

EXHIBIT

1002

FILE HISTORY

US 6,775,235

PATENT: 6,775,235

INVENTORS: Datta, Sanchaita
Bhaskar, Ragula

TITLE: Tools and techniques for directing packets
over disparate networks

APPLICATION
NO: US2003361837A

FILED: 07 FEB 2003

ISSUED: 10 AUG 2004

COMPILED: 15 MAY 2015

02/07/03

AUG 10 2004

PATENT NUMBER and
ISSUE DATE



6775235

U.S. UTILITY Patent Application

APPL. NUMBER 10361837	FILING DATE 02/07/2003	CLASS 370	SUBCLASS 401	GAU 2663	EXAMINER <i>Nguyen</i>
APPLICANTS: Datta Sanchaita; Bhaskar Ragula;					
<p><i>OK MM</i></p> <p>CONTINUING DATA VERIFIED: This application is a CIP of 10/034,197 12/28/2001 which claims benefit of 60/259,269 12/29/2000 This application 10/361,837 claims benefit of 60/355,509 02/08/2002</p>					
<p><i>Non mm</i></p> <p>** FOREIGN APPLICATIONS VERIFIED:</p>					
PG-PUB DO NOT PUBLISH <input type="checkbox"/>		RESCIND <input type="checkbox"/>			
Foreign priority claimed <input type="checkbox"/> yes <input checked="" type="checkbox"/> no			ATTORNEY DOCKET NO.		
35 USC 119 conditions met <input type="checkbox"/> yes <input checked="" type="checkbox"/> no			2003.2.11A		
Verified and Acknowledged Examiners's initials <i>MM</i>					
TITLE : Tools and techniques for directing packets over disparate networks					
<small>U.S. DEPT. OF COMM. PAT. & TM-PTO-436L (R.v. 12-94)</small>					

7/12/04 Formal Drawings (6 sheets) *2/7/03*

NOTICE OF ALLOWANCE MAILED <i>5/26/04</i>		Assistant Examiner <i>Melvin Marcelo</i>	CLAIMS ALLOWED	
ISSUE FEE			Total Claims <i>24</i>	Print Claim for O.G <i>1</i>
Amount Due <i>\$665.</i>	Date Paid <i>6/17/04 JM</i>	MELVIN MARCELO PRIMARY EXAMINER Primary Examiner	DRAWING	
<input type="checkbox"/> TERMINAL			Sheets Drwg. <i>6</i>	Figs. Drwg. <i>11</i>
ISSUE FEE IN FILE		PREPARED FOR ISSUE	<i>Jim Dabbs 5/29/04</i> Application Examiner	
<p>DISCLAIMER</p> <p>WARNING: The information disclosed herein may be restricted. Unauthorized disclosure may be prohibited by the United States Code Title 35, Sections 122, 181 and 368, Possession outside the U.S. Patent & Trademark Office is restricted to authorized employees and contractors only.</p>				

FILED WITH: DISK (CRF) CD-ROM
(Attached in pocket on right inside flap)

6,775,235

**TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER
DISPARATE NETWORKS**

Transaction History

Date	Transaction Description
02-07-2003	Workflow - Drawings Finished
02-07-2003	Workflow - Drawings Matched with File at Contractor
02-07-2003	Initial Exam Team nn
02-26-2003	IFW Scan & PACR Auto Security Review
04-02-2003	Notice Mailed--Application Incomplete--Filing Date Assigned
04-15-2003	Payment of additional filing fee/Preexam
05-01-2003	Application Is Now Complete
05-02-2003	Application Dispatched from OIPE
09-26-2003	Case Docketed to Examiner in GAU
12-10-2003	Petition Entered
01-26-2004	Mail-Record Petition Decision of Granted to Make Special
01-28-2004	Case Docketed to Examiner in GAU
02-23-2004	Non-Final Rejection
02-25-2004	Mail Non-Final Rejection
04-05-2004	Information Disclosure Statement (IDS) Filed
04-05-2004	Information Disclosure Statement (IDS) Filed
05-18-2004	Response after Non-Final Action
05-18-2004	Workflow incoming amendment IFW
05-21-2004	Date Forwarded to Examiner
05-26-2004	Mail Notice of Allowance
05-26-2004	Notice of Allowance Data Verification Completed
06-01-2004	Workflow - File Sent to Contractor
06-01-2004	Receipt into Pubs
06-02-2004	Receipt into Pubs
06-14-2004	Issue Fee Payment Verified
06-14-2004	Issue Fee Payment Received
06-29-2004	Receipt into Pubs
07-09-2004	Correspondence Address Change
07-09-2004	Change in Power of Attorney (May Include Associate POA)
07-12-2004	Dispatch to FDC
07-12-2004	Application Is Considered Ready for Issue
07-14-2004	Receipt into Pubs
07-22-2004	Issue Notification Mailed
08-10-2004	Recordation of Patent Grant Mailed
08-10-2004	Patent Issue Date Used in PTA Calculation
11-09-2005	Correspondence Address Change

10361837

10/361837
02/07/03

INITIALS *JK*

CONTENTS

	Date Received (incl. C. of M.) or Date Mailed	Date Received (incl. C. of M.) or Date Mailed
1. Application <u>6</u> papers.		31.
2. <u>DD</u>	<u>2/13/03</u>	32.
3. <u>the new Fees</u>	<u>4/1/03</u>	33.
4. <u>FEE</u>	<u>04-15-03</u>	34.
5. <u>Patition-Special</u> <u>Accpt</u> <u>Exam</u>	<u>12-10-03</u>	35.
6. <u>Decision-Granted</u>	<u>1-26-04</u>	36.
2/23/04 <u>Key (3)</u>	<u>2/25/04</u>	37.
7. <u>F.D.S.</u>	<u>4.5.04</u>	38.
8. <u>mdtA</u>	<u>5/18/04</u>	39.
5/21/04 <u>Notice of Allowance</u>	<u>5/26/04</u>	40.
7-8-04 <u>Revoc / P.O.A.</u>	<u>6-7-04</u>	41.
12 <u>Notice of Revoc / Accept</u>	<u>7-8-04</u>	42.
13.		43.
14.		44.
15.		45.
16.		46.
17.		47.
18.		48.
19.		49.
20.		50.
21.		51.
22.		52.
23.		53.
24.		54.
25.		55.
26.		56.
27.		57.
28.		58.
29.		59.
30.		60.

ISSUE SLIP STAPLE AREA (for additional cross-references)

ORIGINAL		CROSS REFERENCE(S)					
CLASS	SUBCLASS	CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)				
370	238	370	252	352			
INTERNATIONAL CLASSIFICATION							
H	04L	12	164				
		1					
		1					
		1					
		1					

^ Continued on Issue Slip Inside File Jacket

INDEX OF CLAIMS

✓ Rejected - (Through numeral) ... Canceled N Non-elected A Appeal
 - Allowed + Restricted I Interference O Objected

Claim	Date	Claim	Date	Claim	Date
Final Original		Final Original		Final Original	
1	04/29/03	51		101	
2	04/25/03	52		102	
3	04/04/03	53		103	
4	04/04/03	54		104	
5	04/04/03	55		105	
6	04/04/03	56		106	
7	04/04/03	57		107	
8	04/04/03	58		108	
9	04/04/03	59		109	
10	04/04/03	60		110	
11	04/04/03	61		111	
12	04/04/03	62		112	
13	04/04/03	63		113	
14	04/04/03	64		114	
15	04/04/03	65		115	
16	04/04/03	66		116	
17	04/04/03	67		117	
18	04/04/03	68		118	
19	04/04/03	69		119	
20	04/04/03	70		120	
21	04/04/03	71		121	
22	04/04/03	72		122	
23	04/04/03	73		123	
24	04/04/03	74		124	
25	04/04/03	75		125	
26	04/04/03	76		126	
27	04/04/03	77		127	
28	04/04/03	78		128	
29	04/04/03	79		129	
30	04/04/03	80		130	
31	04/04/03	81		131	
32	04/04/03	82		132	
33	04/04/03	83		133	
34	04/04/03	84		134	
35	04/04/03	85		135	
36	04/04/03	86		136	
37	04/04/03	87		137	
38	04/04/03	88		138	
39	04/04/03	89		139	
40	04/04/03	90		140	
41	04/04/03	91		141	
42	04/04/03	92		142	
43	04/04/03	93		143	
44	04/04/03	94		144	
45	04/04/03	95		145	
46	04/04/03	96		146	
47	04/04/03	97		147	
48	04/04/03	98		148	
49	04/04/03	99		149	
50	04/04/03	100		150	

If more than 150 claims or 9 actions staple additional sheet here

SM 864 4111 03 20876 Reac 57/05/10/03
 2/10/03

SEARCH			
Class	Sub.	Date	Exmr.
370	252	02-10-2004	mm
	352		
	230		
	235		
	238		
done	updated	05-25-2004	mm

INTERFERENCE SEARCHED			
Class	Sub.	Date	Exmr.
370	238	05-25-2004	mm
	252		
	352		

SEARCH NOTES		
(List databases searched. Attach search strategy inside.)		
	Date	Exmr.
cont 1998 explore	02-10-2004	mm
cont	05-25-2004	mm



US006775235B2

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

(10) **Patent No.:** **US 6,775,235 B2**
(45) **Date of Patent:** **Aug. 10, 2004**

- (54) **TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS**
- (75) Inventors: **Sanchaita Datta**, Salt Lake City, UT (US); **Ragula Bhaskar**, Salt Lake City, UT (US)
- (73) Assignee: **Ragula Systems**, Salt Lake City, UT (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.

5,948,069	A	9/1999	Kitai et al.	709/240
6,016,307	A	* 1/2000	Kaplan et al.	370/238
6,119,170	A	* 9/2000	Schoffelman et al.	709/244
6,128,298	A	* 10/2000	Wootton et al.	370/392
6,253,247	B1	6/2001	Bhaskar et al.	709/237
6,295,276	B1	9/2001	Datta et al.	370/218
6,339,595	B1	1/2002	Rekhter et al.	370/392
6,438,100	B1	8/2002	Halpern et al.	370/218
6,449,259	B1	9/2002	Allain et al.	370/253
6,456,594	B1	9/2002	Kaplan et al.	370/238
6,493,341	B1	12/2002	Datta et al.	370/392
6,493,349	B1	12/2002	Casey	370/409
6,665,702	B1	* 12/2003	Zisapel et al.	718/105

- (21) Appl. No.: **10/361,837**
- (22) Filed: **Feb. 7, 2003**
- (65) **Prior Publication Data**
US 2003/0147408 A1 Aug. 7, 2003

Related U.S. Application Data

- (63) Continuation-in-part of application No. 10/034,197, filed on Dec. 28, 2001.
- (60) Provisional application No. 60/355,509, filed on Feb. 8, 2002, and provisional application No. 60/259,269, filed on Dec. 29, 2000.
- (51) **Int. Cl.⁷** **H04L 12/64**
- (52) **U.S. Cl.** **370/238; 370/252; 370/352**
- (58) **Field of Search** **370/252, 352, 370/230, 235, 238**

References Cited

U.S. PATENT DOCUMENTS

5,398,012	A	3/1995	Derby et al.	340/825.03
5,420,862	A	5/1995	Perlman	370/85.13
5,473,599	A	12/1995	Li et al.	370/16
5,737,526	A	4/1998	Periasamy et al.	395/200.06
5,898,673	A	4/1999	Riggan et al.	370/237

OTHER PUBLICATIONS

- 'Radware announces LinkProof: The first IP Load Balancing Solution for networks with multiple ISP connection', Press Release, published Oct. 7, 1999.*
- 'Radware Balances the Network', Internet Traffic Management Center, published Jan. 1, 2000.*
- 'Global Product Spotlight: Radware Linkproof', Network-Magazine.com, published Dec. 1, 1999.*

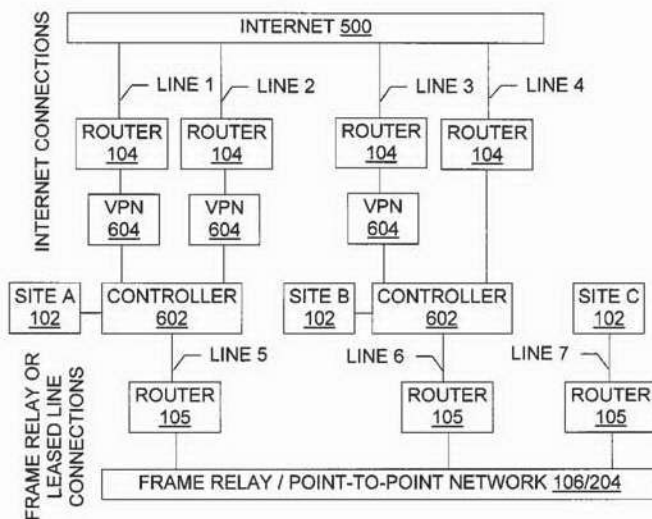
(List continued on next page.)

Primary Examiner—Melvin Marcelo
(74) *Attorney, Agent, or Firm*—Thorpe North & Western LLP

(57) **ABSTRACT**

Methods, configured storage media, and systems are provided for communications using two or more disparate networks in parallel to provide load balancing across network connections, greater reliability, and/or increased security. A controller provides access to two or more disparate networks in parallel, through direct or indirect network interfaces. When one attached network fails, the failure is sensed by the controller and traffic is routed through one or more other disparate networks. When all attached disparate networks are operating, one controller preferably balances the load between them.

24 Claims, 6 Drawing Sheets



OTHER PUBLICATIONS

'Radware Seeks Solutions to Easy-Access Problems', South China Morning Post, published Dec. 7, 1999.*

B. Gleeson et al., "A Framework for IP Based Virtual Private Networks," RFC 2764 (Feb. 2000).

U.S. patent application, Attorney Docket No. 3003.2.9A; see USPTO published application No. US-2002-0087724-A1, Jul. 4, 2002.

T. Liao et al., "Using multiple links to interconnect LANs and public circuit switched data networks," *Proc. Int. Conference on Communications Systems: Towards Global Integration*, vol. 1, Singapore, 59 Nov. 1990, pp. 289-293.

Press release from www.coyotepoint.com, Sep. 8, 1997.

Network Address Translation Technical Discussion, from safety.net; no later than May 7, 1999.

Higginson et al., "Development of Router Clusters to Provide Fast Failover in IP Networks," from www.asia-pacific.digital.com; no later than Sep. 29, 1998.

Pages from www.navpoint.com; no later than Dec. 24, 2001.

"The Basic Guide to Frame Relay Networking", pp. 1-85, copyright date 1998.

"NNI & UNI", pp. 1-2, Nov. 16, 2001.

"Disaster Recovery for Frame Relay Networks", pp. 1-14, no later than Dec. 7, 2001.

T. Nolle, "Watching Your Back", pp. 1-3, Nov. 1, 1999.

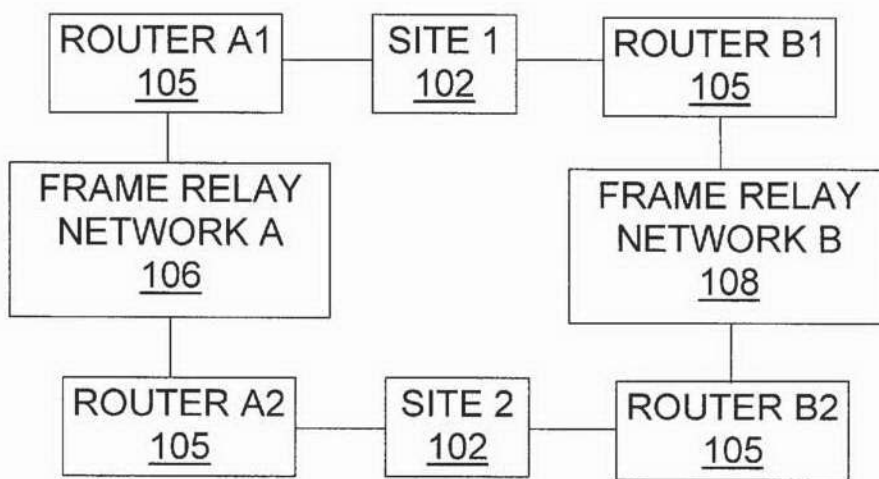
"Multi-Attached and Multi-Homed Dedicated Access", pp. 1-5, no later than Dec. 8, 2001.

Feibel, "Internetwork Link," Novell's® Complete Encyclopedia of Networking, copyright date 1995.

Tanenbaum, *Computer Networks* (3rd Ed.), pp. 396-406; copyright date 1996.


Wexler, "Frame Relay and IP VPNs: Compete Or Coexist?", from www.bcr.com; Jul. 1999.

* cited by examiner




(PRIOR ART)

Fig. 1




(PRIOR ART)

Fig. 2




(PRIOR ART)

Fig. 3



(PRIOR ART)

Fig. 4



(PRIOR ART)

Fig. 5




Fig. 7

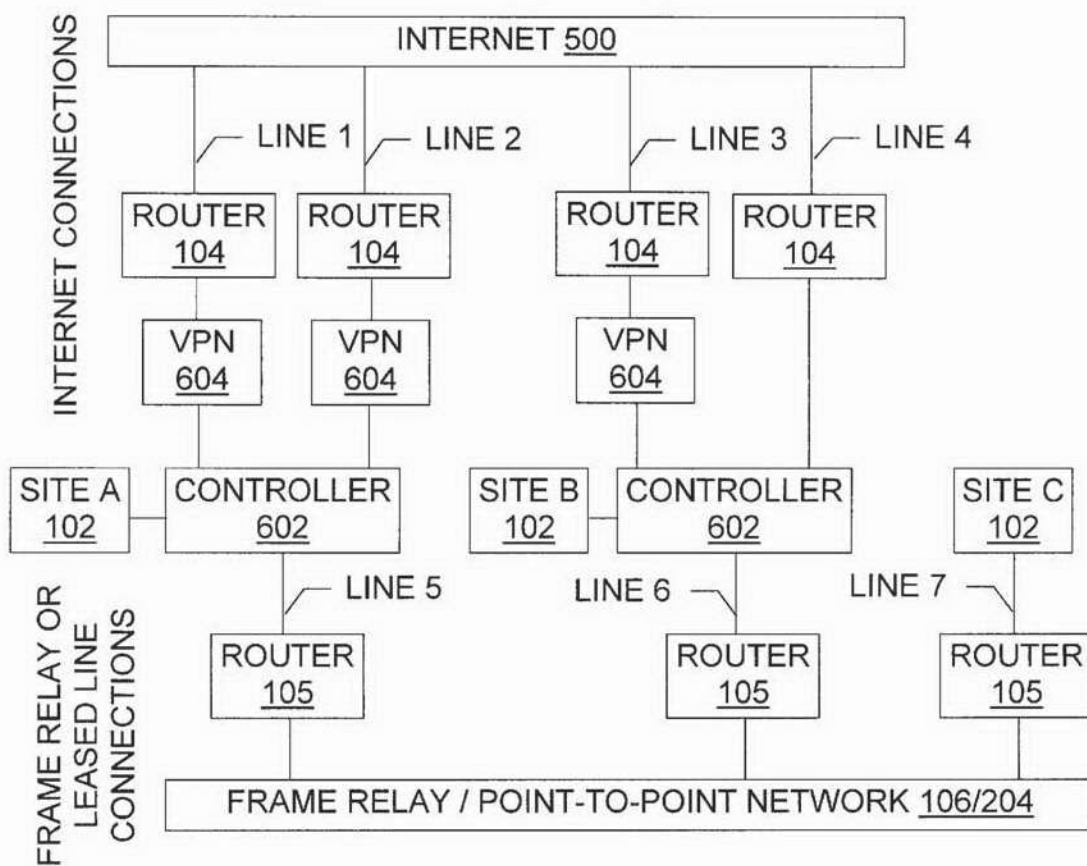


Fig. 6




Fig. 10




Fig. 8

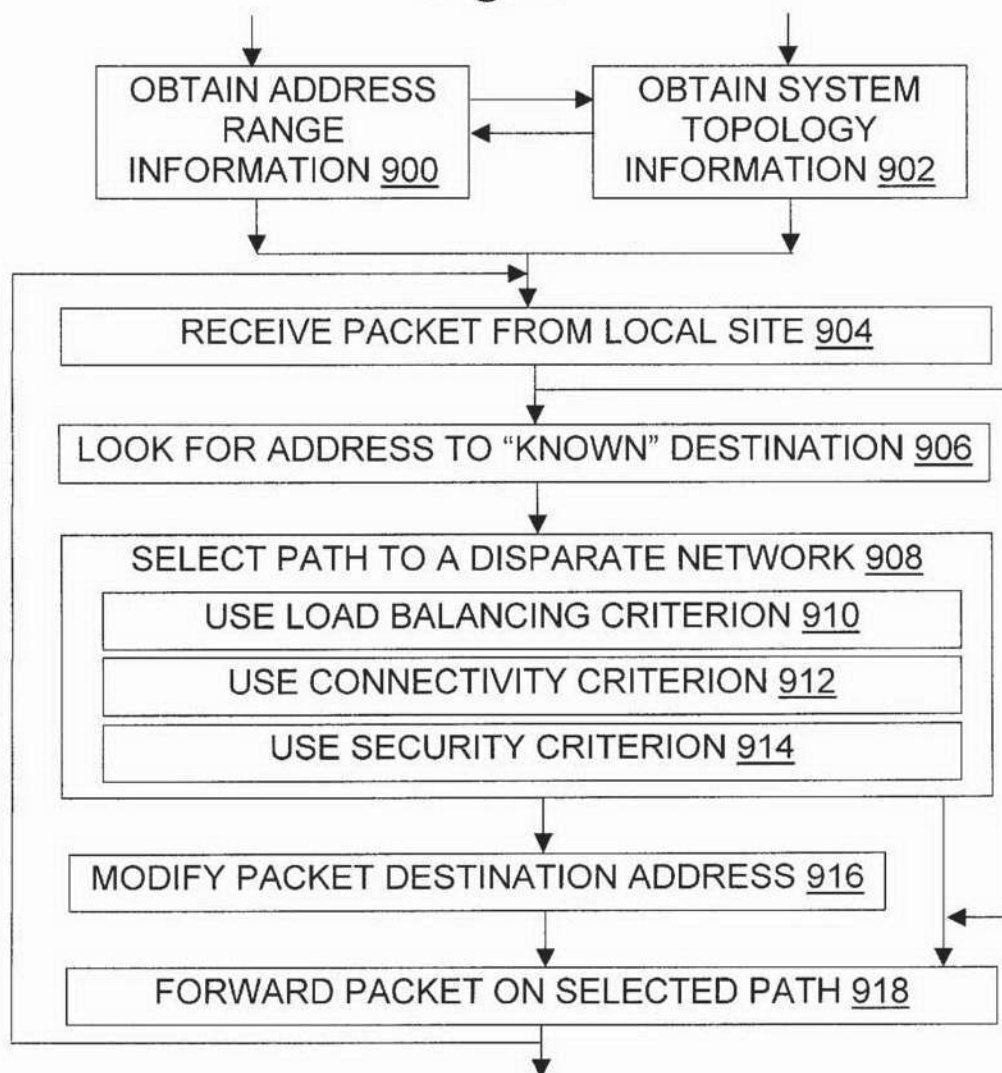


Fig. 9




Fig. 11

1

TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS

RELATED APPLICATIONS

This application claims priority to commonly owned copending U.S. provisional patent application serial No. 60/355,509 filed Feb. 8, 2002, which is also incorporated herein by reference. This application is a continuation-in-part of U.S. patent application Ser. No. 10/034,197 filed Dec. 28, 2001, which claims priority to U.S. provisional patent application serial No. 60/259,269 filed Dec. 29, 2000, each of which is also incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates to computer network data transmission, and more particularly relates to tools and techniques for communications using disparate parallel networks, such as a virtual private network ("VPN") or the Internet in parallel with a point-to-point, leased line, or frame relay network, in order to help provide benefits such as load balancing across network connections, greater reliability, and increased security.

