar. 7, 2006. The subject matter of U.S. application Ser. No. 09/429,643, which is bodily incorporated herein, derives from provisional U.S. Application Nos. 60/106,261 (filed Oct. 30, 1998) and 60/137,704 (filed Jun. 7, 1999). The present application is also related to U.S. application Ser. No. 09/558,209, filed Apr. 26, 2000, now abandoned, and which is incorporated by reference herein. Each of the above-mentioned applications is incorporated herein by reference in its entirety as though fully set forth herein.

BACKGROUND OF THE INVENTION

A tremendous variety of methods have been proposed and implemented to provide security and anonymity for communications over the Internet. The variety stems, in part, from the different needs of different Internet users. A basic heuristic framework to aid in discussing these different security techniques is illustrated in FIG. 1. Two terminals, an originating terminal **100** and a destination terminal **110** are in communication over the Internet. It is desired for the communications to be secure, that is, immune to eavesdropping. For example, terminal **100** may transmit secret information to terminal **110** over the Internet **107**. Also, it may be desired to prevent an eavesdropper from discovering that terminal **100** is in communication with terminal **110**. For example, if terminal **100** is a user and terminal **110** hosts a web site, terminal **100**'s user may not want anyone in the intervening networks to know what web sites he is "visiting." Anonymity would thus be an issue, for example, for companies that want to keep their market research interests private and thus would prefer to prevent outsiders from knowing which websites or other Internet resources they are "visiting." These two security issues may be called data security and anonymity, respectively.

Data security is usually tackled using some form of data encryption. An encryption key **48** is known at both the originating and terminating terminals **100** and **110**. The keys may be private and public at the originating and destination terminals **100** and **110**, respectively or they may be symmetrical keys (the same key is used by both parties to encrypt and decrypt). Many encryption methods are known and usable in this context.

To hide traffic from a local administrator or ISP, a user can employ a local proxy server in communicating over an encrypted channel with an outside proxy such that the local

administrator or ISP only sees the encrypted traffic. Proxy servers prevent destination servers from determining the identities of the originating clients. This system employs an intermediate server interposed between client and destination server. The destination server sees only the Internet Protocol (IP) address of the proxy server and not the originating client. The target server only sees the address of the outside proxy. This scheme relies on a trusted outside proxy server. Also, proxy schemes are vulnerable to traffic analysis methods of determining identities of transmitters and receivers. Another important limitation of proxy servers is that the server knows the identities of both calling and called parties. In many instances, an originating terminal, such as terminal **A**, would prefer to keep its identity concealed from the proxy, for example, if the proxy server is provided by an Internet service provider (ISP).

To defeat traffic analysis, a scheme called Chaum's mixes employs a proxy server that transmits and receives fixed length messages, including dummy messages. Multiple originating terminals are connected through a mix (a server) to multiple target servers. It is difficult to tell which of the originating terminals are communicating to which of the connected target servers, and the dummy messages confuse eavesdroppers' efforts to detect communicating pairs by analyzing traffic. A drawback is that there is a risk that the mix server could be compromised. One way to deal with this risk is to spread the trust among multiple mixes. If one mix is compromised, the identities of the originating and target terminals may remain concealed. This strategy requires a number of alternative mixes so that the intermediate servers interposed between the originating and target terminals are not determinable except by compromising more than one mix. The strategy wraps the message with multiple layers of encrypted addresses. The first mix in a sequence can decrypt only the outer layer of the message to reveal the next destination mix in sequence. The second mix can decrypt the message to reveal the next mix and so on. The target server receives the message and, optionally, a multi-layer encrypted payload containing return information to send data back in the same fashion. The only way to defeat such a mix scheme is to collude among mixes. If the packets are all fixed-length and intermixed with dummy packets, there is no way to do any kind of traffic analysis.

Still another anonymity technique, called 'crowds,' protects the identity of the originating terminal from the intermediate proxies by providing that originating terminals belong to groups of proxies called crowds. The crowd proxies are interposed between originating and target terminals. Each proxy through which the message is sent is randomly chosen by an upstream proxy. Each intermediate proxy can send the message either to another randomly chosen proxy in the "crowd" or to the destination. Thus, even crowd members cannot determine if a preceding proxy is the originator of the message or if it was simply passed from another proxy.

ZKS (Zero-Knowledge Systems) Anonymous IP Protocol allows users to select up to any of five different pseudonyms, while desktop software encrypts outgoing traffic and wraps it in User Datagram Protocol (UDP) packets. The first server in a 2+-hop system gets the UDP packets, strips off one layer of encryption to add another, then sends the traffic to the next server, which strips off yet another layer of encryption and adds a new one. The user is permitted to control the number of hops. At the final server, traffic is decrypted with an untraceable IP address. The technique is called onion-routing. This method can be defeated using traffic analysis. For a simple example, bursts of packets from a user during low-duty periods can reveal the identities of sender and receiver.

Firewalls attempt to protect LANs from unauthorized access and hostile exploitation or damage to computers connected to the LAN. Firewalls provide a server through which all access to the LAN must pass. Firewalls are centralized systems that require administrative overhead to maintain. They can be compromised by virtual-machine applications (“applets”). They instill a false sense of security that leads to security breaches for example by users sending sensitive information to servers outside the firewall or encouraging use of modems to sidestep the firewall security. Firewalls are not useful for distributed systems such as business travelers, extranets, small teams, etc.

#### SUMMARY OF THE INVENTION

A secure mechanism for communicating over the internet, including a protocol referred to as the Tunneled Agile Routing Protocol (TARP), uses a unique two-layer encryption format and special TARP routers. TARP routers are similar in function to regular IP routers. Each TARP router has one or more IP addresses and uses normal IP protocol to send IP packet messages (“packets” or “datagrams”). The IP packets exchanged between TARP terminals via TARP routers are actually encrypted packets whose true destination address is concealed except to TARP routers and servers. The normal or “clear” or “outside” IP header attached to TARP IP packets contains only the address of a next hop router or destination server. That is, instead of indicating a final destination in the destination field of the IP header, the TARP packet’s IP header always points to a next-hop in a series of TARP router hops, or to the final destination. This means there is no overt indication from an intercepted TARP packet of the true destination of the TARP packet since the destination could always be next-hop TARP router as well as the final destination.

Each TARP packet’s true destination is concealed behind a layer of encryption generated using a link key. The link key is the encryption key used for encrypted communication between the hops intervening between an originating TARP terminal and a destination TARP terminal. Each TARP router can remove the outer layer of encryption to reveal the destination router for each TARP packet. To identify the link key needed to decrypt the outer layer of encryption of a TARP packet, a receiving TARP or routing terminal may identify the transmitting terminal by the sender/receiver IP numbers in the cleartext IP header.

Once the outer layer of encryption is removed, the TARP router determines the final destination. Each TARP packet **140** undergoes a minimum number of hops to help foil traffic analysis. The hops may be chosen at random or by a fixed value. As a result, each TARP packet may make random trips among a number of geographically disparate routers before reaching its destination. Each trip is highly likely to be different for each packet composing a given message because each trip is independently randomly determined. This feature is called agile routing. The fact that different packets take different routes provides distinct advantages by making it difficult for an interloper to obtain all the packets forming an entire multi-packet message. The associated advantages have to do with the inner layer of encryption discussed below. Agile routing is combined with another feature that furthers this purpose; a feature that ensures that any message is broken into multiple packets.

The IP address of a TARP router can be changed, a feature called IP agility. Each TARP router, independently or under direction from another TARP terminal or router, can change its IP address. A separate, unchangeable identifier or address

is also defined. This address, called the TARP address, is known only to TARP routers and terminals and may be correlated at any time by a TARP router or a TARP terminal using a Lookup Table (LUT). When a TARP router or terminal changes its IP address, it updates the other TARP routers and terminals which in turn update their respective LUTs.

The message payload is hidden behind an inner layer of encryption in the TARP packet that can only be unlocked using a session key. The session key is not available to any of the intervening TARP routers. The session key is used to decrypt the payloads of the TARP packets permitting the data stream to be reconstructed.

Communication may be made private using link and session keys, which in turn may be shared and used according to any desired method. For example, public/private keys or symmetric keys may be used.

To transmit a data stream, a TARP originating terminal constructs a series of TARP packets from a series of IP packets generated by a network (IP) layer process. (Note that the terms “network layer,” “data link layer,” “application layer,” etc. used in this specification correspond to the Open Systems Interconnection (OSI) network terminology.) The payloads of these packets are assembled into a block and chain-block encrypted using the session key. This assumes, of course, that all the IP packets are destined for the same TARP terminal. The block is then interleaved and the interleaved encrypted block is broken into a series of payloads, one for each TARP packet to be generated. Special TARP headers  $IP_T$  are then added to each payload using the IP headers from the data stream packets. The TARP headers can be identical to normal IP headers or customized in some way. They should contain a formula or data for deinterleaving the data at the destination TARP terminal, a time-to-live (TTL) parameter to indicate the number of hops still to be executed, a data type identifier which indicates whether the payload contains, for example, TCP or UDP data, the sender’s TARP address, the destination TARP address, and an indicator as to whether the packet contains real or decoy data or a formula for filtering out decoy data if decoy data is spread in some way through the TARP payload data.

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

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

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

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

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

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

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

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

In various other embodiments of the invention, a scalable version of the system may be constructed in which a plurality of IP addresses are preassigned to each pair of communicating nodes in the network. Each pair of nodes agrees upon an algorithm for "hopping" between IP addresses (both sending and receiving), such that an eavesdropper sees apparently

continuously random IP address pairs (source and destination) for packets transmitted between the pair. Overlapping or "reusable" IP addresses may be allocated to different users on the same subnet, since each node merely verifies that a particular packet includes a valid source/destination pair from the agreed-upon algorithm. Source/destination pairs are preferably not reused between any two nodes during any given end-to-end session, though limited IP block sizes or lengthy sessions might require it.

Further improvements described in this continuation-in-part application include: (1) a load balancer that distributes packets across different transmission paths according to transmission path quality; (2) a DNS proxy server that transparently creates a virtual private network in response to a domain name inquiry; (3) a large-to-small link bandwidth management feature that prevents denial-of-service attacks at system chokepoints; (4) a traffic limiter that regulates incoming packets by limiting the rate at which a transmitter can be synchronized with a receiver; and (5) a signaling synchronizer that allows a large number of nodes to communicate with a central node by partitioning the communication function between two separate entities.