TECHNICAL BACKGROUND OF THE INVENTION

Organizations have used frame relay networks and point-to-point leased line networks for interconnecting geographically dispersed offices or locations. These networks have been implemented in the past and are currently in use for interoffice communication, data exchange and file sharing. Such networks have advantages, some of which are noted below. But these networks also tend to be expensive, and there are relatively few options for reliability and redundancy. As networked data communication becomes critical to the day-to-day operation and functioning of an organization, the need for lower cost alternatives for redundant back-up for wide area networks becomes important.

Frame relay networking technology offers relatively high throughput and reliability. Data is sent in variable length frames, which are a type of packet. Each frame has an address that the frame relay network uses to determine the frame's destination. The frames travel to their destination through a series of switches in the frame relay network, which is sometimes called a network "cloud"; frame relay is an example of packet-switched networking technology. The transmission lines in the frame relay cloud must be essentially error-free for frame relay to perform well, although error handling by other mechanisms at the data source and destination can compensate to some extent for lower line reliability. Frame relay and/or point-to-point network services are provided or have been provided by various carriers, such as AT&T, Qwest, XO, and MCI WorldCom.

Frame relay networks are an example of a network that is "disparate" from the Internet and from Internet-based virtual private networks for purposes of the present invention. Another example of such a "disparate" network is a point-to-point network, such as a T1 or T3 connection. Although the underlying technologies differ somewhat, for purposes of the present invention frame relay networks and point-to-point networks are generally equivalent in important ways, such as the conventional reliance on manual switchovers when traffic must be redirected after a connection fails, and their implementation distinct from the Internet. A frame relay permanent virtual circuit is a virtual point-to-point

2

connection. Frame relays are used as examples throughout this document, but the teachings will also be understood in the context of point-to-point networks.

A frame relay or point-to-point network may become suddenly unavailable for use. For instance, both MCI WorldCom and AT&T users have lost access to their respective frame relay networks during major outages. During each outage, the entire network failed. Loss of a particular line or node in a network is relatively easy to work around. But loss of an entire network creates much larger problems.

Tools and techniques to permit continued data transmission after loss of an entire frame relay network that would normally carry data are discussed in U.S. patent application Ser. No. 10/034,197 filed Dec. 28, 2001 and incorporated herein. The '197 application focuses on architectures involving two or more "private" networks in parallel, whereas the present application focuses on architectures involving disparate networks in parallel, such as a proprietary frame relay network and the Internet. Note that the term "private network" is used herein in a manner consistent with its use in the '197 application (which comprises frame relay and point-to-point networks), except that a "virtual private network" as discussed herein is not a "private network". Virtual private networks are Internet-based, and hence disparate from private networks, i.e., from frame relay and point-to-point networks. To reduce the risk of confusion that might arise from misunderstanding "private network" to comprise "virtual private network" herein, virtual private networks will be henceforth referred to as VPNs. Other differences and similarities between the present application and the '197 application will also be apparent to those of skill in the art on reading the two applications.

Various architectures involving multiple networks are known in the art. For instance, FIG. 1 illustrates prior art configurations involving two frame relay networks for increased reliability; similar configurations involve one or more point-to-point network connections. Two sites 102 transmit data to each other (alternately, one site might be only a data source, while the other is only a data destination). Each site has two border routers 105. Two frame relay networks 106, 108 are available to the sites 102 through the routers 105. The two frame relay networks 106, 108 have been given separate numbers in the figure, even though each is a frame relay network, to emphasize the incompatibility of frame relay networks provided by different carriers. An AT&T frame relay network, for instance, is incompatible—in details such as maximum frame size or switching capacity—with an MCI WorldCom frame relay network, even though they are similar when one takes the broader view that encompasses disparate networks like those discussed herein. The two frame relay providers have to agree upon information rates, switching capacities, frame sizes, etc. before the two networks can communicate directly with each other.

A configuration like that shown in FIG. 1 may be actively and routinely using both frame relay networks A and B. For instance, a local area network (LAN) at site 1 may be set up to send all traffic from the accounting and sales departments to router A1 and send all traffic from the engineering department to router B1. This may provide a very rough balance of the traffic load between the routers, but it does not attempt to balance router loads dynamically in response to actual traffic and thus is not "load-balancing" as that term is used herein.

Alternatively, one of the frame relay networks may be a backup which is used only when the other frame relay

3

network becomes unavailable. In that case, it may take even skilled network administrators several hours to perform the steps needed to switch the traffic away from the failed network and onto the backup network, unless the invention of the '197 application is used. In general, the necessary Private Virtual Circuits (PVCs) must be established, routers at each site **102** must be reconfigured to use the correct serial links and PVCs, and LANs at each site **102** must be reconfigured to point at the correct router as the default gateway.

Although two private networks are shown in FIG. 1, three or more such networks could be employed, with similar considerations coming into play as to increased reliability, limits on load-balancing, the efforts needed to switch traffic when a network fails, and so on. Likewise, for clarity of illustration FIG. 1 shows only two sites, but three or more sites could communicate through one or more private networks.

FIG. 2 illustrates a prior art configuration in which data is normally sent between sites **102** over a private network **106**. A failover box **202** at each site **102** can detect failure of the network **106** and, in response to such a failure, will send the data instead over an ISDN link **204** while the network **106** is down. Using an ISDN link **204** as a backup is relatively easier and less expensive than using another private network **106** as the backup, but generally provides lower throughput. The ISDN link is an example of a point-to-point or leased line network link.

FIG. 3 illustrates prior art configurations involving two private networks for increased reliability, in the sense that some of the sites in a given government agency or other entity **302** can continue communicating even after one network goes down. For instance, if a frame relay network A goes down, sites **1**, **2**, and **3** will be unable to communicate with each other but sites **4**, **5**, and **6** will still be able to communicate amongst themselves through frame relay network B. Likewise, if network B goes down, sites **1**, **2**, and **3** will still be able to communicate through network A. Only if both networks go down at the same time would all sites be completely cut off. Like the FIG. 1 configurations, the FIG. 3 configuration uses two private networks. Unlike FIG. 1, however, there is no option for switching traffic to another private network when one network **106** goes down, although either or both of the networks in FIG. 3 could have an ISDN backup like that shown in FIG. 2. Note also that even when both private networks are up, sites **1**, **2**, and **3** communicate only among themselves; they are not connected to sites **4**, **5**, and **6**. Networks A and B in FIG. 3 are therefore not in "parallel" as that term is used herein, because all the traffic between each pair of sites goes through at most one of the networks A, B.

FIG. 4 illustrates a prior art response to the incompatibility of frame relay networks of different carriers. A special "network-to-network interface" (NNI) **402** is used to reliably transmit data between the two frame relay networks A and B. NNIs are generally implemented in software at carrier offices. Note that the configuration in FIG. 4 does not provide additional reliability by using two frame relay networks **106**, because those networks are in series rather than in parallel. If either of the frame relay networks A, B in the FIG. 4 configuration fails, there is no path between site **1** and site **2**; adding the second frame relay network has not increased reliability. By contrast, FIG. 1 increases reliability by placing the frame relay networks in parallel, so that an alternate path is available if either (but not both) of the frame relay networks fails. Someone of skill in the art who was looking for ways to improve reliability by putting networks

4

in parallel would probably not consider NNIs pertinent, because they were used for serial configurations rather than parallel ones, and adding networks in a serial manner does not improve reliability.

Internet-based communication solutions such as VPNs and Secure Sockets Layer (SSL) offer alternatives to frame relay **106** and point-to-point leased line networks such as those using an ISDN link **204**. These Internet-based solutions are advantageous in the flexibility and choice they offer in cost, in service providers, and in vendors. Accordingly, some organizations have a frame relay **106** or leased line connection (a.k.a. point-to-point) for intranet communication and also have a connection for accessing the Internet **500**, using an architecture such as that shown in FIG. 5.

But better tools and techniques are needed for use in architectures such as that shown in FIG. 5. In particular, prior approaches for selecting which network to use for which packet(s) are coarse. For instance, all packets from department X might be sent over the frame relay connection **106** while all packets from department Y are sent over the Internet **500**. Or the architecture might send all traffic over the frame relay network unless that network fails, and then be manually reconfigured to send all traffic over a VPN **502**.

Organizations are still looking for better ways to use Internet-based redundant connections to backup the primary frame relay networks. Also, organizations wanting to change from frame relay and point-to-point solutions to Internet-based solutions have not had the option of transitioning in a staged manner. They have had to decide instead between the two solutions, and deploy the solution in their entire network communications system in one step. This is a barrier for deployment of Internet-based solutions **500/502**, since an existing working network would be replaced by a yet-untested new network. Also, for organizations with several geographically distributed locations a single step conversion is very complex. Some organizations may want a redundant Internet-based backup between a few locations while maintaining the frame relay network for the entire organization.

It would be an advancement in the art to provide new tools and techniques for configuring disparate networks (e.g., frame relay/point-to-point WANs and Internet-based VPNs) in parallel, to obtain benefits such as greater reliability, improved security, and/or load-balancing. Such improvements are disclosed and claimed herein.

BRIEF SUMMARY OF THE INVENTION

The present invention provides tools and techniques for directing packets over multiple parallel disparate networks, based on addresses and other criteria. This helps organizations make better use of frame relay networks and/or point-to-point (e.g., T1, T3, fiber, OCx, Gigabit, wireless, or satellite based) network connections in parallel with VPNs and/or other Internet-based networks. For instance, some embodiments of the invention allow frame relay and VPN wide area networks to co-exist for redundancy as well as for transitioning from frame relay/point-to-point solutions to Internet-based solutions in a staged manner. Some embodiments operate in configurations which communicate data packets over two or more disparate WAN connections, with the data traffic being dynamically load-balanced across the connections, while some embodiments treat one of the WANs as a backup for use mainly in case the primary connection through the other WAN fails.

Other features and advantages of the invention will become more fully apparent through the following description.

5

BRIEF DESCRIPTION OF THE DRAWINGS

To illustrate the manner in which the advantages and features of the invention are obtained, a more particular description of the invention will be given with reference to the attached drawings. These drawings only illustrate selected aspects of the invention and its context. In the drawings:

FIG. 1 is a diagram illustrating a prior art approach having frame relay networks configured in parallel for increased reliability for all networked sites, in configurations that employ manual switchover between the two frame relay networks in case of failure.

FIG. 2 is a diagram illustrating a prior art approach having a frame relay network configured in parallel with an ISDN network link for increased reliability for all networked sites.

FIG. 3 is a diagram illustrating a prior art approach having independent and non-parallel frame relay networks, with each network connecting several sites but no routine or extensive communication between the networks.

FIG. 4 is a diagram illustrating a prior art approach having frame relay networks configured in series through a network-to-network interface, with no consequent increase in reliability because the networks are in series rather than in parallel.

FIG. 5 is a diagram illustrating a prior art approach having a frame relay network configured in parallel with a VPN or other Internet-based network that is disparate to the frame relay network, but without the fine-grained packet routing of the present invention.

FIG. 6 is a diagram illustrating one system configuration of the present invention, in which the Internet and a private network are placed in parallel for increased reliability for all networked sites, without requiring manual traffic switchover, and with the option in some embodiments of load balancing between the networks and/or increasing security by transmitting packets of a single logical connection over disparate networks.

FIG. 7 is a diagram further illustrating a multiple disparate network access controller of the present invention, which comprises an interface component for each network to which the controller connects, and a path selector in the controller which uses one or more of the following as criteria: destination address, network status (up/down), network load, use of a particular network for previous packets in a given logical connection or session.

FIG. 8 is a flowchart illustrating methods of the present invention for sending packets using a controller such as the one shown in FIG. 7.

FIG. 9 is a flowchart illustrating methods of the present invention for combining connections to send traffic over multiple parallel independent disparate networks for reasons such as enhanced reliability, load balancing, and/or security.

FIG. 10 is a diagram illustrating another system configuration of the present invention, in which the Internet and a frame relay network are placed in parallel, with a VPN tunnel originating after the source controller and terminating before the destination controller, and each known site that is accessible through one network is also accessible through the other network unless that other network fails.

FIG. 11 is a diagram illustrating a system configuration similar to FIG. 10, except the VPN tunnel originates before the source controller and terminates after the destination controller.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention relates to methods, systems, and configured storage media for connecting sites over multiple

6

independent parallel disparate networks, such as frame relay networks and/or point-to-point network connections, on the one hand, and VPNs or other Internet-based network connections, on the other hand. "Multiple" networks means two or more such networks. "Independent" means routing information need not be shared between the networks. "Parallel" does not rule out all use of NNIs and serial networks, but it does require that at least two of the networks in the configuration be in parallel at the location where the invention distributes traffic, so that alternate data paths through different networks are present. "Frame relay networks" or "private networks" does not rule out the use of an ISDN link or other backup for a particular frame relay or point-to-point private network, but it does require the presence of multiple such networks; FIG. 2, for instance, does not meet this requirement. A "frame relay network" is unavailable to the general public and thus disparate from the Internet and VPNs (which may be Internet-based), even though some traffic in the Internet may use public frame relay networks once the traffic leaves the location where the invention distributes traffic.

FIG. 6 illustrates one of many possible configurations of the present invention. Comments made here also apply to similar configurations involving only one or more frame relay networks **106**, those involving only one or more point-to-point networks **204**, and those not involving a VPN **604**, for example. Two or more disparate networks are placed in parallel between two or more sites **102**. In the illustrated configuration, the Internet **500** and a VPN **604** are disparate from, and in parallel with, frame relay/point-to-point network **106/204**, with respect to site A and site B. No networks are parallel disparate networks in FIG. 6 with regard to site C as a traffic source, since that site is not connected to the Internet **500**. Access to the disparate networks at site A and site B is through an inventive controller **602** at each site. Additional controllers **602** may be used at each location (i.e., controllers **602** may be placed in parallel to one another) in order to provide a switched connection system with no single point of failure.

With continued attention to the illustrative network topology for one embodiment of the invention shown in FIG. 6, in this topology the three locations A, B, and C are connected to each other via a frame relay **106** or leased line network **204**. Assume, for example, that all three locations are connected via a single frame relay network **106**. Locations A and B are also connected to each other via a VPN connection **604**. VPN tunnels are established between locations A and B in the VPN, which pairs line 1 to line 3 and also pairs line 2 to line 3. There can be only one VPN tunnel between locations A and B. There is no VPN connection between location C and either location A or location B.

Therefore, locations A, B, and C can communicate with each other over the frame relay network **106**, and locations A and B (but not C) can also communicate with each other over the VPN connection **604**. Communication between locations A and C, and communication between locations B and C, can take place over the frame relay network **106** only. Communication between locations A and B can take place over frame relay network **106**. It can also take place over one of the lines 1-and-3 pair, or the lines 2-and-3 pair, but not both at the same time. Traffic can also travel over lines 2 and 4, but without a VPN tunnel. When the source and destination IP address pairs are the same between locations A and B but different types of networks connect those locations, as in FIG. 6 for instance, then a traffic routing decision that selects between network types cannot be made with an existing commercially available device. By contrast, the

7

invention allows an organization to deploy an Internet-based solution between locations A and B while maintaining the frame relay network 106 between locations A, B, and C, and allows traffic routing that selects between the Internet and the frame relay network on a packet-by-packet basis.

The invention may thus be configured to allow the organization to achieve the following goals, in the context of FIG. 6; similar goals are facilitated in other configurations. First, the organization can deploy an Internet-based second connection between only locations A and B, while maintaining frame relay connectivity between locations A, B, and C. Later the organization may deploy an Internet-based solution at location C as well. Second, the organization can use the Internet-based connection between locations A and B for full load-balancing or backup, or a combination of the two. Third, the organization can use the frame relay connection between locations A and B for full load-balancing or backup, or a combination of the two. Fourth, the organization can load-balance traffic in a multi-homing situation between two ISPs or two connections to the Internet at locations A and/or B.

To better understand the invention, consider the operation of controller device 602 at location A. The controller 602 examines the IP data traffic meant to go through it and makes determinations and takes steps such as those discussed below.

If the traffic is destined for the Internet 500, send the traffic over the Internet using lines 1 and/or 2. Load balancing decisions that guide the controller 602 in distributing packets between the lines can be based on criteria such as the load of a given network, router, or connection relative to other networks, routers, or connections, to be performed dynamically in response to actual traffic. Load-balancing may be done through a round-robin algorithm which places the next TCP or UDP session on the next available line, or it may involve more complex algorithms that attempt to measure and track the throughput, latency, and/or other performance characteristics of a given link or path element. Load-balancing is preferably done on a per-packet basis for site-to-site data traffic over the Internet or frame relay net, or done on a TCP or UDP session basis for Internet traffic, as opposed to prior approaches that use a per-department and/or per-router basis for dividing traffic. Load-balancing algorithms in general are well understood, although their application in the context of the present invention is believed to be new.

If the traffic is destined for location B, then there are at least three paths from the current location (A) to location B: frame relay line 5, VPN line 1, or Internet line 2. In some embodiments, the invention determines whether the three connections are in load-balance mode or on-failure backup mode or a combination thereof. For a load-balance mode, the controller 602 chooses the communication line based on load-balancing criteria. For backup mode, it chooses the communication line that is either the preferred line or (if the preferred line is down) the currently functional (backup) line.

By contrast with the preceding, if the traffic is destined for location C, then the controller 602 at site A sends the traffic on the frame relay line, line 5.

Now let us look at the operation of the controller device 602 at location B. The device examines the IP data traffic sent to it and makes determinations like the following:

1. Is the traffic destined for the Internet, as opposed to one of the three "known" locations A, B, and C? If so, send the traffic over the Internet lines (line 3 and/or line 4). Load balancing decisions can be based on the criteria described above.

8

2. Is the traffic destined for location A? If so, then there are at least two paths to location A: the frame relay line 6, or VPN line 3. The controller 602 decides whether the two connections are in load-balance or on-failure backup mode, and chooses line(s) accordingly as discussed above.

3. Is the traffic destined for location C? If so, then send the traffic on the frame relay line, line 6.

To operate as discussed herein, the invention uses information about the IP address ranges in the locations reside as input data. For instance, a packet destined for the Internet 500 is one whose destination address is not in any of the address ranges of the known locations (e.g., locations A, B, and C in the example of FIG. 6). In some configurations, this is the same as saying that a packet destined for the Internet is one whose address is not in the address range of any of the organization's locations. However, although all the known locations may belong to a single organization, that is not a necessary condition for using the invention. Known locations may also belong to multiple organizations or individuals. Likewise, other locations belonging to the organization may be unknown for purposes of a given embodiment of the invention.

Address ranges can be specified and tested by the controller 602 using subnet masks. The subnet masks may be of different lengths (contain a different number of one bits) in different embodiments and/or in different address ranges in a given embodiment. For instance, class B and class C addresses may both be used in some embodiments.

As another example, consider the illustrative network topology shown in FIG. 10. This configuration has two locations A and B which are connected by a frame relay network 106 and by the Internet 500, through a frame relay router 105 and an Internet router 104, at each location. For convenience, all routers are designated similarly in the Figures, but those of skill in the art will appreciate that different router models may be used, and in particular and without limitation, different routers may be used to connect to a private network than are used to connect to the Internet. Also, the controllers 602, routers (and in FIG. 6 the VPN interfaces 604) are shown separately in the Figures for convenience and clarity of illustration, but in various embodiments the software and/or hardware implementing these devices 602, 104, 105, 604 may be housed in a single device and/or reside on a single machine.

Suppose that the address ranges used by the routers in the FIG. 10 configuration are the following:

Location	LAN IP	Internet	Frame Relay
A	192.168.x.x	200.x.x.x	196.x.x.x
B	10.0.x.x	210.x.x.x	198.x.x.x

Without the invention, a topology like FIG. 10 (but without the controllers 602) requires some inflexible method of assigning packets to paths. Thus, consider a packet from location A that is meant for location B that has a destination address in the 10.0.x.x range. The network devices are pre-configured to such that all such packets with the 10.0.x.x destination address must be sent to the frame relay router (router Y), even though there is Internet connectivity between the two locations. Likewise, without the invention a packet from location A meant for location B which has a destination address not in the 10.0.x.x. range must be sent to the Internet router (router X) even though there is frame relay connectivity between the two locations.

9

Traditionally, such necessary match-ups of packets with routers were done by inflexible approaches such as sending all traffic from a given department, building, or local area network to a specified router. Manual and/or tedious reconfiguration was needed to change the destination address used in packets from a given source LAN such as one at site A, so this approach allowed load-balancing only on a very broad granularity, and did not load-balance dynamically in response to actual traffic. In particular, difficult reconfiguration of network parameters was needed to redirect packets to another router when the specified router went down.

By placing inventive modules 602 between locations and their routers as illustrated in FIG. 10, however, the invention allows load-balancing, redundancy, or other criteria to be used dynamically, on a granularity as fine as packet-by-packet, to direct packets to an Internet router and/or a frame relay/point-to-point router according to the criteria. For instance, with reference to the illustrative network topology of FIG. 10, if the inventive module 602 at location A receives a packet with a destination address in the 10.0.x.x range and the Internet router X is either down or overloaded, then the inventive module 602 can change the destination address so that it is in the 198.x.x.x range (the rest of the address may be kept) and then send the modified packet to the frame relay router Y. Similarly, if the frame relay path is down, overloaded, or insecure, then the controller 602 can direct packets to the Internet after making the necessary destination address changes to let the Internet router 104 operate successfully on those packets.

Unlike the configuration shown in FIG. 1, the inventive configurations in FIGS. 6 and 10 do not require manual intervention by network administrators to coordinate traffic flow over parallel networks. The disparate networks are independent of each other. When one attached network fails, the failure is sensed by the controller 602 and traffic is automatically routed through one or more other networks. Unlike the configuration in FIG. 2, the inventive configuration combines two or more disparate networks. Unlike the configuration in FIG. 4, the inventive configuration requires two or more disparate networks be placed in parallel (although additional networks may also be placed in series). Unlike the configuration in FIG. 3, the inventive configuration does not merely partition sites between unconnected networks—with the invention, most or all of the connected sites get the benefit of parallel networks, so they can continue transceiving even if one of the networks goes down.

Another difference between the inventive approach and prior approaches is the narrow focus of some prior art on reliability. The present document takes a broader view, which considers load balancing and security as well as reliability. Configurations like those shown in FIG. 2 are directed to reliability (which is also referred to by terms such as “fault tolerance”, “redundancy”, “backup”, “disaster recovery”, “continuity”, and “failover”). That is, one of the network paths (in this case, the one through the frame relay network) is the primary path, in that it is normally used for most or all of the traffic, while the other path (in this case, the one through the ISDN link) is used only when that primary path fails. Although the inventive configurations can be used in a similar manner, with one network being on a primary path and the other network(s) being used only as a backup when that first network fails, the inventive configurations also permit concurrent use of two or more disparate networks. With concurrent use, elements such as load balancing between disparate networks, and increased security by means of splitting pieces of a given message

10

between disparate networks, which are not considerations in the prior art of FIG. 2, become possibilities in some embodiments of the present invention.

In some embodiments, a network at a location T is connected to a controller 602 for a location R but is not necessarily connected to the controller 602 at another location S. In such cases, a packet from location T addressed to location S can be sent over the network to the controller at location S, which can then redirect the packet to location T by sending it over one or more parallel disparate networks. That is, controllers 602 are preferably, but not necessarily, provided at every location that can send packets over the parallel independent networks of the system.

In some embodiments, the controller 602 at the receiving end of the network connection between two sites A and B has the ability to re-sequence the packets. This means that if the lines are of dissimilar speeds or if out-of-sequence transmission is required by security criteria, the system can send packets out of order and re-sequence them at the other end. Packets may be sent out of sequence to enhance security, to facilitate load-balancing, or both. The TCP/IP packet format includes space for a sequence number, which can be used to determine proper packet sequence at the receiving end (the embodiments are dual-ended, with a controller 602 at the sending end and another controller 602 at the receiving end). The sequence number (and possibly more of the packet as well) can be encrypted at the sending end and then decrypted at the receiving end, for enhanced security. Alternately, an unused field in the TCP/IP header can hold alternate sequence numbers to define the proper packet sequence.

In the operation of some embodiments, the controller 602 on each location is provided with a configuration file or other data structure containing a list of all the LAN IP addresses of the controllers 602 at the locations, and their subnet masks. Each controller 602 keeps track of available and active connections to the remote sites 102. If any of the routes are unavailable, the controller 602 preferably detects and identifies them. When a controller 602 receives IP traffic to any of the distant networks, the data is sent on the active connection to that destination. If all connections are active and available, the data load is preferably balanced across all the routers. If any of the connections are unavailable, or any of the routers are down, the traffic is not forwarded to that router; when the routes become available again, the load balancing across all active routes preferably resumes.

In some embodiments, load balancing is not the only factor considered (or is not a factor considered) when the controller 602 determines which router should receive a given packet. Security may be enhanced by sending packets of a given message over two or more disparate networks. Even if a packet sniffer or other eavesdropping tool is used to illicitly obtain data packets from a given network, the eavesdropper will thus obtain at most an incomplete copy of the message because the rest of the message traveled over a different network. Security can be further enhanced by sending packets out of sequence, particularly if the sequence numbers are encrypted.

FIG. 7 is a diagram further illustrating a multiple disparate network access controller 602 of the present invention. A site interface 702 connects the controller 602 to the LAN at the site 102. This interface 702 can be implemented, for instance, as any local area network interface, like 10/100Base-T ethernet, gigabit ATM or any other legacy or new LAN technology.

The controller 602 also includes a packet path selector 704, which may be implemented in custom hardware, or

11

implemented as software configuring semi-custom or general-purpose hardware. The path selector 704 determines which path to send a given packet on. In the configuration of FIG. 6, for instance, the path selector in the controller at location A selects between a path through the router on line 1 and a path through the router on line 2. In different embodiments and/or different situations, one or more of the following criteria may be used to select a path for a given packet, for a given set of packets, and/or for packets during a particular time period:

Redundancy: do not send the packet(s) to a path through a network, a router, or a connection that is apparently down. Instead, use devices (routers, network switches, bridges, etc.) that will still carry packets after the packets leave the selected network interfaces, when other devices that could have been selected are not functioning. Techniques and tools for detecting network path failures are generally well understood, although their application in the context of the present invention is believed to be new.

Load-balancing: send packets in distributions that balance the load of a given network, router, or connection relative to other networks, routers, or connections available to the controller 602. This promotes balanced loads on one or more of the devices (routers, frame relay switches, etc.) that carry packets after the packets leave the selected network interfaces. Load-balancing may be done through an algorithm as simple as a modified round-robin approach which places the next packet on the next available line, or it may involve more complex algorithms that attempt to measure and track the throughput, latency, and/or other performance characteristics of a given link or path element. Load-balancing is preferably done on a per-packet basis for site-to-site data traffic or on a TCP or UDP session basis for Internet traffic, as opposed to prior art approaches which use a per-department and/or per-router basis for dividing traffic. Load-balancing algorithms in general are well understood, although their application in the context of the present invention is believed to be new.

Security: divide the packets of a given message (session, file, web page, etc.) so they travel over two or more disparate networks, so that unauthorized interception of packets on fewer than all of the networks used to carry the message will not provide the total content of the message. Dividing message packets between networks for better security may be done in conjunction with load balancing, and may in some cases be a side-effect of load-balancing. But load-balancing can be done on a larger granularity scale than security, e.g., by sending one entire message over a first network A and the next entire message over a second, disparate network. Security may thus involve finer granularity than load balancing, and may even be contrary to load balancing in the sense that dividing up a message to enhance security may increase the load on a heavily loaded path even though a more lightly loaded alternate path is available and would be used for the entire message if security was not sought by message-splitting between networks. Other security criteria may also be used, e.g., one network may be viewed as more secure than another, encryption may be enabled, or other security measures may be taken.

The controller 602 also includes two or more disparate network interfaces 706, namely, there is at least one interface 706 per network to which the controller 602 controls access. Each interface 706 can be implemented as a direct interface

12

706 or as an indirect interface 706; a given embodiment may comprise only direct interfaces 706, may comprise only indirect interfaces 706, or may comprise at least one of each type of interface.

An indirect interface 706 may be implemented, for instance, as a direct frame relay connection over land line or wireless or network interfaces to which the frame relay routers can connect, or as a point-to-point interface to a dedicated T1, T3, or wireless connection. One suitable implementation includes multiple standard Ethernet cards, in the controller 602 and in the router, which connect to each other. An external frame relay User-Network Interface (UNI) resides in a router 105 of a network 106; a similar Ethernet card resides in the Internet router 104. Each such Ethernet card will then have a specific IP address assigned to it. The controller can also have a single Ethernet card with multiple IP addresses associated with different routers and LANs. An indirect interface 706 may connect to the network over fiber optic, T1, wireless, or other links.

A direct interface 706 comprises a standard connection to the Internet 500, while another direct interface 706 comprises a standard connection to a VPN. One direct interface 706 effectively makes part of the controller 602 into a UNI by including in the interface 706 the same kind of special purpose hardware and software that is found on the frame relay network side (as opposed to the UNI side) of a frame relay network router. Such a direct frame relay network interface 706 is tailored to the specific timing and other requirements of the frame relay network to which the direct interface 706 connects. For instance, one direct interface 706 may be tailored to a Qwest frame relay network 106, while another direct interface 706 in the same controller 602 is tailored to a UUNet network 106. Another direct interface 706 comprises standard VPN components.

An indirect interface 706 relies on special purpose hardware and connectivity/driver software in a router or other device, to which the indirect interface 706 of the controller 602 connects. By contrast, a direct interface 706 includes such special purpose hardware and connectivity/driver software inside the controller 602 itself. In either case, the controller provides packet switching capabilities for at least redundancy without manual switchover, and preferably for dynamic load-balancing between lines as well. FIG. 7 shows three interfaces 706; other controllers may have a different number of interfaces. The three interfaces 706 (for instance) may be implemented using a single card with three IP addresses, or three cards, each with one IP address. The site interface 702 may or may not be on the same card as interface(s) 706. The controller 602 in each case also optionally includes memory buffers in the site interface 702, in the path selector 704, and/or in the network interfaces 706.

An understanding of methods of the invention will follow from understanding the invention's devices, and vice versa. For instance, from FIGS. 6, 7, 10, and 11 one may ascertain methods of the invention for combining connections for access to multiple disparate networks, as well as systems and devices of the invention. As illustrated in FIG. 8, methods of the invention use a device such as controller 602. The controller 602 comprises (a) a site network interface 702, (b) at least two WAN network interfaces 706 tailored as necessary to particular networks, and (c) a packet path selector 704 which selects between network interfaces 706 according to a specified criterion. Path selection criteria may be specified 800 by configuration files, hardware jacks or switches, ROM values, remote network management tools, or other means. Variations in topology are also possible, e.g., in a variation on FIG. 10 the VPNs could swap position with their respective routers.

13

One then connects the site interface **702** to a site **102** to receive packets from a computer (possibly via a LAN) at the site **102**. Likewise, one connects a first network interface **706** to a first router for routing packets to a first network, and a second network interface **706** to a second router for routing packets to a second network, with the networks being disparate to each other. A third, fourth, etc. network may be similarly connected to the controller **602** in some embodiments and/or situations.

The connected disparate networks are parallel to one another (not serial, although additional networks not directly connected to the controller **602** may be serially connected to the parallel disparate networks). The connected disparate networks are independent of one another, in that no routing information need be shared between them, to make them parallel (NNIs can still be used to connect networks in serial to form a larger independent and parallel network). A mistake in the routing information for one network will thus not affect the other network.

After the connections are made (which may be done in a different order than recited here), one sends **802** a packet to the site interface **702**. The controller **602** then sends the packet through the one (or more—copies can be sent through multiple networks) network interface **706** that was selected by the packet path selector **704**. The packet path selector **704** can maintain a table of active sessions, and use that table to select a path for packets in a given session. The packet path selector **704** does not need a session table to select paths for site-to-site traffic, because the controller **602** on the other site knows where to forward the site-to-site packets.

FIG. 9 is a flowchart further illustrating methods of the present invention, which send packets over multiple parallel independent disparate networks for enhanced reliability, load balancing and/or security; frame relay networks and the Internet are used as an example, but point-to-point networks and VPNs may be similarly employed according to the discussion herein.

During an address range information obtaining step **900**, address ranges for known locations are obtained. Address ranges may be specified as partial addresses, e.g., partial IP addresses in which some but not all of the address is specified. Thus, "198.x.x.x" indicates an IP address in which the first field is 198 and the other three address fields are not pinned down, corresponding to the range of addresses from 198.0.0.0 to 198.255.255.255. Each address range has an associated network; a network may have more than one associated contiguous range of addresses which collectively constitute the address range for that network. The locations reachable through the network have addresses in the address range associate with the network. Since part of the address specifies the network, a location reachable through two networks has two addresses, which differ in their network-identifying bits but are typically the same in their other bits. Address ranges may be obtained **900** by reading a configuration file, querying routers, receiving input from a network administrator, and/or other data gathering means.

During a topology information obtaining step **902**, topology information for the system of parallel disparate networks is obtained. The topology information specifies which one or more networks can be used (if functioning) to reach which known locations. With regard to FIG. 6, for instance, the topology information could be represented by a table, list, or other data structure which specifies that: the VPN connects sites A and B; the Internet connects sites A and B; and the private (frame relay/point-to-point) network connects sites A, B, and C. Topology information may be

14

obtained **902** by reading a configuration file, querying routers, receiving input from a network administrator, and/or other data gathering means.

If necessary, a connection forming step is performed, e.g., to obtain a virtual circuit between two sites **102**. The controller **602** then checks the status of each connection and updates the information for available communication paths.

The controller **602** at each location will go through the address range information obtaining step, topology information obtaining step and connection forming step. More generally, the steps illustrated and discussed in this document may be performed in various orders, including concurrently, except in those cases in which the results of one step are required as input to another step. Likewise, steps may be omitted unless required by the claims, regardless of whether they are expressly described as optional in this Detailed Description. Steps may also be repeated, or combined, or named differently.

During a packet receiving step **904**, the controller **602** at a given source location receives a packet to be sent from that location to the destination site **102**. In some cases, multiple packets may be received in a burst. The packet comes into the controller **602** through the site interface **702**.

During a determining step **906**, the controller **602** (or some other device used in implementing the method) looks at the packet destination address to determine whether the destination address lies within a known address range. That is, the destination address is compared to the known location address ranges that were obtained during step **900**, in order to see whether the destination location is a known location. Only packets destined for known locations are potentially rerouted by the invention to balance loads, improve security, and/or improve reliability. Packets destined for unknown locations are simply sent to the network indicated in their respective destination addresses, which is the Internet **500** in the examples given herein but could also be some other "catch-all" network. Although they are not rerouted, such packets may nonetheless be counted as part of the load balancing calculation.

During a path selecting step **908**, the path selector **704** selects the path over which the packet will be sent; selection is made between at least two paths, each of which goes over a different network **106** than the other. The disparate networks are independent parallel networks. This path selecting step **908** may be performed once per packet, or a given selection may pertain to multiple packets. In some embodiments, selecting a network will also select a path, as in the system shown in FIG. 10. In other cases, there may be more than one path to a given network, as discussed in connection with the line pairs shown in FIG. 6. Packet path selection **908** is shown as following packet receipt **904**, but in some embodiments and/or some situations, it may precede packet receipt **904**. That is, in some cases the path for the next packet may be determined by the packet path selector before the packet arrives, e.g., in a round-robin manner, while in other cases the path is determined after the packet arrives, e.g., using per-packet dynamic load balancing.

As indicated, the path selection may use **910** load balancing as a criterion for selecting a path, use **912** network status (up/down) and other connectivity criteria (e.g., router status, connectivity status) as a criterion for selecting a path, and/or use **914** division of packets between disparate networks for enhanced security as a criterion for selecting a path. These steps may be implemented in a manner consistent with the description above of the path selector **704** given in the discussion of FIG. 7. More generally, unless it is

15

otherwise indicated, the description herein of systems of the present invention extends to corresponding methods, and vice versa.

The description of systems and methods likewise extend to corresponding computer-readable media (e.g., RAM, ROM, other memory chips, disks, tape, Iomega ZIP or other removable media, and the like) which are configured by virtue of containing software to perform an inventive method, or software (including any data structure) which is uniquely suited to facilitate performance of an inventive method. Articles of manufacture within the scope of the present invention thus include a computer-readable storage medium in combination with the specific physical configuration of a substrate of the computer-readable storage medium, when that substrate configuration represents data and/or instructions which cause one or more computers to operate in a specific and predefined manner as described and claimed herein.

No change to packet source IP address or destination IP address need by done by the controller in a topology like that shown in FIG. 10. The controller 602 sends the packet to router X or router Y as determined by the packet path selector. This is illustrated in the following summary example:

Packet location	Packet Source IP Address	Packet Destination IP Address
Leaving site A	Site A's IP address	Site B's IP address
Leaving controller A	Site A's IP address	Site B's IP address
Leaving VPN/Router	VPN/Router/Site A	VPN/Router/Site B
<packet travels over Internet/frame relay net/etc.>		
Arrival VPN/Router	VPN/Router/Site A	VPN/Router/Site B
Arrival controller B	Site A	Site B
<controller may need to resequence packets>		
Arrival at site B	Site A	Site B

However, packet addresses are modified during operation of a configuration like that shown in FIG. 11. An example is provided in the following summary example:

Packet Location	Packet Source IP Address	Packet Destination IP Address
Leaving site A	Site A's IP address	Site B's IP address
Leaving VPN A	VPN A's IP address	VPN B's IP address
Leaving controller A	A controller A IP address	A controller B IP address
<packet travels over Internet/frame relay net/etc.>		
Arrival controller B	The controller A IP address	The controller B IP address
<controller may need to resequence packets>		
Arrival at VPN B	VPN A's IP address	VPN B's IP address
<note that the controllers are transparent to the VPNs>		
Arrival at site B	Site A	Site B
<the VPNs are transparent to the sites>		

During an address modifying step 916, the packet destination address is modified as needed to make it lie within an address range (obtained during step 900) which is associated with the selected path to the selected network (selected during step 908). For instance, if a packet is received 904 with a destination address corresponding to travel through the Internet but the path selection 908 selects a path for the

16

packet through a frame relay network 106 to the same destination, then the packet's destination IP address is modified 916 by replacing the IP address with the IP address of the appropriate interface of the controller at Site B. Also the packet's source IP address is replaced with the IP address of the appropriate interface of the source controller. This modifying step may be viewed as optional, in the sense that it need not be performed in every embodiment. But it is required in the sense that a system embodiment of the invention which is claimed with a limitation directed to destination address modification must be at least capable of performing the modifying step, and a method embodiment which is claimed with a limitation directed to the modifying step must perform the modifying step on at least one packet.

With regard to both FIG. 10 and FIG. 11, during a packet transmission step 918, the packet is sent on the selected 908 path. This is done by sending the packet over the network interface 706 for the path selected. As indicated in FIG. 9, the method may then loop back to receive 904 the next packet, select 908 a network for that packet, send 918 it, and so on. As noted, other specific method instances are also possible. One example is the inventive method in which load balancing or reliability criteria cause an initial path selection to be made 908, and then a loop occurs in which multiple packets are received 904 and then sent 918 over the selected path without repeating the selecting step 908 for each receive 904—send 918 pair. Note that some embodiments of the invention permit packets of a given message to be sent over two or more disparate networks, thereby enhancing 914 security. An ending step may be performed as needed during an orderly shutdown for diagnostic or upgrade work, for instance.

The controller 602 at the destination site goes through the steps described above in reverse order as needed. The controller 602 receives the packet from the source location through one of the network interfaces. Packet resequencing may be needed in either the FIG. 10 or the FIG. 11 configuration, while address changes are needed in the FIG. 11 configuration only.

Conclusion

The present invention provides methods and devices for placing frame relay and other private networks in parallel with VPNs and other Internet-based networks, thereby providing redundancy without requiring manual switchover in the event of a network failure. Load-balancing between lines and/or between networks may also be performed. For instance, the invention can be used to provide reliable, efficient, and secure point-to-point connections for private networks 106 in parallel with a VPN and an SSL Internet connection. Some prior art approaches require network reconfiguration each time a frame relay circuit fails, and some have complex router configurations to handle load balancing and network failures. This requires substantial effort by individual network customers to maintain connectivity, and they will often receive little or no help from the frame relay carriers, or not receive prompt service from a VPN provider. Instead, well-trained staff are needed at each location, as are expensive routers. By contrast, these requirements are not imposed by the present invention.

As used herein, terms such as "a" and "the" and item designations such as "connection" or "network" are generally inclusive of one or more of the indicated item. In particular, in the claims a reference to an item normally means at least one such item is required.

The invention may be embodied in other specific forms without departing from its essential characteristics. The described embodiments are to be considered in all respects

17

only as illustrative and not restrictive. Headings are for convenience only. The claims form part of the specification. The scope of the invention is indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed and desired to be secured by patent is:

1. A controller which controls access to multiple independent disparate networks in a parallel network configuration, the disparate networks comprising at least one private network and at least one network based on the Internet, the controller comprising:

- a sit interface connecting the controller to a site;
- at least two network interfaces which send packets toward the disparate networks; and
- a packet path selector which selects between network interfaces according to at least:
 - a destination of the packet, an optional presence of alternate paths to that destination, and at least one specified criterion for selecting between alternate paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector; and

wherein the packet path selector selects between network interfaces according to a security criterion, thereby promoting use of multiple disparate networks to carry different pieces of a given message so that unauthorized interception of packets on fewer than all of the disparate networks used to carry the message will not provide the total content of the message.

2. The controller of claim 1, wherein the controller sends packets out of sequence over the parallel disparate networks.

3. The controller of claim 2, wherein the controller places an encrypted sequence number in at least some of the packets which are sent out of sequence.

4. A controller which controls access to multiple networks in a parallel network configuration, suitable networks comprising Internet-based networks and private networks from at least one more provider, in combination, the controller comprising:

- a site interface connecting the controller to a site;
- at least two network interfaces which send packets toward the networks; and
- a packet path selector which selects between network interfaces on a per-packet basis according to at least: a destination of the packet, an optional presence of alternate paths to that destination, and at least one specified criterion for selecting between alternate paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector.

5. A method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of:

- obtaining at least two known location address ranges which have associated networks;
- obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location;
- receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location;

18

determining whether the destination address lies within a known location address range;

selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location;

forwarding the packet on the selected network path.

6. The method of claim 5, further comprising the step of modifying the packet destination address to lie within a known location address range associated with the selected network before the forwarding step.

7. The method of claim 5, wherein the forwarding step forwards the packet toward the Internet when the packet's destination address does not lie within any known location address range.

8. The method of claim 5, wherein the destination address identifies a destination location to which only a single associated network provides connectivity from the current location, and the forwarding step forwards the packet to that single associated network.

9. The method of claim 5, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

10. The method of claim 5, wherein repeated instances of the selecting step make network path selections on a per session basis.

11. The method of claim 5, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.

12. The method of claim 11, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance line loads by distributing packets between lines.

13. The method of claim 11, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance network loads by distributing packets between disparate networks.

14. The method of claim 5, wherein the selecting step selects the network path at least in part on the basis of a reliability criterion.

15. The method of claim 5, wherein the selecting step selects the network path at least in part on the basis of a security criterion.

16. The method of claim 6, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a private network such that the modified packet destination address lies instead in a known location address range associated with a VPN.

17. The method of claim 6, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a VPN such that the modified packet destination address lies instead in a known location address range associated with a private network.

18. The method of claim 6, wherein the modifying step modifies a packet destination address corresponding to one of: the Internet, a private network, thereby making the modified packet destination address correspond to the other of: the Internet, a private network.

19. A method for combining connections for access to parallel networks, the method comprising the steps of: sending a packet to a site interface of a controller, the controller comprising the site interface which receives packets, at least two network interfaces to parallel

19

networks, and a packet path selector which selects between the network interfaces on a per-session basis to promote load-balancing; and
 forwarding the packet through the network interface selected by the packet path selector;
 wherein the step of sending a packet to the controller site interface is repeated as multiple packets are sent, and the controller sends different packets of a given message to different parallel networks.
 20. A method for combining connections for access to parallel networks, the method comprising the steps of:
 receiving at a first controller a packet which has a first site IP address as source address and a second site IP address as destination address;
 modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address; and
 forwarding the modified packet along a selected path toward the second site.
 21. A method for combining connections for access to parallel networks, the method comprising the steps of:
 receiving at a first controller a packet which has a first VPN IP address as source address and a second VPN IP address as destination address;
 modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address; and
 forwarding the modified packet along a selected path toward the second VPN.
 22. A computer storage medium having a configuration that represents data and instructions which will cause per-

20

formance of a method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of:
 obtaining at least two known location address ranges which have associated networks;
 obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location;
 receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location;
 determining whether the destination address lies within a known location address range; selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location;
 modifying the packet destination address to lie within a known location address range associated with the selected network if it does not already do so; and
 forwarding the packet on the selected network path.
 23. The configured storage medium of claim 22, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.
 24. The configured storage medium of claim 22, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

* * * * *



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

BIBDATASHEET

CONFIRMATION NO. 3845

Bib Data Sheet

SERIAL NUMBER 10/361,837	FILING DATE 02/07/2003 RULE	CLASS 370	GROUP ART UNIT 2663	ATTORNEY DOCKET NO. 3003.2.11A
-----------------------------	-----------------------------------	--------------	------------------------	--------------------------------------

APPLICANTS

Sanchaita Datta, Salt Lake City, UT;
Ragula Bhaskar, Salt Lake City, UT;

** CONTINUING DATA ^{MM} *****
This application is a CIP of 10/034,197 12/28/2001
which claims benefit of 60/259,269 12/29/2000
This application 10/361,837
claims benefit of 60/255,509 02/08/2002

** FOREIGN APPLICATIONS ^{None MM} *****

IF REQUIRED, FOREIGN FILING LICENSE GRANTED ** SMALL ENTITY **
** 03/31/2007

Foreign Priority claimed <input type="checkbox"/> yes <input checked="" type="checkbox"/> no	35 USC 119 (a-d) conditions met <input type="checkbox"/> yes <input checked="" type="checkbox"/> no <input type="checkbox"/> Met after Allowance	STATE OR COUNTRY UT	SHEETS DRAWING 6	TOTAL CLAIMS 24 35	INDEPENDENT CLAIMS 7 8
Verified and Acknowledged <i>[Signature]</i> 05.25.2004 Examiner's Signature Initials					

ADDRESS
23484
JOHN W. L. OGILVIE
1320 EAST LAIRD AVENUE
SALT LAKE CITY, UT
84105

TITLE
Tools and techniques for directing packets over disparate networks

FILING FEE RECEIVED 827	FEES: Authority has been given in Paper No. _____ to charge/credit DEPOSIT ACCOUNT No. _____ for following:	<input type="checkbox"/> All Fees <input type="checkbox"/> 1.16 Fees (Filing) <input type="checkbox"/> 1.17 Fees (Processing Ext. of time) <input type="checkbox"/> 1.18 Fees (Issue) <input type="checkbox"/> Other <input type="checkbox"/> Credit
----------------------------	---	---

PATENT APPLICATION SERIAL NO. 10 361 837

U.S. DEPARTMENT OF COMMERCE
PATENT AND TRADEMARK OFFICE
FEE RECORD SHEET

PTO-1556
(5/87)

U.S. GPO: 1986-433-214/80404

2-11-03


10361837-020703

COMPUTER LAW++[®]

Software patents, copyrights, trademarks, licenses and related legal services

John W.L. Ogilvie
Registered Patent Attorney
M.S. Computer Science
jwlo@LawPlusPlus.com

1211 East Yale Avenue
Salt Lake City, Utah 84105
Voice: 801-582-2724
Fax: 801-583-1984
www.LawPlusPlus.com



Express Mail Label No. EV047149870US

PATENT APPLICATION
Docket No. 3003.2.11A

February 7, 2003

Commissioner for Patents
Box Patent Application
P.O. Box 2327
Arlington, VA 22202

Commissioner:

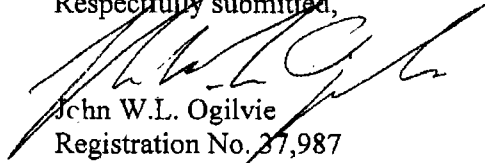
Filed herewith is an application for letters patent for TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS, in the name of inventors Sanchaita Datta and Ragula Bhaskar, comprising a title page, 44 pages of specification and claims, and 6 sheets of drawings. The following are also enclosed:

- An Application Data Sheet pursuant to 37 C.F.R. § 1.76;
 - An Inventors' Declaration (including a Power of Attorney);
 - An Assignment to Ragula Systems, and recordation cover sheet;
 - A Certificate of Mailing by Express Mail and self-addressed stamped postcard.
- Please address all future communications to the undersigned.