The present invention provides key technologies for implementing a secure virtual Internet by using a new agile network protocol that is built on top of the existing Internet protocol (IP). The secure virtual Internet works over the existing Internet infrastructure, and interfaces with client applications the same way as the existing Internet. The key technologies provided by the present invention that support the secure virtual Internet include a "one-click" and "no-click" technique to become part of the secure virtual Internet, a secure domain name service (SDNS) for the secure virtual Internet, and a new approach for interfacing specific client applications onto the secure virtual Internet. According to the invention, the secure domain name service interfaces with existing applications, in addition to providing a way to register and serve domain names and addresses.

According to one aspect of the present invention, a user can conveniently establish a VPN using a "one-click" or a "no-click" technique without being required to enter user identification information, a password and/or an encryption key for establishing a VPN. The advantages of the present invention are provided by a method for establishing a secure communication link between a first computer and a second computer over a computer network, such as the Internet. In one embodiment, a secure communication mode is enabled at a first computer without a user entering any cryptographic information for establishing the secure communication mode of communication, preferably by merely selecting an icon displayed on the first computer. Alternatively, the secure communication mode of communication can be enabled by entering a command into the first computer. Then, a secure communication link is established between the first computer and a second computer over a computer network based on the enabled secure communication mode of communication. According to the invention, it is determined whether a secure communication software module is stored on the first computer in response to the step of enabling the secure communication mode of communication. A predetermined computer network address is then accessed for loading the secure communication software module when the software module is not stored on the first computer. Subsequently, the proxy software module is stored in the first computer. The secure communication link is a virtual private network communication link over the computer network. Preferably, the virtual private network can be based on inserting into each data packet one or more data values that vary according to a pseudo-random

sequence. Alternatively, the virtual private network can be based on a computer network address hopping regime that is used to pseudorandomly change computer network addresses or other data values in packets transmitted between the first computer and the second computer, such that the second computer compares the data values in each data packet transmitted between the first computer and the second computer to a moving window of valid values. Yet another alternative provides that the virtual private network can be based on a comparison between a discriminator field in each data packet to a table of valid discriminator fields maintained for the first computer.

According to another aspect of the invention, a command is entered to define a setup parameter associated with the secure communication link mode of communication. Consequently, the secure communication mode is automatically established when a communication link is established over the computer network.

The present invention also provides a computer system having a communication link to a computer network, and a display showing a hyperlink for establishing a virtual private network through the computer network. When the hyperlink for establishing the virtual private network is selected, a virtual private network is established over the computer network. A non-standard top-level domain name is then sent over the virtual private network communication to a predetermined computer network address, such as a computer network address for a secure domain name service (SDNS).

The present invention provides a domain name service that provides secure computer network addresses for secure, non-standard top-level domain names. The advantages of the present invention are provided by a secure domain name service for a computer network that includes a portal connected to a computer network, such as the Internet, and a domain name database connected to the computer network through the portal. According to the invention, the portal authenticates a query for a secure computer network address, and the domain name database stores secure computer network addresses for the computer network. Each secure computer network address is based on a non-standard top-level domain name, such as .scm, .sorg, .snet, .snet, .sedu, .smil and .sint.

The present invention provides a way to encapsulate existing application network traffic at the application layer of a client computer so that the client application can securely communicate with a server protected by an agile network protocol. The advantages of the present invention are provided by a method for communicating using a private communication link between a client computer and a server computer over a computer network, such as the Internet. According to the invention, an information packet is sent from the client computer to the server computer over the computer network. The information packet contains data that is inserted into the payload portion of the packet at the application layer of the client computer and is used for forming a virtual private connection between the client computer and the server computer. The modified information packet can be sent through a firewall before being sent over the computer network to the server computer and by working on top of existing protocols (i.e., UDP, ICMP and TCP), the present invention more easily penetrates the firewall. The information packet is received at a kernel layer of an operating system on the server side. It is then determined at the kernel layer of the operating system on the host computer whether the information packet contains the data that is used for forming the virtual private connection. The server side replies by sending an information packet to the client computer that has been

modified at the kernel layer to containing virtual private connection information in the payload portion of the reply information packet. Preferably, the information packet from the client computer and the reply information packet from the server side are each a UDP protocol information packet. Alternative, both information packets could be a TCP/IP protocol information packet, or an ICMP protocol information packet.

#### BRIEF DESCRIPTION OF THE DRAWINGS

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

FIG. 20 shows a system employing distributed transmission paths.

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

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

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

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

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

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

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

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

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

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

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

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

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

FIG. 33 shows a system block diagram of a computer network in which the "one-click" secure communication link of the present invention is suitable for use.

FIG. 34 shows a flow diagram for installing and establishing a "one-click" secure communication link over a computer network according to the present invention.

FIG. 35 shows a flow diagram for registering a secure domain name according to the present invention.

FIG. 36 shows a system block diagram of a computer network in which a private connection according to the present invention can be configured to more easily traverse a firewall between two computer networks.

FIG. 37 shows a flow diagram for establishing a virtual private connection that is encapsulated using an existing network protocol.

#### DETAILED DESCRIPTION OF THE INVENTION

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

Each TARP packet's true destination is concealed behind an outer layer of encryption generated using a link key 146. The link key 146 is the encryption key used for encrypted communication between the end points (TARP terminals or TARP routers) of a single link in the chain of hops connecting the originating TARP terminal 100 and the destination TARP terminal 110. Each TARP router 122-127, using the link key 146 it uses to communicate with the previous hop in a chain,

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

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

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

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



## 11

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

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

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

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

1. A window sequence number—an identifier that indicates where the packet belongs in the original message sequence.
2. An interleave sequence number—an identifier that indicates the interleaving sequence used to form the packet so that the packet can be deinterleaved along with other packets in the interleave window.
3. A time-to-live (TTL) datum—indicates the number of TARP-router-hops to be executed before the packet reaches its destination. Note that the TTL parameter may provide a datum to be used in a probabilistic formula for determining whether to route the packet to the destination or to another hop.
4. Data type identifier—indicates whether the payload contains, for example, TCP or UDP data.
5. Sender's address—indicates the sender's address in the TARP network.
6. Destination address—indicates the destination terminal's address in the TARP network.
7. Decoy/Real—an indicator of whether the packet contains real message data or dummy decoy data or a combination.

## 12

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

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

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

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

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

The above scheme may be implemented entirely by processes operating between the data link layer and the network layer of each server or terminal participating in the TARP system. Referring to FIG. **4**, a TARP transceiver **405** can be an originating terminal **100**, a destination terminal **110**, or a TARP router **122-127**. In each TARP Transceiver **405**, a transmitting process is generated to receive normal packets from the Network (IP) layer and generate TARP packets for communication over the network. A receiving process is generated to receive normal IP packets containing TARP packets and generate from these normal IP packets which are "passed

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

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

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

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

Upon detection of an attack, the TARP process may also create a subprocess that maintains the original IP address and continues interacting with the attacker. The latter may provide an opportunity to trace the attacker or study the attacker’s methods (called “fishbowling” drawing upon the analogy of a small fish in a fish bowl that “thinks” it is in the ocean but is actually under captive observation). A history of the communication between the attacker and the abandoned (fish-bowled) IP address can be recorded or transmitted for human analysis or further synthesized for purposes of responding in some way.

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

Decoy packets may be generated by each TARP terminal 100, 110 or each router 122-127 on some basis determined by an algorithm. For example, the algorithm may be a random one which calls for the generation of a packet on a random basis when the terminal is idle. Alternatively, the algorithm may be responsive to time of day or detection of low traffic to generate more decoy packets during low traffic times. Note that packets are preferably generated in groups, rather than one by one, the groups being sized to simulate real messages. In addition, so that decoy packets may be inserted in normal TARP message streams, the background loop may have a latch that makes it more likely to insert decoy packets when a message stream is being received. That is, when a series of messages are received, the decoy packet generation rate may be increased. Alternatively, if a large number of decoy packets is received along with regular TARP packets, the algorithm may increase the rate of dropping of decoy packets rather than forwarding them. The result of dropping and generating decoy packets in this way is to make the apparent incoming message size different from the apparent outgoing message size to help foil traffic analysis. The rate of reception of packets, decoy or otherwise, may be indicated to the decoy packet dropping and generating processes through perishable decoy and regular packet counters. (A perishable counter is one that resets or decrements its value in response to time so that it contains a high value when it is incremented in rapid succession and a small value when incremented either slowly or a small number of times in rapid succession.) Note that destination TARP terminal 110 may generate decoy packets equal in number and size to those TARP packets received to make it appear it is merely routing packets and is therefore not the destination terminal.

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

- S0. A background loop operation is performed which applies an algorithm which determines the generation of decoy IP packets. The loop is interrupted when an encrypted TARP packet is received.
- S2. The TARP packet may be probed in some way to authenticate the packet before attempting to decrypt it using the link key. That is, the router may determine that the packet is an authentic TARP packet by performing a selected operation on some data included with the clear IP header attached to the encrypted TARP packet contained in the payload. This makes it possible to avoid performing decryption on packets that are not authentic TARP packets.

15

S3. The TARP packet is decrypted to expose the destination TARP address and an indication of whether the packet is a decoy packet or part of a real message.

S4. If the packet is a decoy packet, the perishable decoy counter is incremented.

S5. Based on the decoy generation/dropping algorithm and the perishable decoy counter value, if the packet is a decoy packet, the router may choose to throw it away. If the received packet is a decoy packet and it is determined that it should be thrown away (S6), control returns to step S0.

S7. The TTL parameter of the TARP header is decremented and it is determined if the TTL parameter is greater than zero.

S8. If the TTL parameter is greater than zero, a TARP address is randomly chosen from a list of TARP addresses maintained by the router and the link key and IP address corresponding to that TARP address memorized for use in creating a new IP packet containing the TARP packet.

S9. If the TTL parameter is zero or less, the link key and IP address corresponding to the TARP address of the destination are memorized for use in creating the new IP packet containing the TARP packet.

S10. The TARP packet is encrypted using the memorized link key.

S11. An IP header is added to the packet that contains the stored IP address, the encrypted TARP packet wrapped with an IP header, and the completed packet transmitted to the next hop or destination.