10361837,020703

Assistant Commissioner for Patents
February 7, 2003
Page 2

Respectfully submitted,



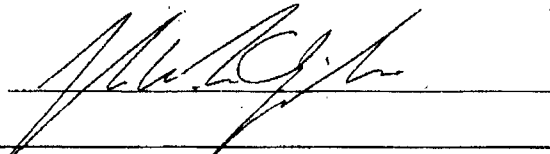
John W.L. Ogilvie
Registration No. 37,987
COMPUTER LAW++
1211 East Yale Avenue
Salt Lake City, Utah 84105
801-582-2724 (voice)
801-583-1984 (fax)

\pxmit11A

CERTIFICATE OF MAILING UNDER 37 CFR 1.10

I hereby certify that the correspondence listed below is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on February 7, 2003 addressed to the Commissioner for Patents, Box Patent Application, P.O. Box 2327, Arlington, VA 22202:

Certificate of Mailing, Postcard
Transmittal Letter
Application Data Sheet (3 pages)
Inventors' Declaration
Assignment w/ Recordation Cover Sheet
Patent Application including title, 44 pages of specification and claims, and 6 drawing sheets



EV047149870US
"Express Mail" label number

1036149870US

Express Mail Label No. EV047149870US
PATENT APPLICATION
DOCKET NO. 3003.2.11A

UNITED STATES
PATENT APPLICATION

OF

SANCHAITA DATTA AND RAGULA BHASKAR

FOR

TOOLS AND TECHNIQUES FOR
DIRECTING PACKETS OVER DISPARATE NETWORKS

**TOOLS AND TECHNIQUES FOR
DIRECTING PACKETS OVER DISPARATE NETWORKS**

5

RELATED APPLICATIONS

This application claims priority to commonly owned copending U.S. provisional patent application serial no. 60/355,509 filed February 8, 2002, which is also incorporated herein by reference. This application is a continuation-in-part of U.S. patent application serial no. 10/034,197 filed December 28, 2001, which claims priority to U.S. provisional patent application serial no. 60/259,269 filed December 29, 2000, each of which is also incorporated herein by reference.

10

FIELD OF THE INVENTION

The present invention relates to computer network data transmission, and more particularly relates to tools and techniques for communications using disparate parallel networks, such as a virtual private network ("VPN") or the Internet in parallel with a point-to-point, leased line, or frame relay network, in order to help provide benefits such as load balancing across network connections, greater reliability, and increased security.

15

20

TECHNICAL BACKGROUND OF THE INVENTION

Organizations have used frame relay networks and point-to-point leased line networks for interconnecting geographically dispersed offices or locations. These networks have been implemented in the past and are currently in use for interoffice communication, data exchange and file sharing. Such networks have advantages, some of which are noted below. But these networks also tend to be expensive, and there are

25

relatively few options for reliability and redundancy. As networked data communication becomes critical to the day-to-day operation and functioning of an organization, the need for lower cost alternatives for redundant back-up for wide area networks becomes important.

5 Frame relay networking technology offers relatively high throughput and reliability. Data is sent in variable length frames, which are a type of packet. Each frame has an address that the frame relay network uses to determine the frame's destination. The frames travel to their destination through a series of switches in the frame relay network, which is sometimes called a network "cloud"; frame relay is an example of
10 packet-switched networking technology. The transmission lines in the frame relay cloud must be essentially error-free for frame relay to perform well, although error handling by other mechanisms at the data source and destination can compensate to some extent for lower line reliability. Frame relay and/or point-to-point network services are provided or have been provided by various carriers, such as AT&T, Qwest, XO, and MCI WorldCom.

15 Frame relay networks are an example of a network that is "disparate" from the Internet and from Internet-based virtual private networks for purposes of the present invention. Another example of such a "disparate" network is a point-to-point network, such as a T1 or T3 connection. Although the underlying technologies differ somewhat, for purposes of the present invention frame relay networks and point-to-point networks
20 are generally equivalent in important ways, such as the conventional reliance on manual switchovers when traffic must be redirected after a connection fails, and their implementation distinct from the Internet. A frame relay permanent virtual circuit is a virtual point-to-point connection. Frame relays are used as examples throughout this

document, but the teachings will also be understood in the context of point-to-point networks.

A frame relay or point-to-point network may become suddenly unavailable for use. For instance, both MCI WorldCom and AT&T users have lost access to their
5 respective frame relay networks during major outages. During each outage, the entire network failed. Loss of a particular line or node in a network is relatively easy to work around. But loss of an entire network creates much larger problems.

Tools and techniques to permit continued data transmission after loss of an entire frame relay network that would normally carry data are discussed in United States Patent
10 Application No. 10/034,197 filed December 28, 2001 and incorporated herein. The '197 application focuses on architectures involving two or more "private" networks in parallel, whereas the present application focuses on architectures involving disparate networks in parallel, such as a proprietary frame relay network and the Internet. Note that the term "private network" is used herein in a manner consistent with its use in the '197 applica-
15 tion (which comprises frame relay and point-to-point networks), except that a "virtual private network" as discussed herein is not a "private network". Virtual private networks are Internet-based, and hence disparate from private networks, i.e., from frame relay and point-to-point networks. To reduce the risk of confusion that might arise from misunderstanding "private network" to comprise "virtual private network" herein, virtual private
20 networks will be henceforth referred to as VPNs. Other differences and similarities between the present application and the '197 application will also be apparent to those of skill in the art on reading the two applications.

Various architectures involving multiple networks are known in the art. For instance, Figure 1 illustrates prior art configurations involving two frame relay networks for increased reliability; similar configurations involve one or more point-to-point network connections. Two sites 102 transmit data to each other (alternately, one site might be only a data source, while the other is only a data destination). Each site has two border routers 105. Two frame relay networks 106, 108 are available to the sites 102 through the routers 105. The two frame relay networks 106, 108 have been given separate numbers in the figure, even though each is a frame relay network, to emphasize the incompatibility of frame relay networks provided by different carriers. An AT&T frame relay network, for instance, is incompatible – in details such as maximum frame size or switching capacity – with an MCI WorldCom frame relay network, even though they are similar when one takes the broader view that encompasses disparate networks like those discussed herein. The two frame relay providers have to agree upon information rates, switching capacities, frame sizes, etc. before the two networks can communicate directly with each other.

A configuration like that shown in Figure 1 may be actively and routinely using both frame relay networks A and B. For instance, a local area network (LAN) at site 1 may be set up to send all traffic from the accounting and sales departments to router A1 and send all traffic from the engineering department to router B1. This may provide a very rough balance of the traffic load between the routers, but it does not attempt to balance router loads dynamically in response to actual traffic and thus is not “load-balancing” as that term is used herein.

Alternatively, one of the frame relay networks may be a backup which is used only when the other frame relay network becomes unavailable. In that case, it may take even skilled network administrators several hours to perform the steps needed to switch the traffic away from the failed network and onto the backup network, unless the invention of the '197 application is used. In general, the necessary Private Virtual Circuits (PVCs) must be established, routers at each site 102 must be reconfigured to use the correct serial links and PVCs, and LANs at each site 102 must be reconfigured to point at the correct router as the default gateway.

Although two private networks are shown in Figure 1, three or more such networks could be employed, with similar considerations coming into play as to increased reliability, limits on load-balancing, the efforts needed to switch traffic when a network fails, and so on. Likewise, for clarity of illustration Figure 1 shows only two sites, but three or more sites could communicate through one or more private networks.

Figure 2 illustrates a prior art configuration in which data is normally sent between sites 102 over a private network 106. A failover box 202 at each site 102 can detect failure of the network 106 and, in response to such a failure, will send the data instead over an ISDN link 204 while the network 106 is down. Using an ISDN link 204 as a backup is relatively easier and less expensive than using another private network 106 as the backup, but generally provides lower throughput. The ISDN link is an example of a point-to-point or leased line network link.

Figure 3 illustrates prior art configurations involving two private networks for increased reliability, in the sense that some of the sites in a given government agency or other entity 302 can continue communicating even after one network goes down. For

10361637 020703

instance, if a frame relay network A goes down, sites 1, 2, and 3 will be unable to communicate with each other but sites 4, 5, and 6 will still be able to communicate amongst themselves through frame relay network B. Likewise, if network B goes down, sites 1, 2, and 3 will still be able to communicate through network A. Only if both
5 networks go down at the same time would all sites be completely cut off. Like the Figure 1 configurations, the Figure 3 configuration uses two private networks. Unlike Figure 1, however, there is no option for switching traffic to another private network when one network goes down, although either or both of the networks in Figure 3 could have an ISDN backup like that shown in Figure 2. Note also that even when both private
10 networks are up, sites 1, 2, and 3 communicate only among themselves; they are not connected to sites 4, 5, and 6. Networks A and B in Figure 3 are therefore not in "parallel" as that term is used herein, because all the traffic between each pair of sites goes through at most one of the networks A, B.

Figure 4 illustrates a prior art response to the incompatibility of frame relay
15 networks of different carriers. A special "network-to-network interface" (NNI) is used to reliably transmit data between the two frame relay networks A and B. NNIs are generally implemented in software at carrier offices. Note that the configuration in Figure 4 does not provide additional reliability by using two frame relay networks, because those networks are in series rather than in parallel. If either of the frame relay
20 networks A, B in the Figure 4 configuration fails, there is no path between site 1 and site 2; adding the second frame relay network has not increased reliability. By contrast, Figure 1 increases reliability by placing the frame relay networks in parallel, so that an alternate path is available if either (but not both) of the frame relay networks fails.

Someone of skill in the art who was looking for ways to improve reliability by putting networks in parallel would probably not consider NNIs pertinent, because they were used for serial configurations rather than parallel ones, and adding networks in a serial manner does not improve reliability.

5 Internet-based communication solutions such as VPNs and Secure Sockets Layer (SSL) offer alternatives to frame relay 106 and point-to-point leased line networks such as those using an ISDN link 204. These Internet-based solutions are advantageous in the flexibility and choice they offer in cost, in service providers, and in vendors. Accordingly, some organizations have a frame relay 106 or leased line connection (a.k.a.
10 point-to-point) for intranet communication and also have a connection for accessing the Internet 500, using an architecture such as that shown in Figure 5.

But better tools and techniques are needed for use in architectures such as that shown in Figure 5. In particular, prior approaches for selecting which network to use for which packet(s) are coarse. For instance, all packets from department X might be sent
15 over the frame relay connection 106 while all packets from department Y are sent over the Internet 500. Or the architecture might send all traffic over the frame relay network unless that network fails, and then be manually reconfigured to send all traffic over a VPN 502.

Organizations are still looking for better ways to use Internet-based redundant
20 connections to backup the primary frame relay networks. Also, organizations wanting to change from frame relay and point-to-point solutions to Internet-based solutions have not had the option of transitioning in a staged manner. They have had to decide instead between the two solutions, and deploy the solution in their entire network communica-

tions system in one step. This is a barrier for deployment of Internet-based solutions 500/502, since an existing working network would be replaced by a yet-untested new network. Also, for organizations with several geographically distributed locations a single step conversion is very complex. Some organizations may want a redundant 5 Internet-based backup between a few locations while maintaining the frame relay network for the entire organization.

It would be an advancement in the art to provide new tools and techniques for configuring disparate networks (e.g., frame relay/point-to-point WANs and Internet-based VPNs) in parallel, to obtain benefits such as greater reliability, improved security, and/or 10 load-balancing. Such improvements are disclosed and claimed herein.

BRIEF SUMMARY OF THE INVENTION

The present invention provides tools and techniques for directing packets over multiple parallel disparate networks, based on addresses and other criteria. This helps 15 organizations make better use of frame relay networks and/or point-to-point (e.g., T1, T3, fiber, OCx, Gigabit, wireless, or satellite based) network connections in parallel with VPNs and/or other Internet-based networks. For instance, some embodiments of the invention allow frame relay and VPN wide area networks to co-exist for redundancy as well as for transitioning from frame relay/point-to-point solutions to Internet-based 20 solutions in a staged manner. Some embodiments operate in configurations which communicate data packets over two or more disparate WAN connections, with the data traffic being dynamically load-balanced across the connections, while some embodiments

treat one of the WANs as a backup for use mainly in case the primary connection through the other WAN fails.

Other features and advantages of the invention will become more fully apparent through the following description.

5

BRIEF DESCRIPTION OF THE DRAWINGS

To illustrate the manner in which the advantages and features of the invention are obtained, a more particular description of the invention will be given with reference to the attached drawings. These drawings only illustrate selected aspects of the invention and its context. In the drawings:

10

Figure 1 is a diagram illustrating a prior art approach having frame relay networks configured in parallel for increased reliability for all networked sites, in configurations that employ manual switchover between the two frame relay networks in case of failure.

15

Figure 2 is a diagram illustrating a prior art approach having a frame relay network configured in parallel with an ISDN network link for increased reliability for all networked sites.

20

Figure 3 is a diagram illustrating a prior art approach having independent and non-parallel frame relay networks, with each network connecting several sites but no routine or extensive communication between the networks.

Figure 4 is a diagram illustrating a prior art approach having frame relay networks configured in series through a network-to-network interface, with no consequent increase in reliability because the networks are in series rather than in parallel.

Figure 5 is a diagram illustrating a prior art approach having a frame relay network configured in parallel with a VPN or other Internet-based network that is disparate to the frame relay network, but without the fine-grained packet routing of the present invention.

5 Figure 6 is a diagram illustrating one system configuration of the present invention, in which the Internet and a private network are placed in parallel for increased reliability for all networked sites, without requiring manual traffic switchover, and with the option in some embodiments of load balancing between the networks and/or increasing security by transmitting packets of a single logical connection over disparate
10 networks.

Figure 7 is a diagram further illustrating a multiple disparate network access controller of the present invention, which comprises an interface component for each network to which the controller connects, and a path selector in the controller which uses one or more of the following as criteria: destination address, network status (up/down),
15 network load, use of a particular network for previous packets in a given logical connection or session.

Figure 8 is a flowchart illustrating methods of the present invention for sending packets using a controller such as the one shown in Figure 7.

Figure 9 is a flowchart illustrating methods of the present invention for combining
20 connections to send traffic over multiple parallel independent disparate networks for reasons such as enhanced reliability, load balancing, and/or security.

Figure 10 is a diagram illustrating another system configuration of the present invention, in which the Internet and a frame relay network are placed in parallel, with a

VPN tunnel originating after the source controller and terminating before the destination controller, and each known site that is accessible through one network is also accessible through the other network unless that other network fails.

Figure 11 is a diagram illustrating a system configuration similar to Figure 10, except the VPN tunnel originates before the source controller and terminates after the destination controller.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention relates to methods, systems, and configured storage media for connecting sites over multiple independent parallel disparate networks, such as frame relay networks and/or point-to-point network connections, on the one hand, and VPNs or other Internet-based network connections, on the other hand. "Multiple" networks means two or more such networks. "Independent" means routing information need not be shared between the networks. "Parallel" does not rule out all use of NNIs and serial networks, but it does require that at least two of the networks in the configuration be in parallel at the location where the invention distributes traffic, so that alternate data paths through different networks are present. "Frame relay networks" or "private networks" does not rule out the use of an ISDN link or other backup for a particular frame relay or point-to-point private network, but it does require the presence of multiple such networks; Figure 2, for instance, does not meet this requirement. A "frame relay network" is unavailable to the general public and thus disparate from the Internet and VPNs (which may be Internet-based), even though some traffic in the Internet may use public frame relay networks once the traffic leaves the location where the invention distributes traffic.

Figure 6 illustrates one of many possible configurations of the present invention. Comments made here also apply to similar configurations involving only one or more frame relay networks 106, those involving only one or more point-to-point networks 204, and those not involving a VPN 604, for example. Two or more disparate networks are placed in parallel between two or more sites 102. In the illustrated configuration, the Internet 500 and a VPN 604 are disparate from, and in parallel with, frame relay / point-to-point network 106/204, with respect to site A and site B. No networks are parallel disparate networks in Figure 6 with regard to site C as a traffic source, since that site is not connected to the Internet 500. Access to the disparate networks at site A and and site B is through an inventive controller 602 at each site. Additional controllers 602 may be used at each location (i.e., controllers 602 may be placed in parallel to one another) in order to provide a switched connection system with no single point of failure.

With continued attention to the illustrative network topology for one embodiment of the invention shown in Figure 6, in this topology the three locations A, B, and C are connected to each other via a frame relay 106 or leased line network 204. Assume, for example, that all three locations are connected via a single frame relay network 106. Locations A and B are also connected to each other via a VPN connection 604. VPN tunnels are established between locations A and B in the VPN, which pairs line 1 to line 3 and also pairs line 2 to line 3. There can be only one VPN tunnel between locations A and B. There is no VPN connection between location C and either location A or location B.

Therefore, locations A, B, and C can communicate with each other over the frame relay network 106, and locations A and B (but not C) can also communicate with each

other over the VPN connection 604. Communication between locations A and C, and communication between locations B and C, can take place over the frame relay network 106 only. Communication between locations A and B can take place over frame relay network 106. It can also take place over one of the lines 1-and-3 pair, or the lines 2-and-3 pair, but not both at the same time. Traffic can also travel over lines 2 and 4, but without a VPN tunnel. When the source and destination IP address pairs are the same between locations A and B but different types of networks connect those locations, as in Figure 6 for instance, then a traffic routing decision that selects between network types cannot be made with an existing commercially available device. By contrast, the invention allows an organization to deploy an Internet-based solution between locations A and B while maintaining the frame relay network 106 between locations A, B, and C, and allows traffic routing that selects between the Internet and the frame relay network on a packet-by-packet basis.

The invention may thus be configured to allow the organization to achieve the following goals, in the context of Figure 6; similar goals are facilitated in other configurations. First, the organization can deploy an Internet-based second connection between only locations A and B, while maintaining frame relay connectivity between locations A, B, and C. Later the organization may deploy an Internet-based solution at location C as well. Second, the organization can use the Internet-based connection between locations A and B for full load-balancing or backup, or a combination of the two. Third, the organization can use the frame relay connection between locations A and B for full load-balancing or backup, or a combination of the two. Fourth, the organization can

load-balance traffic in a multi-homing situation between two ISPs or two connections to the Internet at locations A and/or B.

To better understand the invention, consider the operation of controller device 602 at location A. The controller 602 examines the IP data traffic meant to go through it and
5 makes determinations and takes steps such as those discussed below.

If the traffic is destined for the Internet 500, send the traffic over the Internet using lines 1 and/or 2. Load balancing decisions that guide the controller 602 in distributing packets between the lines can be based on criteria such as the load of a given network, router, or connection relative to other networks, routers, or connections, to be performed
10 dynamically in response to actual traffic. Load-balancing may be done through a round-robin algorithm which places the next TCP or UDP session on the next available line, or it may involve more complex algorithms that attempt to measure and track the throughput, latency, and/or other performance characteristics of a given link or path element. Load-balancing is preferably done on a per-packet basis for site-to-site data
15 traffic over the Internet or frame relay net, or done on a TCP or UDP session basis for Internet traffic, as opposed to prior approaches that use a per-department and/or per-router basis for dividing traffic. Load-balancing algorithms in general are well understood, although their application in the context of the present invention is believed to be new.

20 If the traffic is destined for location B, then there are at least three paths from the current location (A) to location B: frame relay line 5, VPN line 1, or Internet line 2. In some embodiments, the invention determines whether the three connections are in load-balance mode or on-failure backup mode or a combination thereof. For a load-balance

mode, the controller 602 chooses the communication line based on load-balancing criteria. For backup mode, it chooses the communication line that is either the preferred line or (if the preferred line is down) the currently functional (backup) line.

By contrast with the preceding, if the traffic is destined for location C, then the
5 controller 602 at site A sends the traffic on the frame relay line, line 5.

Now let us look at the operation of the controller device 602 at location B. The device examines the IP data traffic sent to it and makes determinations like the following:

1. Is the traffic destined for the Internet, as opposed to one of the three "known" locations A, B, and C? If so, send the traffic over the Internet lines (line 3 and/or line
10 4). Load balancing decisions can be based on the criteria described above.
2. Is the traffic destined for location A? If so, then there are at least two paths to location A: the frame relay line 6, or VPN line 3. The controller 602 decides whether the two connections are in load-balance or on-failure backup mode, and chooses line(s) accordingly as discussed above.
- 15 3. Is the traffic destined for location C? If so, then send the traffic on the frame relay line, line 6.

To operate as discussed herein, the invention uses information about the IP address ranges in the locations reside as input data. For instance, a packet destined for the Internet 500 is one whose destination address is not in any of the address ranges of the
20 known locations (e.g., locations A, B, and C in the example of Figure 6). In some configurations, this is the same as saying that a packet destined for the Internet is one whose address is not in the address range of any of the organization's locations.

However, although all the known locations may belong to a single organization, that is

not a necessary condition for using the invention. Known locations may also belong to multiple organizations or individuals. Likewise, other locations belonging to the organization may be unknown for purposes of a given embodiment of the invention.

Address ranges can be specified and tested by the controller 602 using subnet
5 masks. The subnet masks may be of different lengths (contain a different number of one bits) in different embodiments and/or in different address ranges in a given embodiment. For instance, class B and class C addresses may both be used in some embodiments.

As another example, consider the illustrative network topology shown in Figure
10. This configuration has two locations A and B which are connected by a frame relay network 106 and by the Internet 500, through a frame relay router 105 and an Internet
10 router 104, at each location. For convenience, all routers are designated similarly in the Figures, but those of skill in the art will appreciate that different router models may be used, and in particular and without limitation, different routers may be used to connect to a private network than are used to connect to the Internet. Also, the controllers 602,
15 routers (and in Figure 6 the VPN interfaces 604) are shown separately in the Figures for convenience and clarity of illustration, but in various embodiments the software and/or hardware implementing these devices 602, 104, 105, 604 may be housed in a single device and/or reside on a single machine.

Suppose that the address ranges used by the routers in the Figure 10 configuration
20 are the following:

<u>Location</u>	<u>LAN IP</u>	<u>Internet</u>	<u>Frame Relay</u>
A	192.168.x.x	200.x.x.x	196.x.x.x
B	10.0.x.x	210.x.x.x	198.x.x.x

Without the invention, a topology like Figure 10 (but without the controllers 602)

requires some inflexible method of assigning packets to paths. Thus, consider a packet from location A that is meant for location B that has a destination address in the 10.0.x.x range. The network devices are pre-configured to such that all such packets with the 10.0.x.x destination address must be sent to the frame relay router (router Y), even though there is Internet connectivity between the two locations. Likewise, without the invention a packet from location A meant for location B which has a destination address not in the 10.0.x.x. range must be sent to the Internet router (router X) even though there is frame relay connectivity between the two locations.

Traditionally, such necessary match-ups of packets with routers were done by inflexible approaches such as sending all traffic from a given department, building, or local area network to a specified router. Manual and/or tedious reconfiguration was needed to change the destination address used in packets from a given source LAN such as one at site A, so this approach allowed load-balancing only on a very broad granularity, and did not load-balance dynamically in response to actual traffic. In particular, difficult reconfiguration of network parameters was needed to redirect packets to another router when the specified router went down.

By placing inventive modules 602 between locations and their routers as illustrated in Figure 10, however, the invention allows load-balancing, redundancy, or other criteria to be used dynamically, on a granularity as fine as packet-by-packet, to direct packets to an Internet router and/or a frame relay/point-to-point router according to the criteria. For instance, with reference to the illustrative network topology of Figure 10, if the inventive module 602 at location A receives a packet with a destination address in the 10.0.x.x range and the Internet router X is either down or over-loaded, then the

inventive module 602 can change the destination address so that it is in the 198.x.x.x
range (the rest of the address may be kept) and then send the modified packet to the frame
relay router Y. Similarly, if the frame relay path is down, overloaded, or insecure, then
the controller 602 can direct packets to the Internet after making the necessary destination
5 address changes to let the Internet router 104 operate successfully on those packets.

Unlike the configuration shown in Figure 1, the inventive configurations in
Figures 6 and 10 do not require manual intervention by network administrators to
coordinate traffic flow over parallel networks. The disparate networks are independent of
each other. When one attached network fails, the failure is sensed by the controller 602
10 and traffic is automatically routed through one or more other networks. Unlike the
configuration in Figure 2, the inventive configuration combines two or more disparate
networks. Unlike the configuration in Figure 4, the inventive configuration requires two
or more disparate networks be placed in parallel (although additional networks may also
be placed in series). Unlike the configuration in Figure 3, the inventive configuration
15 does not merely partition sites between unconnected networks – with the invention, most
or all of the connected sites get the benefit of parallel networks, so they can continue
transceiving even if one of the networks goes down.

Another difference between the inventive approach and prior approaches is the
narrow focus of some prior art on reliability. The present document takes a broader view,
20 which considers load balancing and security as well as reliability. Configurations like
those shown in Figure 2 are directed to reliability (which is also referred to by terms such
as “fault tolerance”, “redundancy”, “backup”, “disaster recovery”, “continuity”, and
“failover”). That is, one of the network paths (in this case, the one through the frame

relay network) is the primary path, in that it is normally used for most or all of the traffic, while the other path (in this case, the one through the ISDN link) is used only when that primary path fails. Although the inventive configurations can be used in a similar manner, with one network being on a primary path and the other network(s) being used
5 only as a backup when that first network fails, the inventive configurations also permit concurrent use of two or more disparate networks. With concurrent use, elements such as load balancing between disparate networks, and increased security by means of splitting pieces of a given message between disparate networks, which are not considerations in the prior art of Figure 2, become possibilities in some embodiments of the present
10 invention.

In some embodiments, a network at a location T is connected to a controller 602 for a location R but is not necessarily connected to the controller 602 at another location S. In such cases, a packet from location T addressed to location S can be sent over the network to the controller at location S, which can then redirect the packet to location T by
15 sending it over one or more parallel disparate networks. That is, controllers 602 are preferably, but not necessarily, provided at every location that can send packets over the parallel independent networks of the system.

In some embodiments, the controller 602 at the receiving end of the network connection between two sites A and B has the ability to re-sequence the packets. This
20 means that if the lines are of dissimilar speeds or if out-of-sequence transmission is required by security criteria, the system can send packets out of order and re-sequence them at the other end. Packets may be sent out of sequence to enhance security, to facilitate load-balancing, or both. The TCP/IP packet format includes space for a

sequence number, which can be used to determine proper packet sequence at the receiving end (the embodiments are dual-ended, with a controller 602 at the sending end and another controller 602 at the receiving end). The sequence number (and possibly more of the packet as well) can be encrypted at the sending end and then decrypted at the receiving end, for enhanced security. Alternately, an unused field in the TCP/IP header can hold alternate sequence numbers to define the proper packet sequence.

In the operation of some embodiments, the controller 602 on each location is provided with a configuration file or other data structure containing a list of all the LAN IP addresses of the controllers 602 at the locations, and their subnet masks. Each controller 602 keeps track of available and active connections to the remote sites 102. If any of the routes are unavailable, the controller 602 preferably detects and identifies them. When a controller 602 receives IP traffic to any of the distant networks, the data is sent on the active connection to that destination. If all connections are active and available, the data load is preferably balanced across all the routers. If any of the connections are unavailable, or any of the routers are down, the traffic is not forwarded to that router; when the routes become available again, the load balancing across all active routes preferably resumes.

In some embodiments, load balancing is not the only factor considered (or is not a factor considered) when the controller 602 determines which router should receive a given packet. Security may be enhanced by sending packets of a given message over two or more disparate networks. Even if a packet sniffer or other eavesdropping tool is used to illicitly obtain data packets from a given network, the eavesdropper will thus obtain at most an incomplete copy of the message because the rest of the message traveled over a

different network. Security can be further enhanced by sending packets out of sequence, particularly if the sequence numbers are encrypted.

Figure 7 is a diagram further illustrating a multiple disparate network access controller 602 of the present invention. A site interface 702 connects the controller 602 to the LAN at the site 102. This interface 702 can be implemented, for instance, as any local area network interface, like 10/100Base-T ethernet, gigabit ATM or any other legacy or new LAN technology.

The controller 602 also includes a packet path selector 704, which may be implemented in custom hardware, or implemented as software configuring semi-custom or general-purpose hardware. The path selector 704 determines which path to send a given packet on. In the configuration of Figure 6, for instance, the path selector in the controller at location A selects between a path through the router on line 1 and a path through the router on line 2. In different embodiments and/or different situations, one or more of the following criteria may be used to select a path for a given packet, for a given set of packets, and/or for packets during a particular time period:

- Redundancy: do not send the packet(s) to a path through a network, a router, or a connection that is apparently down. Instead, use devices (routers, network switches, bridges, etc.) that will still carry packets after the packets leave the selected network interfaces, when other devices that could have been selected are not functioning. Techniques and tools for detecting network path failures are generally well understood, although their application in the context of the present invention is believed to be new.

- Load-balancing: send packets in distributions that balance the load of a given network, router, or connection relative to other networks, routers, or connections available to the controller 602. This promotes balanced loads on one or more of the devices (routers, frame relay switches, etc.) that carry packets after the packets leave the selected network interfaces. Load-balancing may be done through an algorithm as simple as a modified round-robin approach which places the next packet on the next available line, or it may involve more complex algorithms that attempt to measure and track the throughput, latency, and/or other performance characteristics of a given link or path element. Load-balancing is preferably done on a per-packet basis for site-to-site data traffic or on a TCP or UDP session basis for Internet traffic, as opposed to prior art approaches which use a per-department and/or per-router basis for dividing traffic. Load-balancing algorithms in general are well understood, although their application in the context of the present invention is believed to be new.
- Security: divide the packets of a given message (session, file, web page, etc.) so they travel over two or more disparate networks, so that unauthorized interception of packets on fewer than all of the networks used to carry the message will not provide the total content of the message. Dividing message packets between networks for better security may be done in conjunction with load balancing, and may in some cases be a side-effect of load-balancing. But load-balancing can be done on a larger granularity scale than security, e.g., by sending one entire message over a first network A and the next entire message over a second, disparate network. Security may thus involve finer granularity than load

balancing, and may even be contrary to load balancing in the sense that dividing up a message to enhance security may increase the load on a heavily loaded path even though a more lightly loaded alternate path is available and would be used for the entire message if security was not sought by message-splitting between networks. Other security criteria may also be used, e.g., one network may be viewed as more secure than another, encryption may be enabled, or other security measures may be taken.

The controller 602 also includes two or more disparate network interfaces 706, namely, there is at least one interface 706 per network to which the controller 602 controls access. Each interface 706 can be implemented as a direct interface 706 or as an indirect interface 706; a given embodiment may comprise only direct interfaces 706, may comprise only indirect interfaces 706, or may comprise at least one of each type of interface.

An indirect interface 706 may be implemented, for instance, as a direct frame relay connection over land line or wireless or network interfaces to which the frame relay routers can connect, or as a point-to-point interface to a dedicated T1, T3, or wireless connection. One suitable implementation includes multiple standard Ethernet cards, in the controller 602 and in the router, which connect to each other. An external frame relay User-Network Interface (UNI) resides in a router 105 of a network 106; a similar Ethernet card resides in the Internet router 104. Each such Ethernet card will then have a specific IP address assigned to it. The controller can also have a single Ethernet card with multiple IP addresses associated with different routers and LANs. An indirect interface 706 may connect to the network over fiber optic, T1, wireless, or other links.

A direct interface 706 comprises a standard connection to the Internet 500, while another direct interface 706 comprises a standard connection to a VPN. One direct interface 706 effectively makes part of the controller 602 into a UNI by including in the interface 706 the same kind of special purpose hardware and software that is found on the frame relay network side (as opposed to the UNI side) of a frame relay network router. Such a direct frame relay network interface 706 is tailored to the specific timing and other requirements of the frame relay network to which the direct interface 706 connects. For instance, one direct interface 706 may be tailored to a Qwest frame relay network 106, while another direct interface 706 in the same controller 602 is tailored to a UUNet network 106. Another direct interface 706 comprises standard VPN components.

An indirect interface 706 relies on special purpose hardware and connectivity/driver software in a router or other device, to which the indirect interface 706 of the controller 602 connects. By contrast, a direct interface 706 includes such special purpose hardware and connectivity/driver software inside the controller 602 itself. In either case, the controller provides packet switching capabilities for at least redundancy without manual switchover, and preferably for dynamic load-balancing between lines as well. Figure 7 shows three interfaces 706; other controllers may have a different number of interfaces. The three interfaces 706 (for instance) may be implemented using a single card with three IP addresses, or three cards, each with one IP address. The site interface 702 may or may not be on the same card as interface(s) 706. The controller 602 in each case also optionally includes memory buffers in the site interface 702, in the path selector 704, and/or in the network interfaces 706.

An understanding of methods of the invention will follow from understanding the invention's devices, and vice versa. For instance, from Figures 6, 7, 10, and 11 one may ascertain methods of the invention for combining connections for access to multiple disparate networks, as well as systems and devices of the invention. As illustrated in
5 Figure 8, methods of the invention use a device such as controller 602. The controller 602 comprises (a) a site network interface 702, (b) at least two WAN network interfaces 706 tailored as necessary to particular networks, and (c) a packet path selector 704 which selects between network interfaces 706 according to a specified criterion. Path selection criteria may be specified 800 by configuration files, hardware jacks or switches, ROM
10 values, remote network management tools, or other means. Variations in topology are also possible, e.g., in a variation on Figure 10 the VPNs could swap position with their respective routers.

One then connects the site interface 702 to a site 102 to receive packets from a computer (possibly via a LAN) at the site 102. Likewise, one connects a first network
15 interface 706 to a first router for routing packets to a first network, and a second network interface 706 to a second router for routing packets to a second network, with the networks being disparate to each other. A third, fourth, etc. network may be similarly connected to the controller 602 in some embodiments and/or situations.

The connected disparate networks are parallel to one another (not serial, although
20 additional networks not directly connected to the controller 602 may be serially connected to the parallel disparate networks). The connected disparate networks are independent of one another, in that no routing information need be shared between them, to make them parallel (NNIs can still be used to connect networks in serial to form a larger independent

and parallel network). A mistake in the routing information for one network will thus not affect the other network.

After the connections are made (which may be done in a different order than recited here), one sends 802 a packet to the site interface 702. The controller 602 then
5 sends the packet through the one (or more – copies can be sent through multiple networks) network interface 706 that was selected by the packet path selector 704. The packet path selector 704 can maintain a table of active sessions, and use that table to select a path for packets in a given session. The packet path selector 704 does not need a session table to select paths for site-to-site traffic, because the controller 602 on the other
10 site knows where to forward the site-to-site-packets.

Figure 9 is a flowchart further illustrating methods of the present invention, which send packets over multiple parallel independent disparate networks for enhanced reliability, load balancing and/or security; frame relay networks and the Internet are used as an example, but point-to-point networks and VPNs may be similarly employed
15 according to the discussion herein.

During an address range information obtaining step 900, address ranges for known locations are obtained. Address ranges may be specified as partial addresses, e.g., partial IP addresses in which some but not all of the address is specified. Thus, “198.x.x.x” indicates an IP address in which the first field is 198 and the other three
20 address fields are not pinned down, corresponding to the range of addresses from 198.0.0.0 to 198.255.255.255. Each address range has an associated network; a network may have more than one associated contiguous range of addresses which collectively constitute the address range for that network. The locations reachable through the

network have addresses in the address range associate with the network. Since part of the address specifies the network, a location reachable through two networks has two addresses, which differ in their network-identifying bits but are typically the same in their other bits. Address ranges may be obtained 900 by reading a configuration file, querying
5 routers, receiving input from a network administrator, and/or other data gathering means.

During a topology information obtaining step 902, topology information for the system of parallel disparate networks is obtained. The topology information specifies which one or more networks can be used (if functioning) to reach which known locations. With regard to Figure 6, for instance, the topology information could be represented by a
10 table, list, or other data structure which specifies that: the VPN connects sites A and B; the Internet connects sites A and B; and the private (frame relay/point-to-point) network connects sites A, B, and C. Topology information may be obtained 902 by reading a configuration file, querying routers, receiving input from a network administrator, and/or other data gathering means.

15 If necessary, a connection forming step is performed, e.g., to obtain a virtual circuit between two sites 102. The controller 602 then checks the status of each connection and updates the information for available communication paths.

The controller 602 at each location will go through the address range information obtaining step, topology information obtaining step and connection forming step. More
20 generally, the steps illustrated and discussed in this document may be performed in various orders, including concurrently, except in those cases in which the results of one step are required as input to another step. Likewise, steps may be omitted unless required

by the claims, regardless of whether they are expressly described as optional in this Detailed Description. Steps may also be repeated, or combined, or named differently.

During a packet receiving step 904, the controller 602 at a given source location receives a packet to be sent from that location to the destination site 102. In some cases, 5 multiple packets may be received in a burst. The packet comes into the controller 602 through the site interface 702.

During a determining step 906, the controller 602 (or some other device used in implementing the method) looks at the packet destination address to determine whether the destination address lies within a known address range. That is, the destination address 10 is compared to the known location address ranges that were obtained during step 900, in order to see whether the destination location is a known location. Only packets destined for known locations are potentially rerouted by the invention to balance loads, improve security, and/or improve reliability. Packets destined for unknown locations are simply sent to the network indicated in their respective destination addresses, which is the 15 Internet 500 in the examples given herein but could also be some other "catch-all" network. Although they are not rerouted, such packets may nonetheless be counted as part of the load balancing calculation.

During a path selecting step 908, the path selector 704 selects the path over which the packet will be sent; selection is made between at least two paths, each of which goes 20 over a different network 106 than the other. The disparate networks are independent parallel networks. This path selecting step 908 may be performed once per packet, or a given selection may pertain to multiple packets. In some embodiments, selecting a network will also select a path, as in the system shown in Figure 10. In other cases, there

may be more than one path to a given network, as discussed in connection with the line pairs shown in Figure 6. Packet path selection 908 is shown as following packet receipt 904, but in some embodiments and/or some situations, it may precede packet receipt 904. That is, in some cases the path for the next packet may be determined by the packet path selector before the packet arrives, e.g., in a round-robin manner, while in other cases the path is determined after the packet arrives, e.g., using per-packet dynamic load balancing.

As indicated, the path selection may use 910 load balancing as a criterion for selecting a path, use 912 network status (up/down) and other connectivity criteria (e.g., router status, connectivity status) as a criterion for selecting a path, and/or use 914 division of packets between disparate networks for enhanced security as a criterion for selecting a path. These steps may be implemented in a manner consistent with the description above of the path selector 704 given in the discussion of Figure 7. More generally, unless it is otherwise indicated, the description herein of systems of the present invention extends to corresponding methods, and vice versa.

The description of systems and methods likewise extend to corresponding computer-readable media (e.g., RAM, ROM, other memory chips, disks, tape, Iomega ZIP or other removable media, and the like) which are configured by virtue of containing software to perform an inventive method, or software (including any data structure) which is uniquely suited to facilitate performance of an inventive method. Articles of manufacture within the scope of the present invention thus include a computer-readable storage medium in combination with the specific physical configuration of a substrate of the computer-readable storage medium, when that substrate configuration represents data

and/or instructions which cause one or more computers to operate in a specific and predefined manner as described and claimed herein.

No change to packet source IP address or destination IP address need by done by the controller in a topology like that shown in Figure 10. The controller 602 sends the
 5 packet to router X or router Y as determined by the packet path selector. This is illustrated in the following summary example:

	<u>Packet location</u>	<u>Packet Source IP Address</u>	<u>Packet Destination IP Address</u>
	Leaving site A	Site A's IP address	Site B's IP address
	Leaving controller A	Site A's IP address	Site B's IP address
10	Leaving VPN/ Router	VPN/ Router/ Site A	VPN/ Router/ Site B
	<packet travels over Internet/ frame relay net/ etc.>		
	Arrival VPN/ Router	VPN/ Router/ Site A	VPN/ Router/ Site B
	Arrival controller B	Site A	Site B
	<controller may need to resequence packets>		
15	Arrival at site B	Site A	Site B

However, packet addresses are modified during operation of a configuration like that shown in Figure 11. An example is provided in the following summary example:

	<u>Packet location</u>	<u>Packet Source IP Address</u>	<u>Packet Destination IP Address</u>
20	Leaving site A	Site A's IP address	Site B's IP address
	Leaving VPN A	VPN A's IP address	VPN B's IP address
	Leaving controller A	A controller A IP address	A controller B IP address
	<packet travels over Internet/ frame relay net/ etc.>		
	Arrival controller B	The controller A IP address	The controller B IP address
25	<controller may need to resequence packets>		
	Arrival at VPN B	VPN A's IP address	VPN B's IP address
	<note that the controllers are transparent to the VPNs>		
	Arrival at site B	Site A	Site B
	<the VPNs are transparent to the sites>		

30

During an address modifying step 916, the packet destination address is modified as needed to make it lie within an address range (obtained during step 900) which is associated with the selected path to the selected network (selected during step 908). For instance, if a packet is received 904 with a destination address corresponding to travel

through the Internet but the path selection 908 selects a path for the packet through a frame relay network 106 to the same destination, then the packet's destination IP address is modified 916 by replacing the IP address with the IP address of the appropriate interface of the controller at Site B. Also the packet's source IP address is replaced with the IP address of the appropriate interface of the source controller. This modifying step may be viewed as optional, in the sense that it need not be performed in every embodiment. But it is required in the sense that a system embodiment of the invention which is claimed with a limitation directed to destination address modification must be at least capable of performing the modifying step, and a method embodiment which is claimed with a limitation directed to the modifying step must perform the modifying step on at least one packet.

With regard to both Figure 10 and Figure 11, during a packet transmission step 918, the packet is sent on the selected 908 path. This is done by sending the packet over the network interface 706 for the path selected. As indicated in Figure 9, the method may then loop back to receive 904 the next packet, select 908 a network for that packet, send 918 it, and so on. As noted, other specific method instances are also possible. One example is the inventive method in which load balancing or reliability criteria cause an initial path selection to be made 908, and then a loop occurs in which multiple packets are received 904 and then sent 918 over the selected path without repeating the selecting step 908 for each receive 904 – send 918 pair. Note that some embodiments of the invention permit packets of a given message to be sent over two or more disparate networks, thereby enhancing 914 security. An ending step may be performed as needed during an orderly shutdown for diagnostic or upgrade work, for instance.

The controller 602 at the destination site goes through the steps described above in reverse order as needed. The controller 602 receives the packet from the source location through one of the network interfaces. Packet resequencing may be needed in either the Figure 10 or the Figure 11 configuration, while address changes are needed in
5 the Figure 11 configuration only.

Conclusion

The present invention provides methods and devices for placing frame relay and other private networks in parallel with VPNs and other Internet-based networks, thereby
10 providing redundancy without requiring manual switchover in the event of a network failure. Load-balancing between lines and/or between networks may also be performed. For instance, the invention can be used to provide reliable, efficient, and secure point-to-point connections for private networks 106 in parallel with a VPN and an SSL Internet connection. Some prior art approaches require network reconfiguration each time a frame
15 relay circuit fails, and some have complex router configurations to handle load balancing and network failures. This requires substantial effort by individual network customers to maintain connectivity, and they will often receive little or no help from the frame relay carriers, or not receive prompt service from a VPN provider. Instead, well-trained staff are needed at each location, as are expensive routers. By contrast, these requirements are
20 not imposed by the present invention.

As used herein, terms such as "a" and "the" and item designations such as "connection" or "network" are generally inclusive of one or more of the indicated item.

In particular, in the claims a reference to an item normally means at least one such item is required.

The invention may be embodied in other specific forms without departing from its essential characteristics. The described embodiments are to be considered in all respects
5 only as illustrative and not restrictive. Headings are for convenience only. The claims form part of the specification. The scope of the invention is indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed and desired to be secured by patent is:

10

1. A controller which controls access to multiple independent disparate networks in a parallel network configuration, the disparate networks comprising at least one private network and at least one network based on the Internet, the controller comprising:

5 a site interface connecting the controller to a site;

at least two network interfaces which send packets toward the disparate networks;

and

a packet path selector which selects between network interfaces according to at

least: a destination of the packet, an optional presence of alternate paths to

10 that destination, and at least one specified criterion for selecting between

alternate paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector.

15 2. The controller of claim 1, wherein the controller controls access to a frame relay private network through a first network interface of the controller, and the controller controls access to the Internet through a second network interface of the controller.

20 3. The controller of claim 1, wherein the packet path selector selects between network interfaces according to a load-balancing criterion, thereby promoting balanced loads on devices that carry packets on the selected path after the packets leave the selected network interfaces.

4. The controller of claim 1, wherein the packet path selector selects between network interfaces according to a reliability criterion, thereby promoting use of devices that will still carry packets on the selected path after the packets leave the selected network interfaces, when other devices on a path not selected are not functioning.

5. The controller of claim 1, wherein the packet path selector selects between network interfaces according to a security criterion, thereby promoting use of multiple disparate networks to carry different pieces of a given message so that unauthorized interception of packets on fewer than all of the disparate networks used to carry the message will not provide the total content of the message.

6. The controller of claim 1, wherein the controller sends packets out of sequence over the parallel disparate networks.

7. The controller of claim 6, wherein the controller places an encrypted sequence number in at least some of the packets which are sent out of sequence.

8. The controller of claim 1, wherein the controller sends packets from a selected network interface to a VPN.

9. The controller of claim 1, wherein the controller sends packets from a selected network interface to a point-to-point private network connection.

10. A controller which controls access to multiple networks in a parallel network configuration, suitable networks comprising Internet-based networks and private networks from at least one more provider, in combination, the controller comprising:

5

a site interface connecting the controller to a site;

at least two network interfaces which send packets toward the networks; and

a packet path selector which selects between network interfaces on granularity

which is at least as fine as session-by-session according to at least: a

destination of the packet, an optional presence of alternate paths to that

10

destination, and at least one specified criterion for selecting between

alternate paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the

packet through the network interface that was selected by the packet path

selector.

15

11. A controller which controls access to multiple networks in a parallel network configuration, suitable networks comprising Internet-based networks and private networks from at least one more provider, in combination, the controller comprising:

a site interface connecting the controller to a site;

20

at least two network interfaces which send packets toward the networks; and

a packet path selector which selects between network interfaces on a per-packet

basis according to at least: a destination of the packet, an optional

presence of alternate paths to that destination, and at least one specified

criterion for selecting between alternate paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector.