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

S20. A background loop operation applies an algorithm that determines the generation of decoy IP packets. The loop is interrupted when a data stream containing IP packets is received for transmission.

S21. The received IP packets are grouped into a set consisting of messages with a constant IP destination address. The set is further broken down to coincide with a maximum size of an interleave window. The set is encrypted, and interleaved into a set of payloads destined to become TARP packets.

S22. The TARP address corresponding to the IP address is determined from a lookup table and stored to generate the TARP header. An initial TTL count is generated and stored in the header. The TTL count may be random with minimum and maximum values or it may be fixed or determined by some other parameter.

S23. The window sequence numbers and interleave sequence numbers are recorded in the TARP headers of each packet.

S24. One TARP router address is randomly chosen for each TARP packet and the IP address corresponding to it stored for use in the clear IP header. The link key corresponding to this router is identified and used to encrypt TARP packets containing interleaved and encrypted data and TARP headers.

S25. A clear IP header with the first hop router's real IP address is generated and added to each of the encrypted TARP packets and the resulting packets.

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

S40. A background loop operation is performed which applies an algorithm which determines the generation of

16

decoy IP packets. The loop is interrupted when an encrypted TARP packet is received.

S42. The TARP packet may be probed to authenticate the packet before attempting to decrypt it using the link key.

S43. The TARP packet is decrypted with the appropriate link key to expose the destination TARP address and an indication of whether the packet is a decoy packet or part of a real message.

S44. If the packet is a decoy packet, the perishable decoy counter is incremented.

S45. Based on the decoy generation/dropping algorithm and the perishable decoy counter value, if the packet is a decoy packet, the receiver may choose to throw it away.

S46. The TARP packets are cached until all packets forming an interleave window are received.

S47. Once all packets of an interleave window are received, the packets are deinterleaved.

S48. The packets block of combined packets defining the interleave window is then decrypted using the session key.

S49. The decrypted block is then divided using the window sequence data and the IP<sub>T</sub> headers are converted into normal IP<sub>C</sub> headers. The window sequence numbers are integrated in the IP<sub>C</sub> headers.

S50. The packets are then handed up to the IP layer processes.

#### 1. Scalability Enhancements

The IP agility feature described above relies on the ability to transmit IP address changes to all TARP routers. The embodiments including this feature will be referred to as "boutique" embodiments due to potential limitations in scaling these features up for a large network, such as the Internet. (The "boutique" embodiments would, however, be robust for use in smaller networks, such as small virtual private networks, for example). One problem with the boutique embodiments is that if IP address changes are to occur frequently, the message traffic required to update all routers sufficiently quickly creates a serious burden on the Internet when the TARP router and/or client population gets large. The bandwidth burden added to the networks, for example in ICMP packets, that would be used to update all the TARP routers could overwhelm the Internet for a large scale implementation that approached the scale of the Internet. In other words, the boutique system's scalability is limited.

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

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

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

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

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

When the router receives the client's packet, it compares the send and receive IP addresses of the packet with the next N predicted send and receive IP address pairs and rejects the packet if it is not a member of this set. Received packets that do not have the predicted source/destination IP addresses falling within the window are rejected, thus thwarting possible hackers. (With the number of possible combinations, even a fairly large window would be hard to fall into at random.) If it is a member of this set, the router accepts the packet and processes it further. This link-based IP-hopping strategy, referred to as "IHOP," is a network element that stands on its own and is not necessarily accompanied by elements of the boutique system described above. If the routing agility feature described in connection with the boutique embodiment is combined with this link-based IP-hopping strategy, the router's next step would be to decrypt the TARP header to determine the destination TARP router for the packet and determine what should be the next hop for the packet. The TARP router would then forward the packet to a random TARP

router or the destination TARP router with which the source TARP router has a link-based IP hopping communication established.

FIG. 8 shows how a client computer **801** and a TARP router **811** can establish a secure session. When client **801** seeks to establish an IHOP session with TARP router **811**, the client **801** sends "secure synchronization" request ("SSYN") packet **821** to the TARP router **811**. This SYN packet **821** contains the client's **801** authentication token, and may be sent to the router **811** in an encrypted format. The source and destination IP numbers on the packet **821** are the client's **801** current fixed IP address, and a "known" fixed IP address for the router **811**. (For security purposes, it may be desirable to reject any packets from outside of the local network that are destined for the router's known fixed IP address.) Upon receipt and validation of the client's **801** SSYN packet **821**, the router **811** responds by sending an encrypted "secure synchronization acknowledgment" ("SSYN ACK") **822** to the client **801**. This SSYN ACK **822** will contain the transmit and receive hopblocks that the client **801** will use when communicating with the TARP router **811**. The client **801** will acknowledge the TARP router's **811** response packet **822** by generating an encrypted SSYN ACK ACK packet **823** which will be sent from the client's **801** fixed IP address and to the TARP router's **811** known fixed IP address. The client **801** will simultaneously generate a SSYN ACK ACK packet; this SSYN ACK packet, referred to as the Secure Session Initiation (SSI) packet **824**, will be sent with the first {sender, receiver} IP pair in the client's transmit table **921** (FIG. 9), as specified in the transmit hopblock provided by the TARP router **811** in the SSYN ACK packet **822**. The TARP router **811** will respond to the SSI packet **824** with an SSI ACK packet **825**, which will be sent with the first {sender, receiver} IP pair in the TARP router's transmit table **923**. Once these packets have been successfully exchanged, the secure communications session is established, and all further secure communications between the client **801** and the TARP router **811** will be conducted via this secure session, as long as synchronization is maintained. If synchronization is lost, then the client **801** and TARP router **802** may re-establish the secure session by the procedure outlined in FIG. 8 and described above.

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

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

19

from its transmit table **923** when constructing a packet to send to the client **901**, and the client **901** will maintain a look-ahead buffer of expected IP pairs on packets that it is receiving. Each TARP router will maintain separate pairs of transmit and receive tables for each client that is currently engaged in a secure session with or through that TARP router.

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

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

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

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

As another extension to the stated architecture, multiple physical paths can be used by a client, in order to provide link redundancy and further thwart attempts at denial of service

20

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

## 2. Further Extensions

The following describes various extensions to the techniques, systems, and methods described above. As described above, the security of communications occurring between computers in a computer network (such as the Internet, an Ethernet, or others) can be enhanced by using seemingly random source and destination Internet Protocol (IP) addresses for data packets transmitted over the network. This feature prevents eavesdroppers from determining which computers in the network are communicating with each other while permitting the two communicating computers to easily recognize whether a given received data packet is legitimate or not. In one embodiment of the above-described systems, an IP header extension field is used to authenticate incoming packets on an Ethernet.

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

### A. Hardware Address Hopping

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

It may be possible for a nefarious listener to acquire information about the contents of a frame and/or its communicants by examining frames on a local network rather than (or in addition to) the IP packets themselves. This is especially true in broadcast media, such as Ethernet, where it is necessary to insert into the frame header the hardware address of the machine that generated the frame and the hardware address of the machine to which frame is being sent. All nodes on the network can potentially “see” all packets transmitted across the network. This can be a problem for secure communica-

21

tions, especially in cases where the communicants do not want for any third party to be able to identify who is engaging in the information exchange. One way to address this problem is to push the address-hopping scheme down to the hardware layer. In accordance with various embodiments of the invention, hardware addresses are “hopped” in a manner similar to that used to change IP addresses, such that a listener cannot determine which hardware node generated a particular message nor which node is the intended recipient.

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

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

The lowest levels of communication software 1204 and 1217 communicate with hardware components 1206 and 1214 respectively, each of which can include one or more registers 1207 and 1215 that allow the hardware to be reconfigured or controlled in accordance with various communication protocols. The hardware components (an Ethernet network interface card, for example) communicate with each other over the communication medium. Each hardware component is typically pre-assigned a fixed hardware address or MAC number that identifies the hardware component to other nodes on the network. One or more interface drivers control the operation of each card and can, for example, be configured to accept or reject packets from certain hardware addresses. As will be described in more detail below, various embodiments of the inventive principles provide for “hopping” different addresses using one or more algorithms and one or more moving windows that track a range of valid addresses to validate received packets. Packets transmitted according to one or more of the inventive principles will be generally referred to as “secure” packets or “secure communications” to differentiate them from ordinary data packets that are transmitted in the clear using ordinary, machine-correlated addresses.

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

This approach, however, runs the risk of using MAC addresses that are currently active on the LAN—which, in turn, could interrupt communications for those machines.

22

Since an Ethernet MAC address is at present 48 bits in length, the chance of randomly misusing an active MAC address is actually quite small. However, if that figure is multiplied by a large number of nodes (as would be found on an extensive LAN), by a large number of frames (as might be the case with packet voice or streaming video), and by a large number of concurrent Virtual Private Networks (VPNs), then the chance that a non-secure machine’s MAC address could be used in an address-hopped frame can become non-trivial. In short, any scheme that runs even a small risk of interrupting communications for other machines on the LAN is bound to receive resistance from prospective system administrators. Nevertheless, it is technically feasible, and can be implemented without risk on a LAN on which there is a small number of machines, or if all of the machines on the LAN are engaging in MAC-hopped communications.

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

One disadvantage of purely-random MAC address hopping is its impact on processing overhead; that is, since every incident frame must be processed, the machine’s CPU is engaged considerably more often than if the network interface driver is discriminating and rejecting packets unilaterally. A compromise approach is to select either a single fixed MAC address or a small number of MAC addresses (e.g., one for each virtual private network on an Ethernet) to use for MAC-hopped communications, regardless of the actual recipient for which the message is intended. In this mode, the network interface driver can check each incident frame against one (or a few) pre-established MAC addresses, thereby freeing the CPU from the task of physical-layer packet discrimination. This scheme does not betray any useful information to an interloper on the LAN; in particular, every secure packet can already be identified by a unique packet type in the outer header. However, since all machines engaged in secure communications would either be using the same MAC address, or be selecting from a small pool of predetermined MAC addresses, the association between a specific machine and a specific MAC address is effectively broken.