5

12. A method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of:

obtaining at least two known location address ranges which have associated networks;

10

obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location;

receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location;

15

determining whether the destination address lies within a known location address range;

selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location;

20

forwarding the packet on the selected network path.

13. The method of claim 12, further comprising the step of modifying the packet destination address to lie within a known location address range associated with the selected network before the forwarding step.

5 14. The method of claim 12, wherein the forwarding step forwards the packet toward the Internet when the packet's destination address does not lie within any known location address range.

10 15. The method of claim 12, wherein the destination address identifies a destination location to which only a single associated network provides connectivity from the current location, and the forwarding step forwards the packet to that single associated network.

15 16. The method of claim 12, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

17. The method of claim 12, wherein repeated instances of the selecting step make network path selections on a per session basis.

20 18. The method of claim 12, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.

19. The method of claim 18, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance line loads by distributing packets between lines.

5 20. The method of claim 18, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance network loads by distributing packets between disparate networks.

10 21. The method of claim 12, wherein the selecting step selects the network path at least in part on the basis of a reliability criterion.

22. The method of claim 12, wherein the selecting step selects the network path at least in part on the basis of a security criterion.

15 23. The method of claim 12, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a private network such that the modified packet destination address lies instead in a known location address range associated with a VPN.

20 24. The method of claim 12, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a VPN

such that the modified packet destination address lies instead in a known location address range associated with a private network.

25. The method of claim 12, wherein the modifying step modifies a packet destination address corresponding to one of: the Internet, a private network, thereby making the modified packet destination address correspond to the other of: the Internet, a private network.

26. A method for combining connections for access to parallel networks, the method comprising the steps of:

sending a packet to a site interface of a controller, the controller comprising the site interface which receives packets, at least two network interfaces to parallel networks, and a packet path selector which selects between the network interfaces on a per-session basis to promote load-balancing; and forwarding the packet, possibly with a modified destination address, through the network interface selected by the packet path selector.

27. The method of claim 26, wherein the step of sending a packet to the controller site interface is repeated as multiple packets are sent, and the controller sends different packets of a given message to different parallel networks.

28. The method of claim 26, wherein the step of sending a packet to the controller site interface is repeated as multiple packets are sent, the network interfaces

include at least two VPN line interfaces and a private network interface, and the packet path selector selects between at least those three interfaces.

29. The method of claim 26, further comprising the step of sensing failure of
5 one of parallel disparate networks and automatically sending traffic through at least one other parallel disparate network.

30. A method for combining connections for access to parallel networks, the method comprising the steps of:

10 receiving at a first controller a packet which has a first site IP address as source address and a second site IP address as destination address;
modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address;
and
15 forwarding the modified packet along a selected path toward the second site.

31. A method for combining connections for access to parallel networks, the method comprising the steps of:

20 receiving at a first controller a packet which has a first VPN IP address as source address and a second VPN IP address as destination address;
modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address;
and

forwarding the modified packet along a selected path toward the second VPN.

32. A method for combining connections for access to disparate parallel networks, the method comprising the steps of:

- 5 receiving at a controller a packet which has a first site IP address as source address and a second site IP address as destination address;
- selecting, within the controller, between a path through an Internet-based network and a path through a private network that is not Internet-based; and
- 10 forwarding the packet along the selected path toward the second site.

33. A computer storage medium having a configuration that represents data and instructions which will cause performance of a method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of:

- 15 obtaining at least two known location address ranges which have associated networks;
- obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location;
- receiving at the current location a packet which identifies a particular destination
- 20 location by specifying a destination address for the destination location;
- determining whether the destination address lies within a known location address range;

selecting a network path from among paths to disparate associated networks, said
networks being in parallel at the current location, each of said networks
specified in the topology information as capable of providing connectivity
between the current location and the destination location;
5 modifying the packet destination address to lie within a known location address
range associated with the selected network if it does not already do so; and
forwarding the packet on the selected network path.

AI
M...


34. The configured storage medium of claim 33, wherein the selecting step
10 selects the network path at least in part on the basis of a dynamic load-balancing criterion.

35. The configured storage medium of claim 33, wherein repeated instances of
the selecting step make network path selections on a packet-by-packet basis.

ABSTRACT


Methods, configured storage media, and systems are provided for communications using two or more disparate networks in parallel to provide load balancing across network connections, greater reliability, and/or increased security. A controller provides access to
5 two or more disparate networks in parallel, through direct or indirect network interfaces. When one attached network fails, the failure is sensed by the controller and traffic is routed through one or more other disparate networks. When all attached disparate networks are operating, one controller preferably balances the load between them.

10 \pc0




(PRIOR ART)

Fig. 1

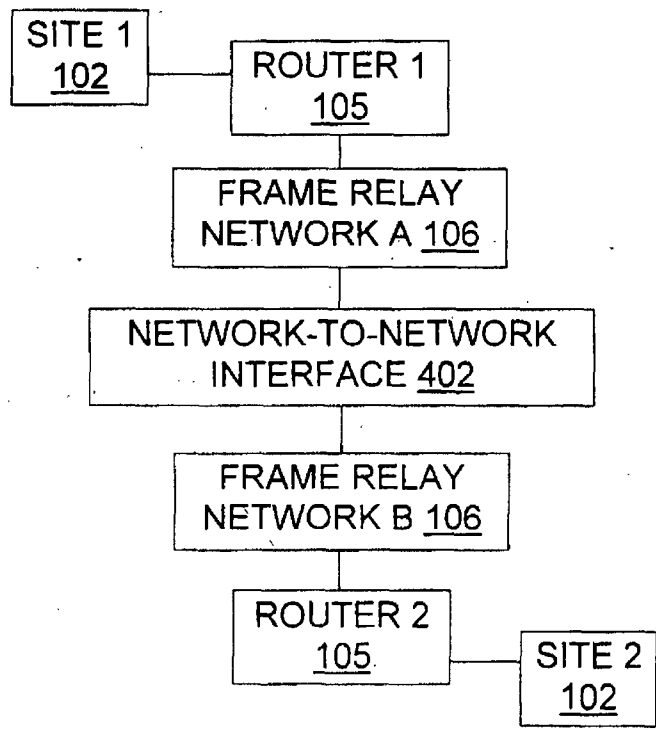


(PRIOR ART)

Fig. 2




(PRIOR ART)
Fig. 3



(PRIOR ART)
Fig. 4

PRINT OF DRAWINGS
AS ORIGINALLY FILED



(PRIOR ART)
Fig. 5




Fig. 7

PRINT OF DRAWINGS
AS ORIGINALLY FILED

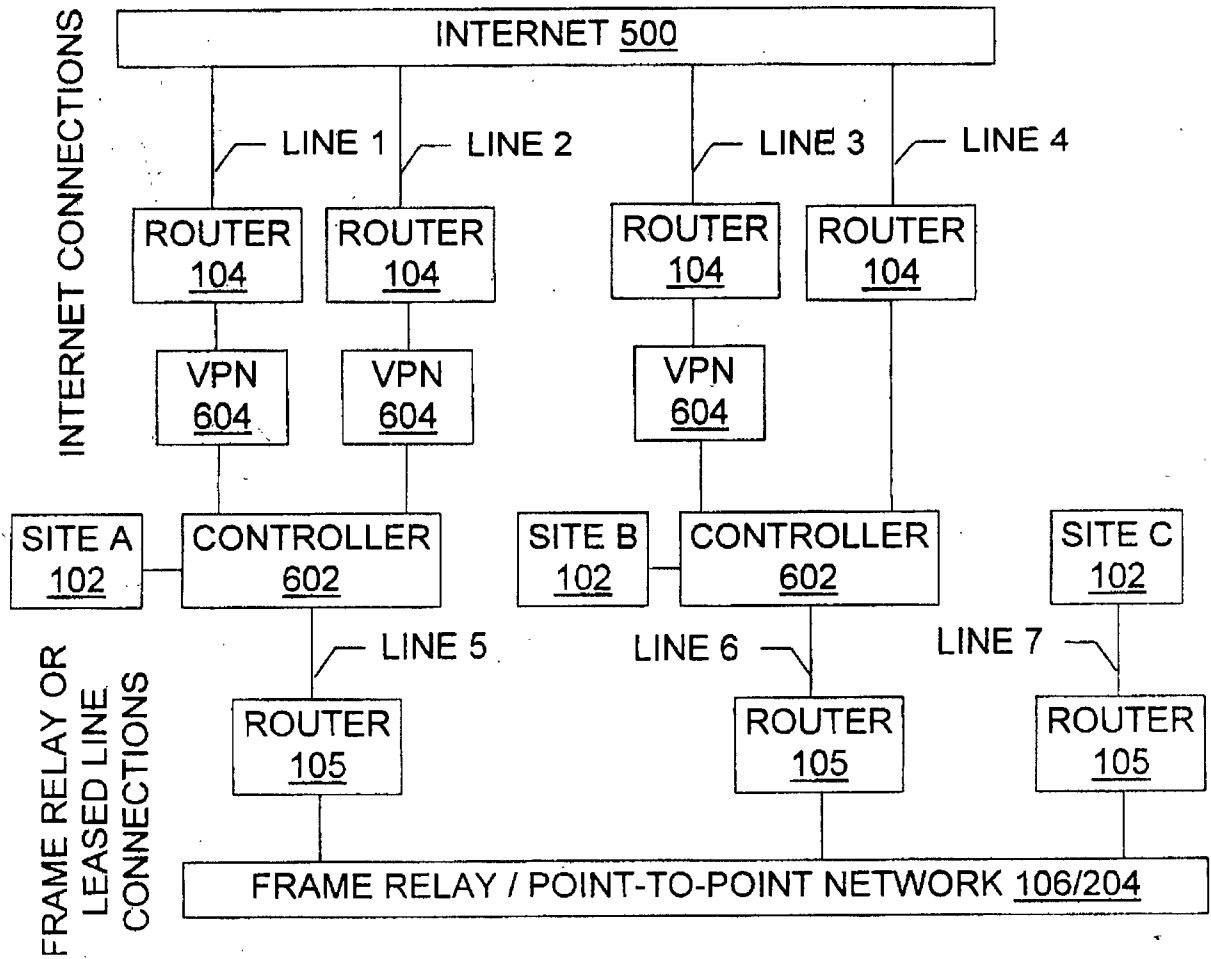


Fig. 6




Fig. 10

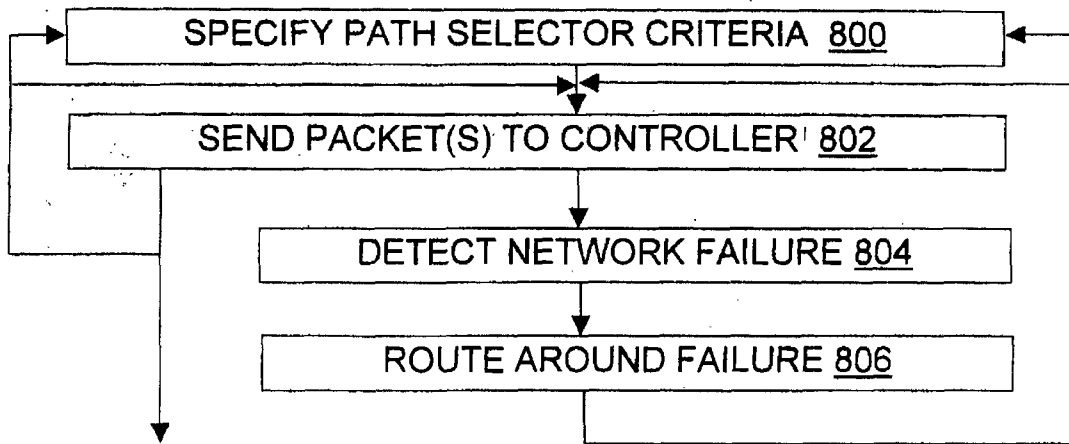


Fig. 8

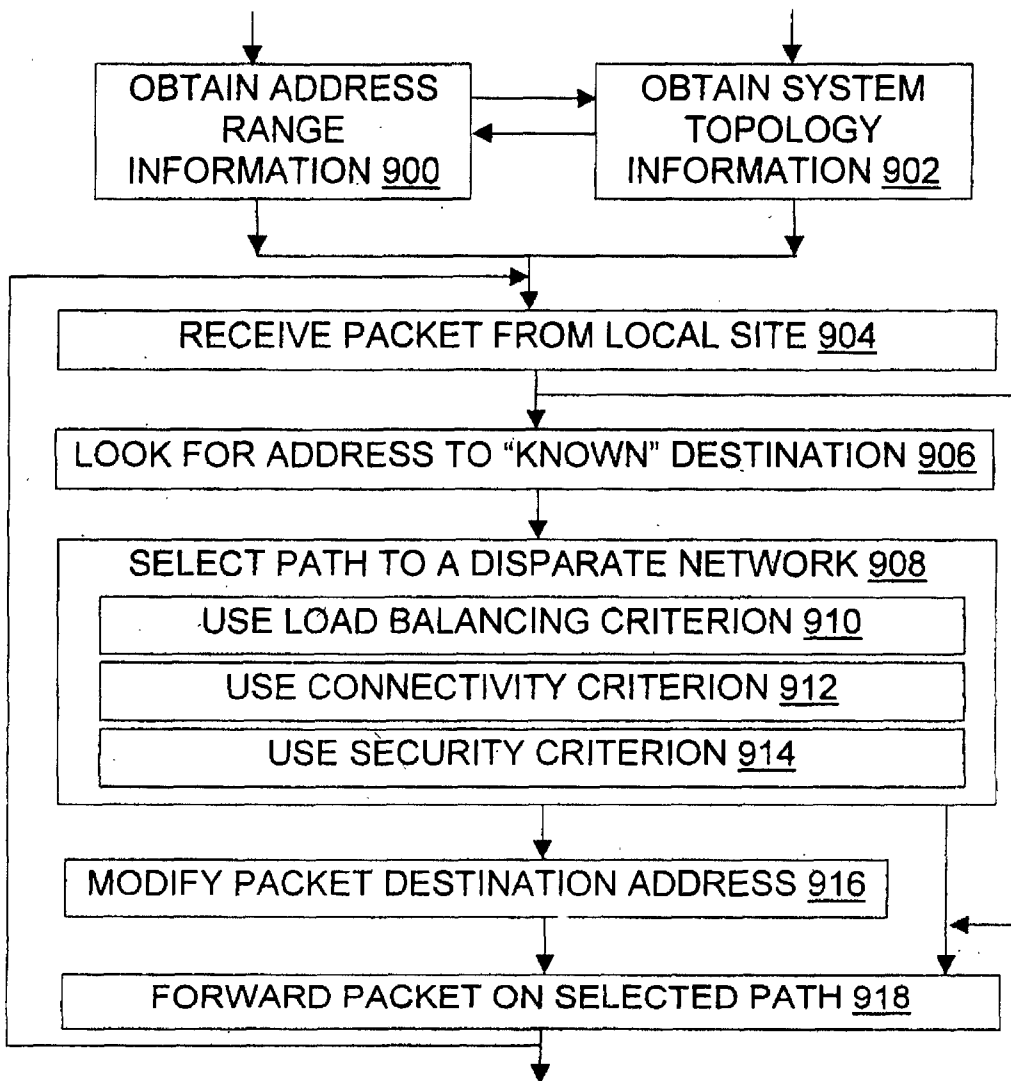




Fig. 9

10361837.020703 ✓
10/36/839




(PRIOR ART)
Fig. 1




(PRIOR ART)
Fig. 2

11




(PRIOR ART)

Fig. 3



(PRIOR ART)

Fig. 4



(PRIOR ART)
Fig. 5

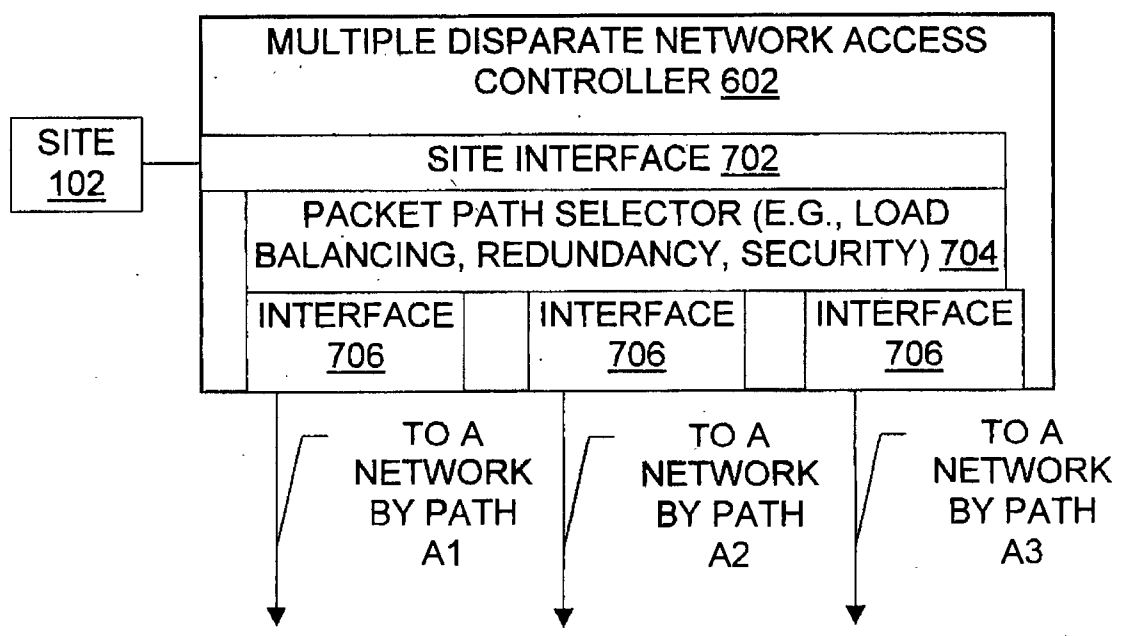


Fig. 7

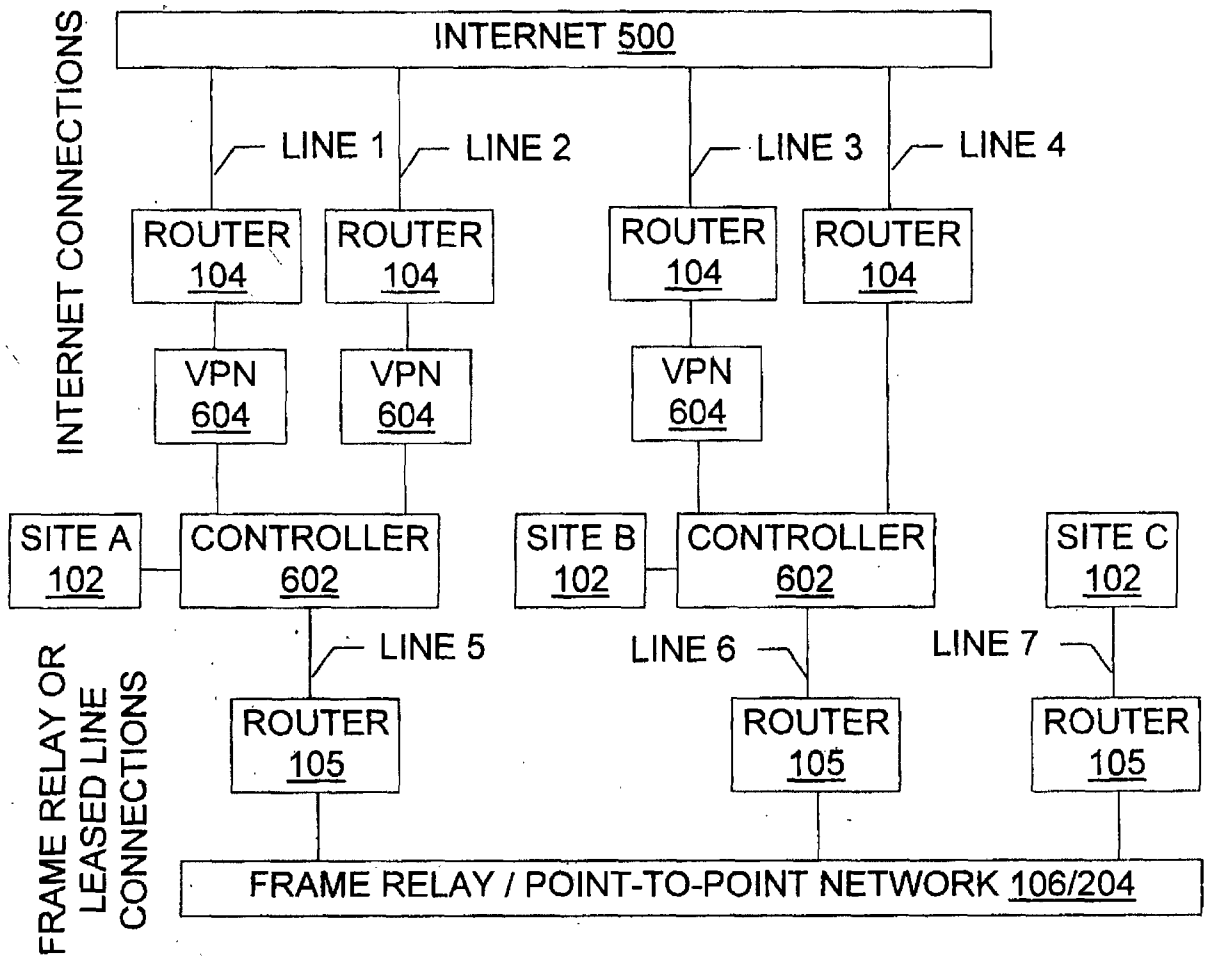


Fig. 6




Fig. 10




Fig. 8




Fig. 9




Fig. 11

Application Data Sheet

Application Information

Application Type:: Regular
Subject Matter:: Utility
Suggested Classification::
Suggested Group Art Unit::
CD-ROM or CD-R? None
✓ Title:: TOOLS AND TECHNIQUES FOR DIRECTING
PACKETS OVER DISPARATE NETWORKS
Attorney Docket Number:: 3003.2.11A
Request for Early Publication?:: No
Request for Non-Publication?:: No
Suggested Drawing Figure:: 9
Total Drawing Sheets:: 6
Small Entity:: Yes
Petition included?:: No
Petition Type::
Secrecy Order in Parent Appl.?:: No

Applicant Information

Applicant Authority type:: Inventor
Primary Citizenship Country:: US
Status:: Full Capacity
Given Name:: Sanchaita
Middle Name::
Family Name:: Datta
Name Suffix::

City of Residence:: Salt Lake City
State of Province of Residence:: UT
Country of Residence:: US
Street of mailing address:: 4540 South Jupiter Drive
City of mailing address:: Salt Lake City
State or Province of mailing address:: UT
Postal or Zip Code of mailing address:: 84124

Applicant Authority type:: Inventor
Primary Citizenship Country:: US
Status:: Full Capacity
Given Name:: Bhaskar
Middle Name::
Family Name:: Ragula
Name Suffix::

City of Residence:: Salt Lake City
State of Province of Residence:: UT
Country of Residence:: US
Street of mailing address:: 4540 South Jupiter Drive
City of mailing address:: Salt Lake City
State or Province of mailing address:: UT
Postal or Zip Code of mailing address:: 84124

Correspondence Information

Correspondence Customer Number:: 23484

Representative Information

Representative Customer Number::	23484
----------------------------------	-------

10361837 .020703

Domestic Priority Information

Application::	Continuity Type::	Parent Application::	Parent Filing Date::
This Application	Continuation-in-part of	10/034,197	12/28/01
10/034,197	Non-Provisional of	60/259,269	12/29/00
This Application	Non-Provisional of	60/355,509	02/08/02

Assignee Information

Assignee Name:: Ragula Systems d/b/a/ FatPipe Networks

Docket No. 3003.2.11A

**INVENTORS' DECLARATION
FOR UTILITY PATENT APPLICATION**

Application of **Sanchaita Datta and Ragula Bhaskar:**

As below named inventors, we hereby declare that:

Our residences, post office addresses, and citizenship are as stated below next to our names; we believe we are the original, first and joint inventors of the subject matter which is claimed and for which a patent is sought on the invention entitled TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS, the specification of which is to be filed concurrently with this Declaration; we have reviewed and understand the contents of said specification, including the claims; and we acknowledge the duty to disclose information which is material to patentability as defined in Title 37 Code of Federal Regulations, § 1.56.

We claim the benefit, to the extent possible under 35 United States Code § 120 and otherwise, of the following United States application:

Serial No. 10/034,197 filed December 28, 2001

In particular, the present application is a continuation-in-part of application no. 10/034,197 filed December 28, 2001, and it claims priority through that '197 application to the earlier application no. 60/259,269 filed December 29, 2000.

We claim the benefit, to the extent possible under 35 United States Code § 119(e) and otherwise, of the following United States provisional application:

Serial No. 60/355,509 filed February 8, 2002

As named inventors, we have appointed the following registered practitioners to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith:

10361557 020703

John W. L. Ogilvie, Reg. No. 37,987
Genie L. Ogilvie, Reg. No. 43,841

Please direct all correspondence to:

John W. L. Ogilvie
Computer Law++
1211 East Yale Avenue
Salt Lake City, Utah 84105
801-582-2724 (voice)
801-583-1984 (fax)

We hereby declare that all statements made herein of our own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Inventors

Sanchaita Datta, residing at and having a post office address of:
4540 South Jupiter Drive, Salt Lake City, Utah 84124
Citizenship: USA

Signed: Sanchaita Datta Date: Feb 5, 03

Ragula Bhaskar, residing at and having a post office address of:
4540 South Jupiter Drive, Salt Lake City, Utah 84124
Citizenship: USA

Signed: R Bhaskar Date: Feb 5, 03

* * * * *

Application Assignment Record

According to the application transmittal letter, an assignment recording ownership was filed with this application; however, a copy of this record was not located in the original file history record obtained from the United States Patent and Trademark Office. Upon your request, we will attempt to obtain the assignment documents from the Assignment Recordation Branch of the United States Patent and Trademark Office or from a related application case (if applicable). Please note that additional charges will apply for this service.

This page is not part of the official USPTO record. It has been determined that content identified on this document is missing from the original file history record.



0430 0780
03/05/03
#2

PATENT APPLICATION
Docket No.: 3003.2.11A

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Sanchaita Datta and Ragula Bhaskar
Serial No.: 10/361837
Filed: February 7, 2003
For: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS
OVER DISPARATE NETWORKS

FIRST INFORMATION DISCLOSURE STATEMENT

Honorable Commissioner of
Patents & Trademarks
Washington, D. C.

Commissioner:

This Information Disclosure Statement is filed in response to the duty of candor described in 37 C.F.R. §§ 1.56, 1.98, MPEP § 2001.06(c), and elsewhere. The references listed on the enclosed Form PTO-1449 (incorporated herein by reference) are respectfully submitted for consideration by the Office.

Dated March 12, 2003.

CERTIFICATE OF MAILING
I hereby certify that the correspondence listed below is being deposited with the United States Postal Service as Priority Mail, postage paid, on March 12, 2003 addressed to the Commissioner for Patents, P.O. Box 2327, Arlington, VA 22202:

Postcard
First IDS w/ PTO-1449 and 25 references

p-ids

Respectfully submitted,

JOHN W.L. OGILVIE
Registration No. 37,987

COMPUTER LAW++
1211 East Yale Avenue
Salt Lake City, Utah 84105
801-582-2724 voice
801-583-1984 fax

Form PTO-1449

Sheet 1 of 2


Applicant: Sanchaita Datta and Ragula Bhaskar

Serial No.: 10/361837

Att'y Docket No. 3003.2.11A

Filing Date: February 7, 2003

For: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS



LIST OF REFERENCES CITED BY APPLICANT {PRIVATE }

U.S. Patent Documents

Examiner Initial*	Document Number	Issue Date	Name	Class	Sub Class	Filing Date
<u>MM</u>	A1	6,493,349	12/10/02	Casey	370 409	11/13/98
<u>MM</u>	A2	6,493,341	12/10/02	Datta et al.	370 392	12/29/00
<u>MM</u>	A3	6,438,100	08/20/02	Halpern et al.	370 218	08/05/99
<u>MM</u>	A4	6,339,595	01/15/02	Rekhter et al.	370 392	12/23/97
<u>MM</u>	A5	6,295,276	09/25/01	Datta et al.	370 218	12/31/99
<u>MM</u>	A6	6,253,247	06/26/01	Bhaskar et al.	709 237	05/20/97
<u>MM</u>	A7	5,948,069	09/07/99	Kitai et al.	709 240	07/19/96
<u>MM</u>	A8	5,737,526	04/07/98	Periasamy et al.	395 200.06	12/30/94
<u>MM</u>	A9	5,473,599	12/05/95	Li et al.	370 16	04/22/94
<u>MM</u>	A10	5,420,862	05/30/95	Perlman	370 85.13	06/14/91
<u>MM</u>	A11	5,398,012	03/14/95	Derby et al.	340 825.03	11/24/92

Other Documents

(including Author, Title, Pertinent Pages, etc.)

- MM A12 U.S. Patent Application, Attorney Docket No. 3003.2.9A; see USPTO published application no. US-2002-0087724-A1, 07/04/2002
- MM A13 T. Liao et al., "Using multiple links to interconnect LANs and public circuit switched data networks," *Proc. Int. Conference on Communications Systems: Towards Global*

Examiner: Melvin Marcelo

Date Considered: 02-09-2004

*EXAMINER: Please initial if reference considered, whether or not citation is in conformance with MPEP 609; draw line through citation if not in conformance and not considered. Please include a copy of this form with the next communication to applicant.

Form PTO-1449

Sheet 2 of 2


Applicant: Sanchaita Datta and Ragula Bhaskar

Serial No.: 10/361837

Att'y Docket No. 3003.2.11A

Filing Date: February 7, 2003

For: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS



Integration, Vol. 1, Singapore, 59 November 1990, pp. 289-293

- mm A14 Press release from www.coyotepoint.com, September 8, 1997,
- mm A15 Network Address Translation Technical Discussion, from safety.net; no later than 05/07/1999
- mm A16 Higginson et al., "Development of Router Clusters to Provide Fast Failover in IP Networks," from www.asia-pacific.digital.com; no later than 9/29/98
- mm A17 Pages from www.navpoint.com; no later than 12/24/2001
- mm A18 "The Basic Guide to Frame Relay Networking", pp. 1-85, copyright date 1998
- mm A19 "NNI & UNI", pp. 1-2, Nov 16, 2001
- mm A20 "Disaster Recovery for Frame Relay Networks", pp. 1-14, no later than 12/7/2001
- mm A21 T. Nolle, "Watching Your Back", pp. 1-3, 11/01/99
- mm A22 "Multi-Attached and Multi-Homed Dedicated Access", pp. 1-5, no later than 12/8/2001
- mm A23 Feibel, "Internetwork Link," Novell's® Complete Encyclopedia of Networking, copyright date 1995
- mm A24 Tanenbaum, Computer Networks (3rd Ed.), pp. 396-406; copyright date 1996
- mm A25 Wexler, "Frame Relay and IP VPNs: Compete Or Coexist?", from www.bcr.com; July 1999

Examiner: Mylo Marcelo

Date Considered: 02-09-2004

*EXAMINER: Please initial if reference considered, whether or not citation is in conformance with MPEP 609; draw line through citation if not in conformance and not considered. Please include a copy of this form with the next communication to applicant.

(A12)

Express Mail Label No. EL855688731US
PATENT APPLICATION
DOCKET NO. 3003.2.9A

See: USPTO published application no. US-2002-0087724-A1 07/24/2002

UNITED STATES
PATENT APPLICATION

OF

SANCHAITA DATTA AND RAGULA BHASKAR

FOR

COMBINING CONNECTIONS FOR PARALLEL ACCESS TO
MULTIPLE FRAME RELAY AND OTHER PRIVATE NETWORKS

COMBINING CONNECTIONS FOR PARALLEL ACCESS TO MULTIPLE FRAME RELAY AND OTHER PRIVATE NETWORKS

5 RELATED APPLICATIONS

This application claims priority to commonly owned copending U.S. provisional patent application serial no. 60/259,269 filed December 29, 2000, which is also incorporated herein by reference.

10 FIELD OF THE INVENTION

The present invention relates to computer network data transmission, and more particularly relates to tools and techniques for point-to-point or switched connection communications such as those using two or more frame relay networks in parallel to provide benefits such as load balancing across network connections, greater reliability,
15 and increased security.

TECHNICAL BACKGROUND OF THE INVENTION

Frame relay networking technology offers relatively high throughput and reliability. Data is sent in variable length frames, which are a type of packet. Each frame
20 has an address that the frame relay network uses to determine the frame's destination. The frames travel to their destination through a series of switches in the frame relay network, which is sometimes called a network "cloud"; frame relay is an example of packet-switched networking technology. The transmission lines in the frame relay cloud must be essentially error-free for frame relay to perform well, although error handling by other
25 mechanisms at the data source and destination can compensate to some extent for lower

line reliability. Frame relay and/or point-to-point network services are provided or have been provided by various carriers, such as AT&T, Qwest, XO, and MCI WorldCom.

Frame relay networks are an example of a "private network". Another example is a point-to-point network, such as a T1 or T3 connection. Although the underlying
5 technologies differ somewhat, for purposes of the present invention frame relay networks and point-to-point networks are generally equivalent in important ways, such as the conventional reliance on manual switchovers when traffic must be redirected after a connection fails. A frame relay permanent virtual circuit is a virtual point-to-point connection. Frame relays are used as examples throughout this document, but the
10 teachings will also be understood in the context of point-to-point networks.

A frame relay or point-to-point network may become suddenly unavailable for use. For instance, both MCI WorldCom and AT&T users have lost access to their respective frame relay networks during major outages. During each outage, the entire network failed. Loss of a particular line or node in a network is relatively easy to work
15 around. But loss of an entire network creates much larger problems. Tools and techniques are needed to permit continued data transmission when the entire frame relay network that would normally carry the data is down.

Figure 1 illustrates prior art configurations involving two frame relay networks for increased reliability; similar configurations involve one or more point-to-point network
20 connections. Two sites 102 transmit data to each other (alternately, one site might be only a data source, while the other is only a data destination). Each site has two border routers 104. Two frame relay networks 106, 108 are available to the sites 102 through the routers 104. The two frame relay networks 106, 108 have been given separate numbers in the

figure, even though each is a frame relay network, to emphasize the incompatibility of frame relay networks provided by different carriers. An AT&T frame relay network, for instance, is incompatible in many details with an MCI WorldCom frame relay network. For instance, two frame relay networks may have different maximum frame sizes or
5 switching capacities. The two providers have to agree upon information rates, switching capacities, frame sizes, etc. before the two networks can communicate directly with each other.

A configuration like that shown in Figure 1 may be actively and routinely using both frame relay networks A and B. For instance, a local area network (LAN) at site 1
10 may be set up to send all traffic from the accounting and sales departments to router A1 and send all traffic from the engineering department to router B1. This may provide a very rough balance of the traffic load between the routers, but it does not attempt to balance router loads dynamically in response to actual traffic and thus is not "load-balancing" as that term is used herein.

15 Alternatively, one of the frame relay networks may be a backup which is used only when the other frame relay network becomes unavailable. In that case, it may take even skilled network administrators several hours to perform the steps needed to switch the traffic away from the failed network and onto the backup network. In general, the necessary Private Virtual Circuits (PVCs) must be established, routers at each site
20 must be reconfigured to use the correct serial links and PVCs, and LANs at each site 102 must be reconfigured to point at the correct router as the default gateway.

Although two private networks are shown in Figure 1, three or more such networks could be employed, with similar considerations coming into play as to increased

reliability, limits on load-balancing, the efforts needed to switch traffic when a network fails, and so on. Likewise, for clarity of illustration Figure 1 shows only two sites, but three or more sites could communicate through one or more private networks.

Figure 2 illustrates a prior art configuration in which data is normally sent
5 between sites 102 over a private network 106. A failover box 202 at each site 102 can detect failure of the network 106 and, in response to such a failure, will send the data instead over an ISDN link 204 while the network 106 is down. Using an ISDN link 204 as a backup is relatively easier and less expensive than using another private network 106 as the backup, but generally provides lower throughput.

10 Figure 3 illustrates prior art configurations involving two private networks for increased reliability, in the sense that some of the sites in a given government agency or other entity 302 can continue communicating even after one network goes down. For instance, if a frame relay network A goes down, sites 1, 2, and 3 will be unable to communicate with each other but sites 4, 5, and 6 will still be able to communicate
15 amongst themselves through frame relay network B. Likewise, if network B goes down, sites 1, 2, and 3 will still be able to communicate through network A. Only if both networks go down at the same time would all sites be completely cut off. Like the Figure 1 configurations, the Figure 3 configuration uses two private networks. Unlike Figure 1, however, there is no option for switching traffic to another private network when one
20 network 106 goes down, although either or both of the networks in Figure 3 could have an ISDN backup like that shown in Figure 2. Note also that even when both private networks are up, sites 1, 2, and 3 communicate only among themselves; they are not connected to sites 4, 5, and 6.

Figure 4 illustrates a prior art response to the incompatibility of frame relay networks of different carriers. A special "network-to-network interface" (NNI) 402 is used to reliably transmit data between the two frame relay networks A and B. NNIs are generally implemented in software at carrier offices. Note that the configuration in Figure 4 does not provide additional reliability by using two frame relay networks 106, because those networks are in series rather than in parallel. If either of the frame relay networks A, B in the Figure 4 configuration fails, there is no path between site 1 and site 2; adding the second frame relay network has not increased reliability. By contrast, Figure 1 increases reliability by placing the frame relay networks in parallel, so that an alternate path is available if either (but not both) of the frame relay networks fails. Someone of skill in the art who was looking for ways to improve reliability by putting networks in parallel would probably not consider NNIs pertinent, because they are used for serial configurations rather than parallel ones, and adding networks in a serial manner does not improve reliability.

It would be an advancement in the art to provide another alternative for increasing reliability by configuring private networks in parallel, especially if other benefits are also provided. Such improvements are disclosed and claimed herein.

BRIEF SUMMARY OF THE INVENTION

The present invention provides tools and techniques for accessing multiple independent frame relay networks and/or point-to-point (e.g., T1 or T3) network connections in a parallel network configuration. In some embodiments a controller according to the invention comprises a site interface connecting the controller to a site, at

least two private network interfaces, and a packet path selector which selects between private network interfaces according to a specified criterion. The controller receives a packet through the site interface and sends the packet through the private network interface that was selected by the packet path selector. The controller's packet path selector selects between private network interfaces according to various criteria, such as (a) a load-balancing criterion that promotes balanced loads on devices that carry packets after the packets leave the selected private network interfaces; (b) a reliability criterion that promotes use of devices that will still carry packets after the packets leave the selected private network interfaces, when other devices that could have been selected are not functioning, and (c) a security criterion that promotes use of multiple private networks to carry different pieces of a given message so that unauthorized interception of packets on fewer than all of the networks used to carry the message will not provide the total content of the message. Some controller embodiments include only two private network interfaces, while others have three or more private network interfaces, each of which is selectable by the packet path selector. The private network interfaces may connect to a User-to-Network Interface, or they may comprise network-specific interface means of the type found in frame relay network routers.

One method of the invention for combining connections for access to multiple parallel frame relay and/or point-to-point networks, comprises the steps of: obtaining a controller, the controller comprising a site interface, at least two private network interfaces, and a packet path selector which selects between private network interfaces according to a specified criterion; connecting the controller site interface to a site to receive packets from a computer at the site; connecting a first private network interface of

the controller to a first private network; connecting a second private network interface of the controller to a second private network which is parallel to and independent of the first private network; and sending a packet to the site interface which then sends the packet through a private network interface selected by the packet path selector. The criterion
5 used by the packet path selector may be a load-balancing criterion, a reliability criterion, and/or a security criterion.

Another method for combining connections for access to multiple independent parallel frame relay or point-to-point networks comprises the steps of: sending a packet to a site interface of a controller, the controller comprising the site interface which receives
10 packets, at least two private network interfaces, and a packet path selector which selects between private network interfaces according to a specified criterion; and specifying the criterion for use by the packet path selector, wherein the specified criterion is one of: a security criterion, a reliability criterion, a load-balancing criterion. In one variation, the step of sending a packet to the controller site interface is repeated as multiple packets are
15 sent, the step of specifying a criterion specifies a security criterion, and the controller sends different packets of a given message to different frame relay networks.

Other features and advantages of the invention will become more fully apparent through the following description.

20

BRIEF DESCRIPTION OF THE DRAWINGS

To illustrate the manner in which the advantages and features of the invention are obtained, a more particular description of the invention will be given with reference to the

attached drawings. These drawings only illustrate selected aspects of the invention and its context. In the drawings:

Figure 1 is a diagram illustrating a prior art approach having frame relay networks configured in parallel for increased reliability for all networked sites, in configurations
5 that employ manual switchover between the two networks in case of failure.

Figure 2 is a diagram illustrating a prior art approach having a frame relay network configured in parallel with an ISDN network link for increased reliability for all networked sites.

Figure 3 is a diagram illustrating a prior art approach having independent frame
10 relay networks, with each network connecting several sites but little or no communication between the networks.

Figure 4 is a diagram illustrating a prior art approach having frame relay networks configured in series through a network-to-network interface, with no consequent increase in reliability because the networks are in series rather than in parallel.

15 Figure 5 is a diagram illustrating generally configurations of the present invention, in which two or more private networks are placed in parallel for increased reliability for all networked sites, without requiring manual traffic switchover, and with the option in some embodiments of load balancing between the networks and/or increasing security by transmitting packets of a single logical connection over different private networks.

20 Figure 6 is a diagram further illustrating the present invention, in which three sites can communicate over two parallel private networks.

Figure 7 is a diagram further illustrating a multiple private network access controller of the present invention, which comprises a component tailored to each private

network to which the controller connects, and a path selector in the controller which uses one or more of the following as criteria: private network status (up/down), private network load, use of a particular private network for previous packets in a given logical connection or session.

5 Figure 8 is a flowchart illustrating methods of the present invention for sending packets over multiple parallel independent private networks for enhanced reliability, load balancing and/or security.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

10 The present invention relates to methods, systems, and configured storage media for connecting sites over multiple independent parallel private networks such as frame relay networks and/or point-to-point network connections. "Multiple" networks means two or more such networks. "Independent" means routing information need not be shared between the networks. "Parallel" does not rule out the use of NNIs and serial networks,
15 but it does require that at least two of the networks in the configuration be in parallel so that alternate data paths through different private networks are present. "Frame relay networks" or "private networks" does not rule out the use of an ISDN link or other backup for a particular frame relay or point-to-point private network, but it does require the presence of multiple such networks – Figure 2, for instance, does not meet this
20 requirement.

Figure 5 illustrates generally configurations of the present invention involving frame relay networks; comments made here also apply to similar configurations involving point-to-point networks, or both types (frame relay and point-to-point) of private network.

Two or more frame relay networks 106 are placed in parallel between two or more sites 102. Access to the frame relay networks 106 at each site is through an inventive controller 502. The system containing the controllers 502 provides point-to-point connectivity between the sites 102. Additional controllers 502 may be used at each location, to provide
5 a switched connection system with no single point of failure.

Unlike the configuration shown in Figure 1, the inventive configuration in Figure 5 does not require manual intervention by network administrators to coordinate traffic flow over the parallel networks 106. The networks 106 are independent of each other. When one attached network fails, the failure is sensed by the controller 502 and traffic is
10 automatically routed through one or more other frame relay networks. Unlike the configuration in Figure 2, the inventive configuration combines two or more frame relay networks 106. Unlike the configuration in Figure 4, the inventive configuration requires two or more frame relay networks 106 be placed in parallel (although additional networks may also be placed in series). Unlike the configuration in Figure 3, the inventive
15 configuration does not merely partition sites between unconnected networks – with the invention, most or all of the connected sites get the benefit of parallel networks, so they can continue transceiving even if one of the networks goes down.

Another difference between the inventive approach and prior approaches may also be noted here, namely, the narrow focus of some prior art on reliability differs from the
20 present document's broader view, which considers load balancing and security as well as reliability. Configurations like those shown in Figure 2 are directed to reliability (which is also referred to by terms such as "fault tolerance", "redundancy", "backup", "disaster recovery", "continuity", and "failover"). That is, one of the network paths (in this case,

the one through the frame relay network) is the primary path, in that it is normally used for most or all of the traffic, while the other path (in this case, the one through the ISDN link) is used only when that primary path fails. Although the inventive configurations can be used in a similar manner, with one frame relay network being on a primary path and
5 the other network(s) being used only as a backup when that first network fails, the inventive configurations also permit concurrent use of two or more frame relay networks. With concurrent use, elements such as load balancing between frame relay networks, and increased security by means of splitting pieces of a given message between frame relay networks, which are not considerations in the prior art of Figure 2, become possibilities in
10 some embodiments of the present invention.

In general, the different frame relay or other private networks 106 will be provided by different carriers (WorldCom, AT&T, Qwest, etc.). In such cases, each frame relay network 106 typically operates on its own distinct clock. In some embodiments, the controller 502 sends traffic over all frame relay networks 106 to which it is connected, for
15 load balancing and/or enhanced security. In other embodiments or situations, the controller 502 prefers a particular network 106, and uses the other network(s) as backup in case the preferred network 106 becomes unavailable.

In some embodiments, a frame relay network C at a location 3 is connected to a controller 502 for a location 1 but is not necessarily connected to the controller 502 at
20 another location 2. In such cases, a packet from location 3 addressed to location 2 can be sent over network C to the controller at location 1, which can then redirect the packet to location 2 by sending it over network A or network B. That is, controllers 502 are

preferably, but not necessarily, provided at every location that can send packets over the parallel independent networks 106 of the system.

In some embodiments, the controller 502 at the receiving end of the network connection between two sites A and B has the ability to re-sequence the packets. This means that if the lines are of dissimilar speeds or if required by security criteria, the system can send packets out of order and re-sequence them at the other end. Packets may be sent out of sequence to enhance security, to facilitate load-balancing, or both. The TCP/IP packet format includes space for a sequence number, which can be used to determine proper packet sequence at the receiving end (the embodiments are dual-ended, with a controller 502 at the sending end and another controller 502 at the receiving end). The sequence number (and possibly more of the packet as well) can be encrypted at the sending end and then decrypted at the receiving end, for enhanced security.

Figure 6 further illustrates the present invention, in a particular configuration in which three sites 102 can communicate over two parallel independent frame relay networks 106; two or more point-to-point networks could be used similarly, as could a mixture of frame relay and point-to-point networks. In one such configuration, sites 1, 2, and 3 are connected via frame relay clouds 106. Routers 1, 2, and 3 are connected to frame relay cloud A, and routers 4, 5, and 6 are connected to frame relay cloud B. The WAN ports of the routers 104 on each frame cloud 106 are configured to form a single subnet. Virtual circuits (VCs) exist between site 1 and site 2, between site 2 and site 3, and between site 3 and site 1, on each of the clouds 106. A controller 502 is connected to each pair of routers 104 at each location to provide at least reliability through redundancy.