In this scheme, the CPU will be engaged more often than it would be in non-secure communications (or in synchronized MAC address hopping), since the network interface driver cannot always unilaterally discriminate between secure packets that are destined for that machine, and secure packets from

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

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

Reference will now be made to FIGS. 12A and 12B in order to describe the many combinations and features that follow the inventive principles. As explained above, two computer nodes 1201 and 1202 are assumed to be communicating over a network or communication medium such as an Ethernet. A communication protocol in each node (1204 and 1217, respectively) contains a modified element 1205 and 1216 that performs certain functions that deviate from the standard communication protocols. In particular, computer node 1201 implements a first “hop” algorithm 1208X that selects seemingly random source and destination IP addresses (and, in one embodiment, seemingly random IP header discriminator fields) in order to transmit each packet to the other computer node. For example, node 1201 maintains a transmit table 1208 containing triplets of source (S), destination (D), and discriminator fields (DS) that are inserted into outgoing IP packet headers. The table is generated through the use of an appropriate algorithm (e.g., a random number generator that is seeded with an appropriate seed) that is known to the recipient node 1202. As each new IP packet is formed, the next sequential entry out of the sender’s transmit table 1208 is used to populate the IP source, IP destination, and IP header extension field (e.g., discriminator field). It will be appreciated that the transmit table need not be created in advance but could instead be created on-the-fly by executing the algorithm when each packet is formed.

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

as invalid. The length of window W3 can be adjusted as necessary to reflect network delays or other factors.

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

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

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

In a second mode referred to as “promiscuous per VPN” mode, a small set of fixed hardware addresses are used, with a fixed source/destination hardware address used for all nodes communicating over a virtual private network. For example, if there are six nodes on an Ethernet, and the network is to be split up into two private virtual networks such that nodes on one VPN can communicate with only the other two nodes on its own VPN, then two sets of hardware addresses could be used: one set for the first VPN and a second set for the second VPN. This would reduce the amount of overhead involved in

checking for valid frames since only packets arriving from the designated VPN would need to be checked. IP addresses and one or more discriminator fields could still be hopped as before for secure communication within the VPN. Of course, this solution compromises the anonymity of the VPNs (i.e., an outsider can easily tell what traffic belongs in which VPN, though he cannot correlate it to a specific machine/person). It also requires the use of a discriminator field to mitigate the vulnerability to certain types of DoS attacks, (For example, without the discriminator field, an attacker on the LAN could stream frames containing the MAC addresses being used by the VPN; rejecting those frames could lead to excessive processing overhead. The discriminator field would provide a low-overhead means of rejecting the false packets.)

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

### B. Extending the Address Space

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

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

### C. Synchronization Techniques

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

One possible technique is to require that each node provide an acknowledgment upon successful receipt of each packet and, if no acknowledgment is received within a certain period

of time, to re-send the unacknowledged packet. This approach, however, drives up overhead costs and may be prohibitive in high-throughput environments such as streaming video or audio, for example.

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

In one embodiment, a “sync field” is added to the header of each packet sent out by the sender. This sync field could appear in the clear or as part of an encrypted portion of the packet. Assuming that a sender and receiver have selected a random-number generator (RNG) and seed value, this combination of RNG and seed can be used to generate a random-number sequence (RNS). The RNS is then used to generate a sequence of source/destination IP pairs (and, if desired, discriminator fields and hardware source and destination addresses), as described above. It is not necessary, however, to generate the entire sequence (or the first N-1 values) in order to generate the Nth random number in the sequence; if the sequence index N is known, the random value corresponding to that index can be directly generated (see below). Different RNGs (and seeds) with different fundamental periods could be used to generate the source and destination IP sequences, but the basic concepts would still apply. For the sake of simplicity, the following discussion will assume that IP source and destination address pairs (only) are hopped using a single RNG sequencing mechanism.

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

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



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

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

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

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

An important consideration in this architecture is the concept of “future” and “past” where the public sync values are concerned. Though the sync values, themselves, should be random to prevent spoofing attacks, it may be important that

the receiver be able to quickly identify a sync value that has already been sent—even if the packet containing that sync value was never actually received by the receiver. One solution is to hash a time stamp or sequence number into the public sync portion, which could be quickly extracted, checked, and discarded, thereby validating the public sync portion itself.

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

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

#### D. Other Synchronization Schemes

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

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

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

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

1. In the transmitter, ckpt\_o (“checkpoint old”) is the IP pair that was used to re-send the last SYNC\_REQ packet to the receiver. In the receiver, ckpt\_o (“checkpoint old”) is the IP pair that receives repeated SYNC\_REQ packets from the transmitter.
2. In the transmitter, ckpt\_n (“checkpoint new”) is the IP pair that will be used to send the next SYNC\_REQ packet to the receiver. In the receiver, ckpt\_n (“checkpoint new”) is the IP pair that receives a new SYNC\_REQ packet from the transmitter and which causes the receiver’s window to be re-aligned, ckpt\_o set to ckpt\_n, a new ckpt\_n to be generated and a new ckpt\_r to be generated.
3. In the transmitter, ckpt\_r is the IP pair that will be used to send the next SYNC\_ACK packet to the receiver. In the receiver, ckpt\_r is the IP pair that receives a new SYNC\_C\_ACK packet from the transmitter and which causes a new ckpt\_n to be generated. Since SYNC\_ACK is transmitted from the receiver ISP to the sender ISP, the trans-

mmitter ckpt\_r refers to the ckpt\_r of the receiver and the receiver ckpt\_r refers to the ckpt\_r of the transmitter (see FIG. 14).

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

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

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

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

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

E. Random Number Generator with a Jump-Ahead Capability

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

Linear congruential random number generators (LCRs) are fast, simple and well characterized random number generators that can be made to jump ahead n steps efficiently. An LCR generates random numbers X<sub>1</sub>, X<sub>2</sub>, X<sub>3</sub> . . . X<sub>k</sub> starting with seed X<sub>0</sub> using a recurrence

X<sub>i</sub>=(a X<sub>i-1</sub>+b) mod c, (1)

where a, b and c define a particular LCR. Another expression for X<sub>i</sub>,

X<sub>i</sub>=(a<sup>i</sup>(X<sub>0</sub>+b)-b)/(a-1) mod c (2)

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

X<sub>i</sub>=(a<sup>i</sup>(X<sub>0</sub>(a-1)+b)-b)/(a-1) mod c. (3)

It can be shown that:

(a<sup>i</sup>(X<sub>0</sub>(a-1)+b)-b)/(a-1) mod c=((a<sup>i</sup> mod ((a-1)c)(X<sub>0</sub> (a-1)+b)-b)/(a-1)) mod c (4).

(X<sub>0</sub>(a-1)+b) can be stored as (X<sub>0</sub>(a-1)+b) mod c, b as b mod c and compute a<sup>i</sup> mod ((a-1)c) (this requires O(log (i)) steps).

A practical implementation of this algorithm would jump a fixed distance, n, between synchronizations; this is tantamount to synchronizing every n packets. The window would commence n IP pairs from the start of the previous window. Using X<sub>j</sub><sup>w</sup>, the random number at the j<sup>th</sup> checkpoint, as X<sub>0</sub> and n as i, a node can store a<sup>n</sup> mod ((a-1)c) once per LCR and set

X<sub>j+1</sub><sup>w</sup>=X<sub>n(j+i)</sub>=(a<sup>n</sup> mod ((a-1)c)(X<sub>j</sub><sup>w</sup>(a-1)+b)-b)/(a-1) mod c, (5)

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

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

Random number generators can be cryptographically insecure. An adversary can derive the RNG parameters by examining the output or part of the output. This is true of LCGs. This vulnerability can be mitigated by incorporating an encryptor, designed to scramble the output as part of the random number generator. The random number generator prevents an adversary from mounting an attack—e.g., a known plaintext attack—against the encryptor.

F. Random Number Generator Example

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

X<sub>i</sub>=(31 X<sub>i-1</sub>+4) mod 15. (6)

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

(91 (X<sub>0</sub>30+4)-4)/30 mod 15 (7).

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

TABLE 1

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

G. Fast Packet Filter

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

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

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

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

H. Presence Vector Algorithm

A presence vector is a bit vector of length  $2^n$  that can be indexed by n-bit numbers (each ranging from 0 to  $2^n - 1$ ). One can indicate the presence of k n-bit numbers (not necessarily unique), by setting the bits in the presence vector indexed by

each number to 1. Otherwise, the bits in the presence vector are 0. An n-bit number, x, is one of the k numbers if and only if the  $x^{th}$  bit of the presence vector is 1. A fast packet filter can be implemented by indexing the presence vector and looking for a 1, which will be referred to as the “test.”

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

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

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

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

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

I. Further Synchronization Enhancements

A slightly modified form of the synchronization techniques described above can be employed. The basic principles of the previously described checkpoint synchronization scheme remain unchanged. The actions resulting from the reception of the checkpoints are, however, slightly different. In this variation, the receiver will maintain between OoO (“Out of Order”) and  $2 \times \text{WINDOW\_SIZE} + \text{OoO}$  active addresses ( $1 \leq \text{OoO} \leq \text{WINDOW\_SIZE}$  and  $\text{WINDOW\_SIZE} \geq 1$ ). OoO and WINDOW\_SIZE are engineerable parameters, where OoO is the minimum number of addresses needed to accommodate lost packets due to events in the network or out of order arrivals and WINDOW\_SIZE is the number of packets

transmitted before a SYNC\_REQ is issued. FIG. 17 depicts a storage array for a receiver's active addresses.

The receiver starts with the first  $2 \times \text{WINDOW\_SIZE}$  addresses loaded and active (ready to receive data). As packets are received, the corresponding entries are marked as "used" and are no longer eligible to receive packets. The transmitter maintains a packet counter, initially set to 0, containing the number of data packets transmitted since the last initial transmission of a SYNC\_REQ for which SYNC\_ACK has been received. When the transmitter packet counter equals WINDOW\_SIZE, the transmitter generates a SYNC\_REQ and does its initial transmission. When the receiver receives a SYNC\_REQ corresponding to its current CKPT\_N, it generates the next WINDOW\_SIZE addresses and starts loading them in order starting at the first location after the last active address wrapping around to the beginning of the array after the end of the array has been reached. The receiver's array might look like FIG. 18 when a SYNC\_REQ has been received. In this case a couple of packets have been either lost or will be received out of order when the SYNC\_REQ is received.

FIG. 19 shows the receiver's array after the new addresses have been generated. If the transmitter does not receive a SYNC\_ACK, it will re-issue the SYNC\_REQ at regular intervals. When the transmitter receives a SYNC\_ACK, the packet counter is decremented by WINDOW\_SIZE. If the packet counter reaches  $2 \times \text{WINDOW\_SIZE} - \text{OoO}$  then the transmitter ceases sending data packets until the appropriate SYNC\_ACK is finally received. The transmitter then resumes sending data packets. Future behavior is essentially a repetition of this initial cycle. The advantages of this approach are:

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

#### J. Distributed Transmission Path Variant

Another embodiment incorporating various inventive principles is shown in FIG. 20. In this embodiment, a message transmission system includes a first computer 2001 in communication with a second computer 2002 through a network 2011 of intermediary computers. In one variant of this embodiment, the network includes two edge routers 2003 and 2004 each of which is linked to a plurality of Internet Service Providers (ISPs) 2005 through 2010. Each ISP is coupled to a plurality of other ISPs in an arrangement as shown in FIG. 20, which is a representative configuration only and is not intended to be limiting. Each connection between ISPs is labeled in FIG. 20 to indicate a specific physical transmission path (e.g., AD is a physical path that links ISP A (element 2005) to ISP D (element 2008)). Packets arriving at each edge router are selectively transmitted to one of the ISPs to which the router is attached on the basis of a randomly or quasi-randomly selected basis.

As shown in FIG. 21, computer 2001 or edge router 2003 incorporates a plurality of link transmission tables 2100 that identify, for each potential transmission path through the network, valid sets of IP addresses that can be used to transmit the packet. For example, AD table 2101 contains a plurality of IP source/destination pairs that are randomly or quasi-randomly generated. When a packet is to be transmitted from first

computer 2001 to second computer 2002, one of the link tables is randomly (or quasi-randomly) selected, and the next valid source/destination address pair from that table is used to transmit the packet through the network. If path AD is randomly selected, for example, the next source/destination IP address pair (which is pre-determined to transmit between ISP A (element 2005) and ISP B (element 2008)) is used to transmit the packet. If one of the transmission paths becomes degraded or inoperative, that link table can be set to a "down" condition as shown in table 2105, thus preventing addresses from being selected from that table. Other transmission paths would be unaffected by this broken link.

### 3. Continuation-in-Part Improvements

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

#### A. Load Balancer

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

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

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

When the quality of a path degrades to where the transmitter is turned off by the synchronization function (i.e., no packets are arriving at the destination), the weight is set to zero. If all transmitters are turned off, no packets are sent.

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

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

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

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

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

Beginning in step 2201, the transmission quality of a given transmission path is measured. As described above, this measurement can be based on a comparison between the number of packets transmitted over a particular link to the number of packet acknowledgements received over the link (e.g., per unit time, or in absolute terms). Alternatively, the quality can be evaluated by comparing the number of packets that are acknowledged within the transmission window to the number of packets that were transmitted within that window. In yet another variation, the number of missed synchronization messages can be used to indicate link quality. Many other variations are of course possible.

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

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

If in step 2203 the quality of the path was greater than or equal to the threshold, then in step 2204 a check is made to

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

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

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

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

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

FIG. 23 shows a computer node 2301 employing various principles of the above-described embodiments. It is assumed that two computer nodes of the type shown in FIG. 23 communicate over a plurality of separate physical transmission paths. As shown in FIG. 23, four transmission paths X1 through X4 are defined for communicating between the two nodes. Each node includes a packet transmitter 2302 that operates in accordance with a transmit table 2308 as described above. (The packet transmitter could also operate without using the IP-hopping features described above, but the following description assumes that some form of hopping is employed in conjunction with the path selection mechanism.) The computer node also includes a packet receiver 2303 that operates in accordance with a receive table 2309, including a moving window W that moves as valid packets are received. Invalid packets having source and destination addresses that do not fall within window W are rejected.

As each packet is readied for transmission, source and destination IP addresses (or other discriminator values) are selected from transmit table 2308 according to any of the various algorithms described above, and packets containing these source/destination address pairs, which correspond to the node to which the four transmission paths are linked, are generated to a transmission path switch 2307. Switch 2307, which can comprise a software function, selects from one of the available transmission paths according to a weight distri-

bution table 2306. For example, if the weight for path X1 is 0.2, then every fifth packet will be transmitted on path X1. A similar regime holds true for the other paths as shown. Initially, each link's weight value can be set such that it is proportional to its bandwidth, which will be referred to as its "steady-state" value.

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

Link quality measurement function 2304 can be made to operate as part of a synchronizer function as described above. That is, if resynchronization occurs and the receiver detects that synchronization has been lost (e.g., resulting in the synchronization window W being advanced out of sequence), that fact can be used to drive link quality measurement function 2304. According to one embodiment, load balancing is performed using information garnered during the normal synchronization, augmented slightly to communicate link health from the receiver to the transmitter. The receiver maintains a count, MESS\_R(W), of the messages received in synchronization window W. When it receives a synchronization request (SYNC\_REQ) corresponding to the end of window W, the receiver includes counter MESS\_R in the resulting synchronization acknowledgement (SYNC\_ACK) sent back to the transmitter. This allows the transmitter to compare messages sent to messages received in order to assess the health of the link.

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

When the transmitter receives a SYNC\_ACK, the MESS\_R is compared with the number of messages transmitted in a window (MESS\_T). When the transmitter receives a SYNC\_ACK, the traffic probabilities will be examined and adjusted if necessary. MESS\_R is compared with the number of messages transmitted in a window (MESS\_T). There are two possibilities:

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

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

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

2. If MESS\_R for a link is greater than or equal to THRESH, the link will be deemed healthy. If the weight P for

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

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

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

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

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

Suppose that a first link L1 can sustain a transmission bandwidth of 100 Mb/s and has a window size of 32; link L2 can sustain 75 Mb/s and has a window size of 24; and link L3 can sustain 25 Mb/s and has a window size of 8. The combined links can thus sustain 200 Mb/s. The steady state traffic weights are 0.5 for link L1; 0.375 for link L2, and 0.125 for link L3. MIN=1 Mb/s, THRESH=0.8 MESS\_T for each link,  $\alpha=0.75$  and  $\beta=0.5$ . These traffic weights will remain stable until a link stops for synchronization or reports a number of packets received less than its THRESH. Consider the following sequence of events:

1. Link L1 receives a SYNC\_ACK containing a MESS\_R of 24, indicating that only 75% of the MESS\_T (32) messages transmitted in the last window were successfully received. Link L1 would be below THRESH (0.8). Consequently, link L1's traffic weight value would be reduced to 0.12825, while link L2's traffic weight value would be increased to 0.65812 and link L3's traffic weight value would be increased to 0.217938.
2. Link L2 and L3 remained healthy and link L1 stopped to synchronize. Then link L1's traffic weight value would be set to 0, link L2's traffic weight value would be set to 0.75, and link L3's traffic weight value would be set to 0.25.
3. Link L1 finally received a SYNC\_ACK containing a MESS\_R of 0 indicating that none of the MESS\_T (32) messages transmitted in the last window were successfully received. Link L1 would be below THRESH. Link L1's traffic weight value would be increased to 0.005, link L2's traffic weight value would be decreased to 0.74625, and link L3's traffic weight value would be decreased to 0.24875.
4. Link L1 received a SYNC\_ACK containing a MESS\_R of 32 indicating that 100% of the MESS\_T (32) messages transmitted in the last window were successfully received. Link L1 would be above THRESH. Link L1's traffic weight value would be increased to 0.2525, while link L2's traffic weight value would be decreased to 0.560625 and link L3's traffic weight value would be decreased to 0.186875.
5. Link L1 received a SYNC\_ACK containing a MESS\_R of 32 indicating that 100% of the MESS\_T (32) messages transmitted in the last window were successfully received. Link L1 would be above THRESH. Link L1's traffic weight value would be increased to 0.37625; link L2's traffic weight value would be decreased to 0.4678125, and link L3's traffic weight value would be decreased to 0.1559375.
6. Link L1 remains healthy and the traffic probabilities approach their steady state traffic probabilities.

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

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

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

This conventional scheme is shown in FIG. 25. A user's computer 2501 includes a client application 2504 (for example, a web browser) and an IP protocol stack 2505. When the user enters the name of a destination host, a request DNS REQ is made (through IP protocol stack 2505) to a DNS 2502 to look up the IP address associated with the name. The DNS returns the IP address DNS RESP to client application 2504, which is then able to use the IP address to communicate with the host 2503 through separate transactions such as PAGE REQ and PAGE RESP.

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

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

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

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

FIG. 26 shows a system employing various principles summarized above. A user's computer 2601 includes a conventional client (e.g., a web browser) 2605 and an IP protocol stack 2606 that preferably operates in accordance with an IP hopping function 2607 as outlined above. A modified DNS

server 2602 includes a conventional DNS server function 2609 and a DNS proxy 2610. A gatekeeper server 2603 is interposed between the modified DNS server and a secure target site 2704. An "unsecure" target site 2611 is also accessible via conventional IP protocols.

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

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

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

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

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

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

level can also be determined by transmitting a request message back to the user's computer requiring that it prove that it has sufficient privileges.

If the user is not authorized to access the secure site, then a "host unknown" message is returned (step 2705). If the user has sufficient security privileges, then in step 2706 a secure VPN is established between the user's computer and the secure target site. As described above, this is preferably done by allocating a hopping regime that will be carried out between the user's computer and the secure target site, and is preferably performed transparently to the user (i.e., the user need not be involved in creating the secure link). As described in various embodiments of this application, any of various fields can be "hopped" (e.g., IP source/destination addresses; a field in the header; etc.) in order to communicate securely.

Some or all of the security functions can be embedded in gatekeeper 2603, such that it handles all requests to connect to secure sites. In this embodiment, DNS proxy 2610 communicates with gatekeeper 2603 to determine (preferably over a secure administrative VPN) whether the user has access to a particular web site. Various scenarios for implementing these features are described by way of example below:

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

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

Scenario #3: Client has permission to connect using a normal non-VPN link, and the gatekeeper does not have a rule to set up a VPN for the client to the target site. In this scenario, the client's DNS request is received by DNS proxy server 2610, which would check its rules and determine that no VPN is needed. Gatekeeper 2603 would then inform the DNS proxy server to forward the request to conventional DNS server 2609, which would resolve the request and return the result to the DNS proxy server and then back to the client.

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

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

One feature of the basic architecture is the ability to prevent so-called "denial of service" attacks that can occur if a computer hacker floods a known Internet node with packets, thus preventing the node from communicating with other nodes.

Because IP addresses or other fields are "hopped" and packets arriving with invalid addresses are quickly discarded, Internet nodes are protected against flooding targeted at a single IP address.

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

Suppose that a computer hacker is able to transmit a large quantity of dummy packets addressed to first host computer 2801 across high bandwidth link HIGH BW. Normally, host computer 2801 would be able to quickly reject the packets since they would not fall within the acceptance window permitted by the IP address hopping scheme. However, because the packets must travel across low bandwidth link LOW BW, the packets overwhelm the lower bandwidth link before they are received by host computer 2801. Consequently, the link to host computer 2801 is effectively flooded before the packets can be discarded.

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

In one embodiment, the ISP distinguishes between VPN and non-VPN packets using the protocol of the packet. In the case of IPSEC [rfc 2401], the packets have IP protocols 420 and 421. In the case of the TARP VPN, the packets will have an IP protocol that is not yet defined. The ISP's link guard, 2805, maintains a table of valid VPNs which it uses to validate whether VPN packets are cryptographically valid. According to one embodiment, packets that do not fall within any hop windows used by nodes on the low-bandwidth link are rejected, or are sent with a lower quality of service. One approach for doing this is to provide a copy of the IP hopping tables used by the low-bandwidth nodes to the high-bandwidth node, such that both the high-bandwidth and low-bandwidth nodes track hopped packets (e.g., the high-bandwidth node moves its hopping window as valid packets are received). In such a scenario, the high-bandwidth node discards packets that do not fall within the hopping window before they are transmitted over the low-bandwidth link. Thus, for example, ISP 2903 maintains a copy 2910 of the receive table used by host computer 2901. Incoming packets that do not fall within this receive table are discarded. According to a different embodiment, link guard 2805 validates each VPN packet using a keyed hashed message authentication code (HMAC) [rfc 2104].

According to another embodiment, separate VPNs (using, for example, hopblocks) can be established for communicat-



ing between the low-bandwidth node and the high-bandwidth node (i.e., packets arriving at the high-bandwidth node are converted into different packets before being transmitted to the low-bandwidth node).

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

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

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

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

#### D. Traffic Limiter

In a system in which multiple nodes are communicating using “hopping” technology, a treasonous insider could inter-

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

A simple modification to the checkpoint synchronizer will serve to protect a receiver from accidental or deliberate overload from an internally treasonous client. This modification is based on the observation that a receiver will not update its tables until a SYNC\_REQ is received on hopped address CKPT\_N. It is a simple matter of deferring the generation of a new CKPT\_N until an appropriate interval after previous checkpoints.

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

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

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

1. The receiver rate should be slightly higher than the permitted rate in order to allow for statistical fluctuations in traffic arrival times and non-uniform load balancing.
2. Since a transmitter will rightfully continue to transmit for a period after a SYNC\_REQ is transmitted, the algorithm above can artificially reduce the transmitter’s bandwidth. If events prevent a compliant transmitter from synchronizing for a period (e.g. the network dropping a SYNC\_REQ or a SYNC\_ACK) a SYNC\_REQ will be accepted later than expected. After this, the transmitter will transmit fewer than expected messages before encountering the next checkpoint. The new checkpoint will not have been activated and the transmitter will have to retransmit the SYNC\_REQ. This will appear to the receiver as if the transmitter is not compliant. Therefore, the next checkpoint will be accepted late from the transmitter’s perspective. This has the effect of reducing the transmitter’s allowed packet rate until the

45

transmitter transmits at a packet rate below the agreed upon rate for a period of time.

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

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

As described above, receiving computer 3000 maintains a receive table 3002 including a window  $W$  that defines valid IP address pairs that will be accepted when appearing in incoming data packets. Transmitting computer 3001 maintains a transmit table 3003 from which the next IP address pairs will be selected when transmitting a packet to receiving computer 3000. (For the sake of illustration, window  $W$  is also illustrated with reference to transmit table 3003). As transmitting computer moves through its table, it will eventually generate a SYNC\_REQ message as illustrated in function 3010. This is a request to receiver 3000 to synchronize the receive table 3002, from which transmitter 3001 expects a response in the form of a CKPT\_N (included as part of a SYNC\_ACK message). If transmitting computer 3001 transmits more messages than its allotment, it will prematurely generate the SYNC\_REQ message. (If it has been altered to remove the SYNC\_REQ message generation altogether, it will fall out of synchronization since receiver 3000 will quickly reject packets that fall outside of window  $W$ , and the extra packets generated by transmitter 3001 will be discarded).

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

Otherwise, if the rate has not been exceeded, then in step 3109 the next CKPT\_N value is calculated and inserted into the receiver's hopping table prior to the next SYNC\_REQ from the transmitter 3101. Transmitter 3101 then processes the SYNC\_REQ in the normal manner.

#### E. Signaling Synchronizer

In a system in which a large number of users communicate with a central node using secure hopping technology, a large

46

amount of memory must be set aside for hopping tables and their supporting data structures. For example, if one million subscribers to a web site occasionally communicate with the web site, the site must maintain one million hopping tables, thus using up valuable computer resources, even though only a small percentage of the users may actually be using the system at any one time. A desirable solution would be a system that permits a certain maximum number of simultaneous links to be maintained, but which would "recognize" millions of registered users at any one time. In other words, out of a population of a million registered users, a few thousand at a time could simultaneously communicate with a central server, without requiring that the server maintain one million hopping tables of appreciable size.

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

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

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

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

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

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

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

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

The meaning and behaviors of CKPT\_N, CKPT\_O and CKPT\_R remain the same from the previous description, except that CKPT\_N can receive a combined data and SYNC\_REQ message or a SYNC\_REQ message without the data.

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

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

1. The client sends the message marked as a data message on the inner header using the transmitter's CKPT\_N address. It turns the transmitter off and starts a timer T1 noting CKPT\_O. Messages can be one of three types: DATA, SYNC\_REQ and SYNC\_ACK. In the normal algorithm, some potential problems can be prevented by identifying each message type as part of the encrypted inner header field. In this algorithm, it is important to distinguish a data packet and a SYNC\_REQ in the signaling synchronizer since the data and the SYNC\_REQ come in on the same address.
2. When the server receives a data message on its CKPT\_N, it verifies the message and passes it up the stack. The message can be verified by checking message type and other information (i.e., user credentials) contained in the inner header. It replaces its CKPT\_O with CKPT\_N and generates the next CKPT\_N. It updates its transmitter side CKPT\_R to correspond to the client's receiver side CKPT\_R and transmits a SYNC\_ACK containing CKPT\_O in its payload.

3. When the client side receiver receives a SYNC\_ACK on its CKPT\_R with a payload matching its transmitter side CKPT\_O and the transmitter is off, the transmitter is turned on and the receiver side CKPT\_R is updated. If the SYNC\_ACK's payload does not match the transmitter side CKPT\_O or the transmitter is on, the SYNC\_ACK is simply discarded.
4. T1 expires: If the transmitter is off and the client's transmitter side CKPT\_O matches the CKPT\_O associated with the timer, it starts timer T1 noting CKPT\_O again, and a SYNC\_REQ is sent using the transmitter's CKPT\_O address. Otherwise, no action is taken.
5. When the server receives a SYNC\_REQ on its CKPT\_N, it replaces its CKPT\_O with CKPT\_N and generates the next CKPT\_N. It updates its transmitter side CKPT\_R to correspond to the client's receiver side CKPT\_R and transmits a SYNC\_ACK containing CKPT\_O in its payload.
6. When the server receives a SYNC\_REQ on its CKPT\_O, it updates its transmitter side CKPT\_R to correspond to the client's receiver side CKPT\_R and transmits a SYNC\_ACK containing CKPT\_O in its payload.

FIG. **32** shows message flows to highlight the protocol.

Reading from top to bottom, the client sends data to the server using its transmitter side CKPT\_N. The client side transmitter is turned off and a retry timer is turned off. The transmitter will not transmit messages as long as the transmitter is turned off. The client side transmitter then loads CKPT\_N into CKPT\_O and updates CKPT\_N. This message is successfully received and a passed up the stack. It also synchronizes the receiver i.e., the server loads CKPT\_N into CKPT\_O and generates a new CKPT\_N, it generates a new CKPT\_R in the server side transmitter and transmits a SYNC\_ACK containing the server side receiver's CKPT\_O the server. The SYNC\_ACK is successfully received at the client. The client side receiver's CKPT\_R is updated, the transmitter is turned on and the retry timer is killed. The client side transmitter is ready to transmit a new data message.

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

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

The above-described procedures allow a client to be authenticated at signaling server **3201** while maintaining the ability of signaling server **3201** to quickly reject invalid packets, such as might be generated by hacker computer **3205**. In various embodiments, the signaling synchronizer is really a

derivative of the synchronizer. It provides the same protection as the hopping protocol, and it does so for a large number of low bandwidth connections.

#### F. One-Click Secure on-Line Communications and Secure Domain Name Service

The present invention provides a technique for establishing a secure communication link between a first computer and a second computer over a computer network. Preferably, a user enables a secure communication link using a single click of a mouse, or a corresponding minimal input from another input device, such as a keystroke entered on a keyboard or a click entered through a trackball. Alternatively, the secure link is automatically established as a default setting at boot-up of the computer (i.e., no click). FIG. 33 shows a system block diagram 3300 of a computer network in which the one-click secure communication method of the present invention is suitable. In FIG. 33, a computer terminal or client computer 3301, such as a personal computer (PC), is connected to a computer network 3302, such as the Internet, through an ISP 3303. Alternatively, computer 3301 can be connected to computer network 3302 through an edge router. Computer 3301 includes an input device, such as a keyboard and/or mouse, and a display device, such as a monitor. Computer 3301 can communicate conventionally with another computer 3304 connected to computer network 3302 over a communication link 3305 using a browser 3306 that is installed and operates on computer 3301 in a well-known manner.

Computer 3304 can be, for example, a server computer that is used for conducting e-commerce. In the situation when computer network 3302 is the Internet, computer 3304 typically will have a standard top-level domain name such as .com, .net, .org, .edu, .mil or .gov.

FIG. 34 shows a flow diagram 3400 for installing and establishing a "one-click" secure communication link over a computer network according to the present invention. At step 3401, computer 3301 is connected to server computer 3304 over a non-VPN communication link 3305. Web browser 3306 displays a web page associated with server 3304 in a well-known manner. According to one variation of the invention, the display of computer 3301 contains a hyperlink, or an icon representing a hyperlink, for selecting a virtual private network (VPN) communication link ("go secure" hyperlink) through computer network 3302 between terminal 3301 and server 3304. Preferably, the "go secure" hyperlink is displayed as part of the web page downloaded from server computer 3304, thereby indicating that the entity providing server 3304 also provides VPN capability.

By displaying the "go secure" hyperlink, a user at computer 3301 is informed that the current communication link between computer 3301 and server computer 3304 is a non-secure, non-VPN communication link. At step 3402, it is determined whether a user of computer 3301 has selected the "go secure" hyperlink. If not, processing resumes using a non-secure (conventional) communication method (not shown). If, at step 3402, it is determined that the user has selected the "go secure" hyperlink, flow continues to step 3403 where an object associated with the hyperlink determines whether a VPN communication software module has already been installed on computer 3301. Alternatively, a user can enter a command into computer 3301 to "go secure."

If, at step 3403, the object determines that the software module has been installed, flow continues to step 3407. If, at step 3403, the object determines that the software module has not been installed, flow continues to step 3404 where a non-VPN communication link 3307 is launched between com-

puter 3301 and a website 3308 over computer network 3302 in a well-known manner. Website 3308 is accessible by all computer terminals connected to computer network 3302 through a non-VPN communication link. Once connected to website 3308, a software module for establishing a secure communication link over computer network 3302 can be downloaded and installed. Flow continues to step 3405 where, after computer 3301 connects to website 3308, the software module for establishing a communication link is downloaded and installed in a well-known manner on computer terminal 3301 as software module 3309. At step 3405, a user can optionally select parameters for the software module, such as enabling a secure communication link mode of communication for all communication links over computer network 3302. At step 3406, the communication link between computer 3301 and website 3308 is then terminated in a well-known manner.

By clicking on the "go secure" hyperlink, a user at computer 3301 has enabled a secure communication mode of communication between computer 3301 and server computer 3304. According to one variation of the invention, the user is not required to do anything more than merely click the "go secure" hyperlink. The user does not need to enter any user identification information, passwords or encryption keys for establishing a secure communication link. All procedures required for establishing a secure communication link between computer 3301 and server computer 3304 are performed transparently to a user at computer 3301.

At step 3407, a secure VPN communications mode of operation has been enabled and software module 3309 begins to establish a VPN communication link. In one embodiment, software module 3309 automatically replaces the top-level domain name for server 3304 within browser 3406 with a secure top-level domain name for server computer 3304. For example, if the top-level domain name for server 3304 is .com, software module 3309 replaces the .com top-level domain name with a .scom top-level domain name, where the "s" stands for secure. Alternatively, software module 3409 can replace the top-level domain name of server 3304 with any other non-standard top-level domain name.

Because the secure top-level domain name is a non-standard domain name, a query to a standard domain name service (DNS) will return a message indicating that the universal resource locator (URL) is unknown. According to the invention, software module 3409 contains the URL for querying a secure domain name service (SDNS) for obtaining the URL for a secure top-level domain name. In this regard, software module 3309 accesses a secure portal 3310 that interfaces a secure network 3311 to computer network 3302. Secure network 3311 includes an internal router 3312, a secure domain name service (SDNS) 3313, a VPN gatekeeper 3314 and a secure proxy 3315. The secure network can include other network services, such as e-mail 3316, a plurality of chatrooms (of which only one chatroom 3317 is shown), and a standard domain name service (STD DNS) 3318. Of course, secure network 3311 can include other resources and services that are not shown in FIG. 33.

When software module 3309 replaces the standard top-level domain name for server 3304 with the secure top-level domain name, software module 3309 sends a query to SDNS 3313 at step 3408 through secure portal 3310 preferably using an administrative VPN communication link 3319. In this configuration, secure portal 3310 can only be accessed using a VPN communication link. Preferably, such a VPN communication link can be based on a technique of inserting a source and destination IP address pair into each data packet that is selected according to a pseudo-random sequence; an IP

51

address hopping regime that pseudorandomly changes IP addresses in packets transmitted between a client computer and a secure target computer; periodically changing at least one field in a series of data packets according to a known sequence; an Internet Protocol (IP) address in a header of each data packet that is compared to a table of valid IP addresses maintained in a table in the second computer; and/or a comparison of the IP address in the header of each data packet to a moving window of valid IP addresses, and rejecting data packets having IP addresses that do not fall within the moving window. Other types of VPNs can alternatively be used. Secure portal 3310 authenticates the query from software module 3309 based on the particular information hopping technique used for VPN communication link 3319.

SDNS 3313 contains a cross-reference database of secure domain names and corresponding secure network addresses. That is, for each secure domain name, SDNS 3313 stores a computer network address corresponding to the secure domain name. An entity can register a secure domain name in SDNS 3313 so that a user who desires a secure communication link to the website of the entity can automatically obtain the secure computer network address for the secure website. Moreover, an entity can register several secure domain names, with each respective secure domain name representing a different priority level of access in a hierarchy of access levels to a secure website. For example, a securities trading website can provide users secure access so that a denial of service attack on the website will be ineffectual with respect to users subscribing to the secure website service. Different levels of subscription can be arranged based on, for example, an escalating fee, so that a user can select a desired level of guarantee for connecting to the secure securities trading website. When a user queries SDNS 3313 for the secure computer network address for the securities trading website, SDNS 3313 determines the particular secure computer network address based on the user's identity and the user's subscription level.

At step 3409, SDNS 3313 accesses VPN gatekeeper 3314 for establishing a VPN communication link between software module 3309 and secure server 3320. Server 3320 can only be accessed through a VPN communication link. VPN gatekeeper 3314 provisions computer 3301 and secure web server computer 3320, or a secure edge router for server computer 3320, thereby creating the VPN. Secure server computer 3320 can be a separate server computer from server computer 3304, or can be the same server computer having both non-VPN and VPN communication link capability, such as shown by server computer 3322. Returning to FIG. 34, in step 3410, SDNS 3313 returns a secure URL to software module 3309 for the .com server address for a secure server 3320 corresponding to server 3304.

Alternatively, SDNS 3313 can be accessed through secure portal 3310 "in the clear", that is, without using an administrative VPN communication link. In this situation, secure portal 3310 preferably authenticates the query using any well-known technique, such as a cryptographic technique, before allowing the query to proceed to SDNS 3313. Because the initial communication link in this situation is not a VPN communication link, the reply to the query can be "in the clear." The querying computer can use the clear reply for establishing a VPN link to the desired domain name. Alternatively, the query to SDNS 3313 can be in the clear, and SDNS 3313 and gatekeeper 3314 can operate to establish a VPN communication link to the querying computer for sending the reply.

At step 3411, software module 3309 accesses secure server 3320 through VPN communication link 3321 based on the

52

VPN resources allocated by VPN gatekeeper 3314. At step 3412, web browser 3306 displays a secure icon indicating that the current communication link to server 3320 is a secure VPN communication link. Further communication between computers 3301 and 3320 occurs via the VPN, e.g., using a "hopping" regime as discussed above. When VPN link 3321 is terminated at step 3413, flow continues to step 3414 where software module 3309 automatically replaces the secure top-level domain name with the corresponding non-secure top-level domain name for server 3304. Browser 3306 accesses a standard DNS 3325 for obtaining the non-secure URL for server 3304. Browser 3306 then connects to server 3304 in a well-known manner. At step 3415, browser 3306 displays the "go secure" hyperlink or icon for selecting a VPN communication link between terminal 3301 and server 3304. By again displaying the "go secure" hyperlink, a user is informed that the current communication link is a non-secure, non-VPN communication link.

When software module 3309 is being installed or when the user is off-line, the user can optionally specify that all communication links established over computer network 3302 are secure communication links. Thus, anytime that a communication link is established, the link is a VPN link. Consequently, software module 3309 transparently accesses SDNS 3313 for obtaining the URL for a selected secure website. In other words, in one embodiment, the user need not "click" on the secure option each time secure communication is to be effected.

Additionally, a user at computer 3301 can optionally select a secure communication link through proxy computer 3315. Accordingly, computer 3301 can establish a VPN communication link 3323 with secure server computer 3320 through proxy computer 3315. Alternatively, computer 3301 can establish a non-VPN communication link 3324 to a non-secure website, such as non-secure server computer 3304.

FIG. 35 shows a flow diagram 3500 for registering a secure domain name according to the present invention. At step 3501, a requester accesses website 3308 and logs into a secure domain name registry service that is available through website 3308. At step 3502, the requestor completes an online registration form for registering a secure domain name having a top-level domain name, such as .com, .net, .org, .edu, .mil or .gov. Of course, other secure top-level domain names can also be used. Preferably, the requestor must have previously registered a non-secure domain name corresponding to the equivalent secure domain name that is being requested. For example, a requester attempting to register secure domain name "website.scom" must have previously registered the corresponding non-secure domain name "website.com".

At step 3503, the secure domain name registry service at website 3308 queries a non-secure domain name server database, such as standard DNS 3322, using, for example, a who is query, for determining ownership information relating to the non-secure domain name corresponding to the requested secure domain name. At step 3504, the secure domain name registry service at website 3308 receives a reply from standard DNS 3322 and at step 3505 determines whether there is conflicting ownership information for the corresponding non-secure domain name. If there is no conflicting ownership information, flow continues to step 3507, otherwise flow continues to step 3506 where the requestor is informed of the conflicting ownership information. Flow returns to step 3502.

When there is no conflicting ownership information at step 3505, the secure domain name registry service (website 3308) informs the requestor that there is no conflicting ownership information and prompts the requestor to verify the information entered into the online form and select an approved form

of payment. After confirmation of the entered information and appropriate payment information, flow continues to step 3508 where the newly registered secure domain name sent to SDNS 3313 over communication link 3326.

If, at step 3505, the requested secure domain name does not have a corresponding equivalent non-secure domain name, the present invention informs the requestor of the situation and prompts the requestor for acquiring the corresponding equivalent non-secure domain name for an increased fee. By accepting the offer, the present invention automatically registers the corresponding equivalent non-secure domain name with standard DNS 3325 in a well-known manner. Flow then continues to step 3508.

#### G. Tunneling Secure Address Hopping Protocol Through Existing Protocol Using Web Proxy

The present invention also provides a technique for implementing the field hopping schemes described above in an application program on the client side of a firewall between two computer networks, and in the network stack on the server side of the firewall. The present invention uses a new secure connectionless protocol that provides good denial of service rejection capabilities by layering the new protocol on top of an existing IP protocol, such as the ICMP, UDP or TCP protocols. Thus, this aspect of the present invention does not require changes in the Internet infrastructure.

According to the invention, communications are protected by a client-side proxy application program that accepts unencrypted, unprotected communication packets from a local browser application. The client-side proxy application program tunnels the unencrypted, unprotected communication packets through a new protocol, thereby protecting the communications from a denial of service at the server side. Of course, the unencrypted, unprotected communication packets can be encrypted prior to tunneling.

The client-side proxy application program is not an operating system extension and does not involve any modifications to the operating system network stack and drivers. Consequently, the client is easier to install, remove and support in comparison to a VPN. Moreover, the client-side proxy application can be allowed through a corporate firewall using a much smaller "hole" in the firewall and is less of a security risk in comparison to allowing a protocol layer VPN through a corporate firewall.

The server-side implementation of the present invention authenticates valid field-hopped packets as valid or invalid very early in the server packet processing, similar to a standard virtual private network, for greatly minimizing the impact of a denial of service attempt in comparison to normal TCP/IP and HTTP communications, thereby protecting the server from invalid communications.

FIG. 36 shows a system block diagram of a computer network 3600 in which a virtual private connection according to the present invention can be configured to more easily traverse a firewall between two computer networks. FIG. 37 shows a flow diagram 3700 for establishing a virtual private connection that is encapsulated using an existing network protocol.

In FIG. 36 a local area network (LAN) 3601 is connected to another computer network 3602, such as the Internet, through a firewall arrangement 3603. Firewall arrangement operates in a well-known manner to interface LAN 3601 to computer network 3602 and to protect LAN 3601 from attacks initiated outside of LAN 3601.

A client computer 3604 is connected to LAN 3601 in a well-known manner. Client computer 3604 includes an oper-

ating system 3605 and a web browser 3606. Operating system 3605 provides kernel mode functions for operating client computer 3604. Browser 3606 is an application program for accessing computer network resources connected to LAN 3601 and computer network 3602 in a well-known manner. According to the present invention, a proxy application 3607 is also stored on client computer 3604 and operates at an application layer in conjunction with browser 3606. Proxy application 3607 operates at the application layer within client computer 3604 and when enabled, modifies unprotected, unencrypted message packets generated by browser 3606 by inserting data into the message packets that are used for forming a virtual private connection between client computer 3604 and a server computer connected to LAN 3601 or computer network 3602. According to the invention, a virtual private connection does not provide the same level of security to the client computer as a virtual private network. A virtual private connection can be conveniently authenticated so that, for example, a denial of service attack can be rapidly rejected, thereby providing different levels of service that can be subscribed to by a user.

Proxy application 3607 is conveniently installed and uninstalled by a user because proxy application 3607 operates at the application layer within client computer 3604. On installation, proxy application 3607 preferably configures browser 3606 to use proxy application for all web communications. That is, the payload portion of all message packets is modified with the data for forming a virtual private connection between client computer 3604 and a server computer. Preferably, the data for forming the virtual private connection contains field-hopping data, such as described above in connection with VPNs. Also, the modified message packets preferably conform to the UDP protocol. Alternatively, the modified message packets can conform to the TCP/IP protocol or the ICMP protocol. Alternatively, proxy application 3606 can be selected and enabled through, for example, an option provided by browser 3606. Additionally, proxy application 3607 can be enabled so that only the payload portion of specially designated message packets is modified with the data for forming a virtual private connection between client computer 3604 and a designated host computer. Specially designated message packets can be, for example, selected predetermined domain names.

Referring to FIG. 37, at step 3701, unprotected and unencrypted message packets are generated by browser 3606. At step 3702, proxy application 3607 modifies the payload portion of all message packets by tunneling the data for forming a virtual private connection between client computer 3604 and a destination server computer into the payload portion. At step, 3703, the modified message packets are sent from client computer 3604 to, for example, website (server computer) 3608 over computer network 3602.

Website 3608 includes a VPN guard portion 3609, a server proxy portion 3610 and a web server portion 3611. VPN guard portion 3609 is embedded within the kernel layer of the operating system of website 3608 so that large bandwidth attacks on website 3608 are rapidly rejected. When client computer 3604 initiates an authenticated connection to website 3608, VPN guard portion 3609 is keyed with the hopping sequence contained in the message packets from client computer 3604, thereby performing a strong authentication of the client packet streams entering website 3608 at step 3704. VPN guard portion 3609 can be configured for providing different levels of authentication and, hence, quality of service, depending upon a subscribed level of service. That is, VPN guard portion 3609 can be configured to let all message packets through until a denial of service attack is detected, in

55

which case VPN guard portion 3609 would allow only client packet streams conforming to a keyed hopping sequence, such as that of the present invention.

Server proxy portion 3610 also operates at the kernel layer within website 3608 and catches incoming message packets from client computer 3604 at the VPN level. At step 3705, server proxy portion 3610 authenticates the message packets at the kernel level within host computer 3604 using the destination IP address, UDP ports and discriminator fields. The authenticated message packets are then forwarded to the authenticated message packets to web server portion 3611 as normal TCP web transactions.

At step 3705, web server portion 3611 responds to message packets received from client computer 3604 in accordance with the particular nature of the message packets by generating reply message packets. For example, when a client computer requests a webpage, web server portion 3611 generates message packets corresponding to the requested webpage. At step 3706, the reply message packets pass through server proxy portion 3610, which inserts data into the payload portion of the message packets that are used for forming the virtual private connection between host computer 3608 and client computer 3604 over computer network 3602. Preferably, the data for forming the virtual private connection is contains field-hopping data, such as described above in connection with VPNs. Server proxy portion 3610 operates at the kernel layer within host computer 3608 to insert the virtual private connection data into the payload portion of the reply message packets. Preferably, the modified message packets sent by host computer 3608 to client computer 3604 conform to the UDP protocol. Alternatively, the modified message packets can conform to the TCP/IP protocol or the ICMP protocol.

At step 3707, the modified packets are sent from host computer 3608 over computer network 3602 and pass through firewall 3603. Once through firewall 3603, the modified packets are directed to client computer 3604 over LAN 3601 and are received at step 3708 by proxy application 3607 at the application layer within client computer 3604. Proxy application 3607 operates to rapidly evaluate the modified message packets for determining whether the received packets should be accepted or dropped. If the virtual private connection data inserted into the received information packets conforms to expected virtual private connection data, then the received packets are accepted. Otherwise, the received packets are dropped.

While the present invention has been described in connection with the illustrated embodiments, it will be appreciated and understood that modifications may be made without departing from the true spirit and scope of the invention.

What is claimed is:

1. A client device comprising: (a) memory configured and arranged to facilitate a connection of the client device with a target device over a secure communication link created based on (i) interception of a request, generated by the client device, to look up an internet protocol (IP) address of the target device based on a domain name associated with the target device, and (ii) a determination as a result of the request that the target device is a device with which a secure communication link can be established; (b) an application program configured and arranged so as to allow participation in audio/video communications with the target device over the secure communication link once the secure communication link is established; and (c) a signal processing configuration arranged to execute the application program.

2. The client device of claim 1, wherein the client device is a computer.

56

3. The client device of claim 1, wherein the client device is a phone.

4. The client device of claim 3, wherein the establishment of the secure communication link is based on a determination being made by a server that the target device is a device with which a secure communication link can be established.

5. The client device of claim 1, wherein the interception of the request consists of receiving the request to determine whether the target device is available for the secure communications service.

6. The client device of claim 3, wherein the secure communication link is a virtual private network link.

7. The client device of claim 3, wherein the domain name is a secure domain name.

8. The client device of claim 3, wherein the establishment of the secure communication link is based on the determination being made by a proxy module that the target device is a device with which a secure communication link can be established when the domain name corresponds to a target device identified in a DNS lookup table.

9. The client device of claim 1, wherein the interception of the request occurs within another device that is separate from the client device.

10. The client device of claim 3, wherein the secure communication link is a modulated transmission link.

11. The client device of claim 3, where the secure communication link supports at least one of the following: FTM, TDM and CDMA.

12. The client device of claim 3, wherein the secure communication link supports TCP/IP protocols.

13. The client device of claim 3, wherein the target device is a server.

14. The client device of claim 3, wherein the target device is a computer.

15. The client device of claim 3, wherein the target device is a phone.

16. A method executed by a client device for communicating with a target device, the method comprising:

(a) facilitating a connection with the target device over a secure communication link created based on (i) interception of a request, generated by the client device, to look up an internet protocol (IP) address of the target device based on a domain name associated with the target device, and (ii) a determination as a result of the request that the target device is a device with which a secure communication link can be established; and

(b) Allowing participation in audio/video communications with the target device over the secure communication link once the secure communication link is established.

17. The method of claim 16, wherein the client device is a computer.

18. The method of claim 16, wherein the client device is a phone.

19. The method of claim 16, wherein the interception of the request consists of receiving the request to determine whether the target device is available for the secure communications service.

20. The method of claim 16, wherein the establishment of the secure communication link is based on a determination being made by a server that the target device is a device with which a secure communication link can be established.

21. The method of claim 16, wherein the secure communication link is a virtual private network link.

22. The method of claim 16, wherein the domain name is a secure domain name.

23. The method of claim 16, wherein the interception of the request occurs within another device that is separate from the client device.

24. The method of claim 16, wherein the secure communication link is an unmodulated transmission link. 5

25. The method of claim 16, wherein the secure communication link is a modulated transmission link.

26. The method of claim 16, where the secure communication link supports at least one of the following: FTM, TDM and CDMA. 10

27. The method of claim 16, wherein the secure communication link supports TCP/IP protocols.

28. The method of claim 16, wherein the target device is a server.

29. The method of claim 16, wherein the target device is a computer. 15

30. The method of claim 16, wherein the target device is a phone.

\* \* \* \* \*