In operation, the controller 502 on each location is provided with a configuration file or other data structure containing a list of all the LAN IP addresses of the controllers 502 at the locations, and their subnet masks. Each controller 502 keeps track of available and active connections to the remote sites 102. If any of the routes are unavailable, the controller 502 preferably detects and identifies them. When a controller 502 receives IP traffic to any of the distant networks, the data is sent on the active connection to that destination. If all connections are active and available, the data load is preferably balanced across all the routers 104. If any of the VCs (or point-to-point connections) are unavailable, or any of the routers 104 are down, the traffic is not forwarded to that router; when the routes become available again, the load balancing across all active routes preferably resumes.

In some embodiments, load balancing is not the only factor considered when the controller 502 determines which router 104 should receive a given packet. Security may be enhanced by sending packets of a given message over two or more networks 106. Even if a packet sniffer or other eavesdropping tool is used to illicitly obtain data packets from a given network 106, the eavesdropper will thus obtain at most an incomplete copy of the message because the rest of the message traveled over a different network 106. Security can be further enhanced by sending packets out of sequence, particularly if the sequence numbers are encrypted.

Figure 7 is a diagram further illustrating a multiple frame relay and/or point-to-point network access controller 502 of the present invention. A site interface 702 connects the controller 502 to the LAN at the site 102. This interface 702 can be

implemented, for instance, as any local area network interface, like 10/100Base-T ethernet, gigabit ATM or any other legacy or new LAN technology.

The controller 502 also includes a packet path selector 704, which may implemented in custom hardware, or implemented as software configuring semi-custom or general-purpose hardware. The path selector 704 determines which path to send a given packet on. In the configuration of Figure 6, for instance, the path selector in the controller at location 1 selects between a path through router 1 and a path through router 4. In different embodiments and/or different situations, one or more of the following criteria may be used to select a path for a given packet, for a given set of packets, and/or for packets during a particular time period:

- Redundancy: do not send the packet(s) to a path through a network 106, a router 104, or a connection that is apparently down. Instead, use devices (routers, network switches, bridges, etc.) that will still carry packets after the packets leave the selected network interfaces, when other devices that could have been selected are not functioning. Techniques and tools for detecting network path failures are generally well understood, although their application in the context of the present invention is believed to be new.
- Load-balancing: send packets in distributions that balance the load of a given network, router, or connection relative to other networks, routers, or connections available to the controller 502. This promotes balanced loads on one or more of the devices (routers, frame relay switches) that carry packets after the packets leave the selected network interfaces. Load-balancing may be done through an algorithm as simple as a modified round-robin approach which places the next

packet on the next available line, or it may involve more complex algorithms that attempt to measure and track the throughput, latency, and/or other performance characteristics of a given link or path element. Load-balancing is preferably done on a per-line basis, as opposed to prior art approaches which use a per-department and/or per-router basis for dividing traffic. Load-balancing algorithms in general are well understood, although their application in the context of the present invention is believed to be new.

- Security: divide the packets of a given message (session, file, web page, etc.) so they travel over different networks 106. This promotes the use of multiple frame relay networks to carry different pieces of a given message, so that unauthorized interception of packets on fewer than all of the networks used to carry the message will not provide the total content of the message. Dividing message packets between networks 106 for better security may be done in conjunction with load balancing, and may in some cases be a side-effect of load-balancing. But load-balancing can be done on a larger granularity scale than security, e.g., by sending one entire message over network A and the next entire message over network B. Security may thus involve finer granularity than load balancing, and may even be contrary to load balancing in the sense that dividing up a message to enhance security may increase the load on a heavily loaded path even though a more lightly loaded alternate path is available and would be used for the entire message if security was not sought by message-splitting between networks. Other security criteria may also be used, e.g., one network 106 may be viewed as more secure than another, encryption may be enabled, or other security measures may be taken.

The controller 502 also includes two or more private network interfaces 706, namely, so there is at least one interface 706 per private network 106 to which the controller 502 controls access. Each interface 706 can be implemented as a direct interface 706 or as an indirect interface 706; a given embodiment may comprise only

5 direct interfaces 706, may comprise only indirect interfaces 706, or may comprise at least one of each type of interface. A direct interface 706 may be implemented, for instance, as a direct frame relay connection over land line or wireless or network interfaces to which the frame relay routers can connect, or as a point-to-point interface to a dedicated T1, T3, or wireless connection. One suitable implementation includes a standard Ethernet card,

10 which connects to an external frame relay User-Network Interface (UNI) in a router of a network 106. UNIs generally are known in the art. One indirect interface 706 effectively makes part of the controller 502 into a UNI by including in the interface 706 the same kind of special purpose hardware and software that is found on the frame relay network side (as opposed to the UNI side) of a frame relay network router. Such an indirect frame

15 relay network interface 706 is tailored to the specific timing and other requirements of the frame relay network to which the indirect interface 706 connects. For instance, one indirect interface 706 may be tailored to a Qwest frame relay network 106, while another indirect interface 706 in the same controller 502 is tailored to a UUNet network 106. The indirect interface 706 may connect to the frame relay network 106 over fiber optic, T1,

20 wireless, or other links. In short, a direct interface 706 relies on special purpose hardware and connectivity/driver software in a router, to which the direct interface 706 of the controller 502 connects through a UNI. By contrast, an indirect interface 706 includes such special purpose hardware and connectivity/driver software inside the controller 502

itself. In either case, the controller provides packet switching capabilities for at least redundancy without manual switchover, and preferably for dynamic load-balancing between lines as well. The controller 502 in each case also optionally includes memory buffers in the site interface 702, in the path selector 704, and/or in the network interfaces
5 706.

An understanding of methods of the invention will follow from understanding the invention's devices, and vice versa. For instance, from Figures 5-7, one may ascertain methods of the invention for combining connections for access to multiple parallel private networks 106, such as frame relay networks. One method begins by obtaining a controller
10 502. The controller comprises (a) a site interface 702, (b) at least two network interfaces 706 tailored to particular frame relay networks 106 for operation as though part of a network-to-network interface in a serial network configuration, and (c) a packet path selector 704 which selects between network interfaces 706 according to a specified criterion. Path selection criteria may be specified by configuration files, hardware jacks or
15 switches, ROM values, remote network management tools, or other means. One then connects the site interface 702 to a site 102 to receive packets from a computer (possibly via a LAN) at the site 102. Likewise, one connects a first network interface 706 to a first router 104 for routing packets to a first frame relay network 106, and a second network interface 706 to a second router 104 for routing packets to a second frame relay network
20 106. A third, fourth, etc. frame relay network 106 may be similarly connected to the controller 502 in some embodiments and/or situations. The connected frame relay networks 106 are parallel to one another (not serial, although additional networks not directly connected to the controller 502 may be serially connected to the networks 106).

The connected frame relay networks 106 are independent of one another, in that no routing information need be shared between them, to make them parallel (NNIs can still be used to connect networks in serial to form a larger independent and parallel network). A mistake in the routing information for one network 106 will thus not affect the other
5 network 106. After the connections are made (which may be done in a different order than recited here), one sends a packet to the site interface 702, which then sends the packet through the one (or more – copies can be sent through multiple networks 106) network interface 706 that was selected by the packet path selector 704.

Figure 8 is a flowchart further illustrating methods of the present invention, which
10 send packets over multiple parallel independent private networks 106 for enhanced reliability, load balancing and/or security; frame relay networks are used as an example, but point-to-point networks may be similarly employed. During a connection forming step 802, at least one virtual circuit is obtained between two sites 102. If the frame relay networks 106 will be used concurrently, the controllers 502 provide a connection which
15 comprises multiple conventional virtual circuits, since two or more networks may (or will) carry packets during the step 802 connection. The controller 502 then checks the status of each connection and updates the information for available communication paths.

During a packet receiving step 804, the controller 502 at a given location receives a packet to be sent from that location to another site 102. In some cases, multiple packets
20 may be received in a burst. The packet comes into the controller 502 through the site interface 702.

During a path selecting step 806, the path selector 704 selects the path over which the packet will be sent; selection is made between at least two paths, each of which goes

over a different network 106 than the other. The networks 106 are independent parallel frame relay networks. This path selecting step 806 may be performed once per packet, or a given selection may pertain to multiple packets. Path selection 806 is shown as following packet receipt 804, but in some embodiments and/or some situations, it may precede packet receipt 804. More generally, the steps illustrated and discussed in this document may be performed in various orders, including concurrently, except in those cases in which the results of one step are required as input to another step. Likewise, steps may be omitted unless required by the claims, regardless of whether they are expressly described as optional in this Detailed Description. Steps may also be repeated, or combined, or named differently.

As indicated, the path selection may use 808 load balancing as a criterion for selecting a path, use 810 network 106 status (up/down) and other connectivity criteria (e.g., router status, connectivity status) as a criterion for selecting a path, and/or use 812 division of packets between networks 106 for enhanced security as a criterion for selecting a path. These steps may be implemented in a manner consistent with the description above of the path selector 704 given in the discussion of Figure 7. More generally, unless it is otherwise indicated, the description herein of systems of the present invention extends to corresponding methods, and vice versa.

The description of systems and methods likewise extend to corresponding computer-readable media (e.g., RAM, ROM, other memory chips, disks, tape, Iomega ZIP or other removable media, and the like) which are configured by virtue of containing software to perform an inventive method, or software (including any data structure) which is uniquely suited to facilitate performance of an inventive method. Articles of

manufacture within the scope of the present invention thus include a computer-readable storage medium in combination with the specific physical configuration of a substrate of the computer-readable storage medium, when that substrate configuration represents data and/or instructions which cause one or more computers to operate in a specific and
5 predefined manner as described and claimed herein.

During a packet transmission step 814, the packet is sent on the selected 806 path. This is done by sending the packet over the network interface 706 for the path selected. As indicated in Figure 8, the method may then loop back to receive 804 the next packet, select 806 its path, send 814 it, and so on. As noted, other specific method instances are
10 also possible. One example is the inventive method in which load balancing or reliability criteria cause an initial path selection to be made 806, and then a loop occurs in which multiple packets are received 804 and then sent 814 over the selected path without repeating the selecting step 806 for each receive 804 – send 814 pair. Note that some
15 embodiments of the invention permit packets of a given message to be sent over different networks 106, thereby enhancing 812 security. The PVCs are in general always connected, but an ending step 816 may be performed during an orderly shutdown for diagnostic or upgrade work, for instance.

Summary

20 The present invention provides methods and devices for placing frame relay and other private networks in parallel, thereby providing redundancy without requiring manual switchover in the event of a network failure. Load-balancing between lines and/or between networks may also be performed. For instance, the invention can be used to

provide reliable, efficient, and secure point-to-point connections for private networks

102. Some prior art approaches require network reconfiguration each time a frame relay
circuit fails, and some have complex router configurations to handle load balancing and
network failures. This requires substantial effort by individual frame relay network
5 customers to maintain connectivity, and they will often receive little or no help from the
frame relay carriers. Instead, well-trained staff are needed at each location, as are
expensive routers. By contrast, these requirements are not imposed by the present
invention.

As used herein, terms such as "a" and "the" and item designations such as
10 "connection" or "network" are generally inclusive of one or more of the indicated item. In
particular, in the claims a reference to an item normally means at least one such item is
required.

The invention may be embodied in other specific forms without departing from its
essential characteristics. The described embodiments are to be considered in all respects
15 only as illustrative and not restrictive. Headings are for convenience only. The scope of
the invention is, therefore, indicated by the appended claims rather than by the foregoing
description. All changes which come within the meaning and range of equivalency of the
claims are to be embraced within their scope.

What is claimed and desired to be secured by patent is:

20

1. A controller which controls access to multiple independent private networks in a parallel network configuration, the controller comprising:

a site interface connecting the controller to a site;

at least two private network interfaces; and

5 a packet path selector which selects between private network interfaces according to a specified criterion;

wherein the controller receives a packet through the site interface and sends the packet through the private network interface that was selected by the packet path selector.

10

2. The controller of claim 1, wherein the controller control access to multiple independent frame relay networks, and each of the at least two private network interfaces comprises a frame relay network interface.

15 3. The controller of claim 1, wherein the packet path selector selects between private network interfaces according to a load-balancing criterion, thereby promoting balanced loads on devices that carry packets after the packets leave the selected private network interfaces.

20 4. The controller of claim 1, wherein the packet path selector selects between private network interfaces according to a reliability criterion, thereby promoting use of devices that will still carry packets after the packets leave the selected private network interfaces, when other devices that could have been selected are not functioning.

5. The controller of claim 1, wherein the packet path selector selects between private network interfaces according to a security criterion, thereby promoting use of multiple private networks to carry different pieces of a given message so that unauthorized interception of packets on fewer than all of the private networks used to carry the message will not provide the total content of the message.

6. The controller of claim 1, wherein the controller sends packets out of sequence over the parallel private networks.

10

7. The controller of claim 6, wherein the controller places an encrypted sequence number in at least some of the packets which are sent out of sequence.

8. The controller of claim 1, wherein the controller comprises at least three frame relay network interfaces, each of which is selectable by the packet path selector.

15

9. The controller of claim 1, wherein the controller operates in a system providing at least one point-to-point connection.

20

10. The controller of claim 1, wherein the controller operates in a system providing connectivity over at least two frame relay networks from at least two carriers, each frame relay network operating on its own clock which is different from the clock of the other frame relay network.

11. The controller of claim 1, wherein each private network interface is an indirect interface tailored to a particular type of frame relay network.

5 12. The controller of claim 1, wherein each private network interface is a direct interface comprising an Ethernet card.

13. A method for combining connections for access to multiple parallel private networks, the method comprising the steps of:

10 obtaining a controller, the controller comprising a site interface, at least two private network interfaces, and a packet path selector which selects between private network interfaces according to a specified criterion; connecting the controller site interface to a site to receive packets from a computer at the site;

15 connecting a first private network interface of the controller to a first private network;

connecting a second private network interface of the controller to a second private network which is parallel to and independent of the first private network;

and

20 sending a packet to the site interface which then sends the packet through a private network interface selected by the packet path selector.

14. The method of claim 13, wherein the private networks are frame relay networks.

15. The method of claim 13, further comprising the step of specifying the
5 criterion for use by the packet path selector, wherein the specified criterion is a load-balancing criterion.

16. The method of claim 13, further comprising the step of specifying the
10 criterion for use by the packet path selector, wherein the specified criterion is a reliability criterion.

17. The method of claim 13, further comprising the step of specifying the
15 criterion for use by the packet path selector, wherein the specified criterion is a security criterion.

18. The method of claim 13, wherein at least one of the steps connecting a
private network interface of the controller connects the controller to a User-to-Network
Interface in a router of a frame relay network.

20 19. A method for combining connections for access to multiple independent
parallel frame relay networks, the method comprising the steps of:
sending a packet to a site interface of a controller, the controller comprising the
site interface which receives packets, at least two network interfaces, and a

packet path selector which selects between network interfaces according to
a specified criterion; and

specifying the criterion for use by the packet path selector, wherein the specified
criterion is one of: a security criterion, a reliability criterion, a load-
5 balancing criterion.

20. The method of claim 19, wherein the step of sending a packet to the
controller site interface is repeated as multiple packets are sent, the step of specifying a
criterion specifies a security criterion, and the controller sends different packets of a given
10 message to different frame relay networks.

21. The method of claim 19, further comprising the step of sensing failure of
one of the parallel frame relay networks and automatically sending traffic through at least
one other parallel frame relay network.

15

20

ABSTRACT

Methods, configured storage media, and systems are provided for communications using two or more frame relay or point-to-point networks in parallel to provide load balancing across network connections, greater reliability, and/or increased security. A controller provides access to two or more private networks in parallel, through direct or indirect network interfaces. When one attached network fails, the failure is sensed by the controller and traffic is routed through one or more other private networks. When all attached networks are operating, the controller preferably balances the load between them.

10

\pc0

AL3

USING MULTIPLE LINKS TO INTERCONNECT LANs AND PUBLIC CIRCUIT SWITCHED DATA NETWORKS

T. Liao¹, M. Noosong¹, Y. Liang² AND D. Seret¹

¹ ENST, 46 rue Barrault, 75634 Paris, France
² Bull S.A., 94 av. Gambetta, 75020 Paris, France

Abstract - This paper presents the design of a gateway that interconnects a TCP/IP-based LAN and a public circuit switched data network (CSDN) using multilinks. In the interconnection of high speed LANs and relatively slow public networks, performance degradation is a significant issue. In order to improve the throughput, simultaneously controlling n data links for data transfer is an immediate solution. In this way, the transmission rate can be increased roughly by n times higher over a single link. The use of multilinks also enables the gateway to handle multiple communications at the same time. A simple multilink control procedure is defined to control from one to multiple links for data transfer. Some other techniques for the design of the gateway are also presented.

communications at the same time.

This paper is organized in the following way. Section 2 describes the architecture of the gateway. The multilink control problems are discussed in section 3. Section 4 presents the techniques used in the realization. The paper ends with some conclusions.

II. INTERCONNECTION SCHEME

The interconnection scheme is shown in Fig.1. The users of CSDN must follow X.21 procedure and use a call number (8-digit for Transcom) to establish connections at layer 1. From layer 2 to layer 7, the users are free to choose any protocols that meet their needs. Communication between two hosts on a LAN is supported by TCP/IP protocol suite³.

I. INTRODUCTION

Local area networks (LANs) seem to be a major answer to the problem of open systems interconnection in a limited geographic area such as an office building, a campus, a manufacturing plant etc. But in some circumstances they are not adequate. As an example; we can not extend a LAN in the main site to a remote site. The interconnection of LANs and public packet switched networks (PSNs) or public switched telephone networks (PSTNs) provides a way to extend the range of LANs^{1,2}. Another alternative is the interconnection of remote LANs via public circuit switched data networks (CSDNs). The latter also offers the possibility of simultaneously using a number of data links to enhance the performance and allows users to economically establish connections between two sites.

This paper presents the design of a gateway that connects a TCP/IP-based LAN (internet) to a CSDN (e.g. Transcom that offers 64 kb/s links in France) by using multiple physical links. Our main objectives are to describe the techniques used and to discuss the problems encountered in controlling n data links. In addition to the interconnection of two remote LANs, this gateway also allows a host connected to a CSDN to communicate with another host connected to a LAN. The function of multilink control can either increase the transmission rate roughly by n times higher over a single link or allow the gateway to handle multiple


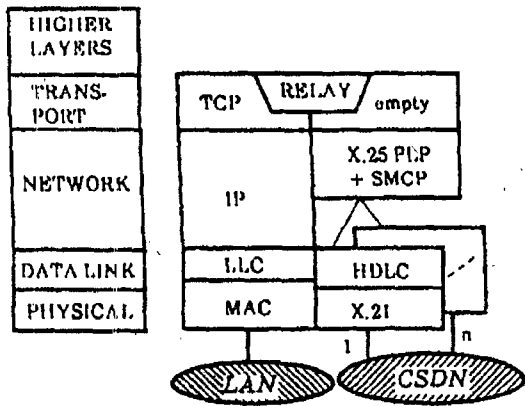


Fig.1. The interconnection scheme

As for the application processes, what they need is a reliable transport connection. First, the relay function is placed on TCP regarding that the services provided by IP are connectionless and not reliable. On the CSDN side, given that n ($n > 1$) physical links are available, the gateway must include the following functions (see Fig.1):

- 1) In order to improve the throughput, the gateway must be able to use multiple physical links to establish a logical connection between two hosts.
- 2) In some cases, one link is sufficient for the quality of service and there may be simultaneously m ($1 < m \leq n$) connection requests for different destinations. By using a single link for a connection, the gateway must be able to establish m independent connections at the same time. For instance, if host A is connected to



SMCP : Simple Multilink Control Procedure

Fig.2. The gateway architecture

host C, host B can still connect to host D.

- 3) The gateway must be able to support a mechanism like virtual circuits. For instance, hosts A and B can connect to host C at the same time.

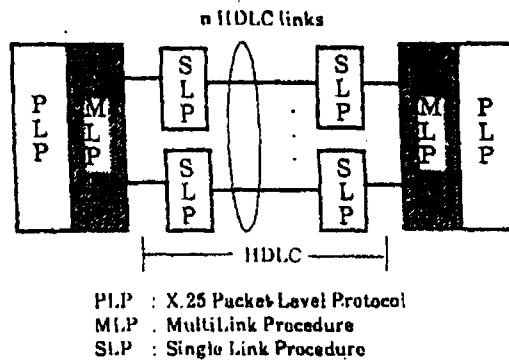
Thus the protocols adopted for the CSDN are X.25, HDLC and X.21. In addition, some multilink control function should be included in the gateway. This function is placed in the layer 3 for several reasons discussed in detail in the next section. The gateway architecture is shown in Fig.2 in which the relay function receives data from TCP and passes them to X.25 and vice versa.

III. MULTILINK CONTROL

A. Motivation and Definition

The interconnection of a high speed LAN and a relatively slow public network often introduces performance degradation to the supported applications. With variegated applications and office automation, e.g., large file transfer, image applications, facsimile, integration of data and voice etc., multimedia communications become indispensable. To tackle the performance problem which is critical for the applications mentioned above, an immediate solution in the interim is the use of multilinks. The problems issue of the multilink control include the connection establishment, packet resequencing and error control. Since the multilink control function can be included in one of the OSI seven layers, several solutions are thus possible.

The first is that in the Physical Layer, n links are grouped as a single link and several additional links, typically $(n+7)/8$, are used for control and signaling. This single link with X.21 interface is transparent for higher layers^[4]. Another existing solution is the MultiLink Procedure (MLP) which is defined as a sub-layer of X.25^[5]. The MLP controls n Single Link Procedures (SLPs) which are described as HDLC procedures (Fig.3). For each packet to be sent, the MLP adds two octets of control information before the header of the packet for sequence control and dispatches the packet to one of SLPs according to a



PLP : X.25 Packet Level Protocol
MLP : MultiLink Procedure
SLP : Single Link Procedure

Fig.3. MultiLink Procedure

specific algorithm. Each Single Link Procedure works independently. The third solution is the ISO Transport Protocol class 4 (TP4) in which the multilink control function is defined as "splitting"^[6]. The last approach is in the dynamic use of multiple ISDN B-channels where the multilink control function is placed in the layer 2^[12].

The performance studies show that the MLP usually offers higher throughput than the others, if an appropriate method for packet distribution is adopted^[7,8]. The MLP is not feasible in certain cases, for example, there is no function of controlling n links independently (i.e., each link is connected to a different destination). Therefore the MLP is not suitable for our gateway. However, using TP4 on top of X.25 is also redundant and consequently causes loss of performance. In short, the existing solutions do not satisfy our needs. A new procedure - Simple Multilink Control Procedure (SMCP) is defined here, as shown in Fig.4, which contains two parts: SMCP-h (high part) procedure and SMCP-l (low part) procedure between which is X.25 PLP. The particular functions of the SMCP are described below:

- a) the SMCP offers multiple network connections over multiple data links to a transport entity.
- b) the SMCP regroups n data links to establish a network connection.
- c) the SMCP also performs packet sequencing and packet resequencing, error processing.
- d) Identification of network connections and so on.

The advantages of the SMCP are: first, it is able to control each link independently; second, it keeps the efficiency of the MLP.

B. Identification of Network Connection

Within the SMCP, an "SMCP context" is maintained which contains information about each network connection. This context, shared by the SMCP-h and SMCP-l, contains a table of the entries consisting of: X.25 source and destination addresses, virtual circuit number, SLP number (each SLP is assigned a number), etc. The internal identification of network connection within the SMCP is a combination of SLP numbers and virtual circuit number. Two working modes are defined in the SMCP. The first is the multilink (ML) mode in which multiple data links

SMCP

SM

are u
second
data l
in the

C. Co

At
estab
layer
reque
the
inform
PLP f
data l

Th
proces
numb
l proc
to the
descri


1

CM

sei

SF

FO
AS CI
NPDI



SMCP-h : Simple Multilink Control Procedure - h (high) part
 SMCP-l : Simple Multilink Control Procedure - l (low) part

Fig.4. Multi-Link Control Protocol

are used to support a network connection and the second is the single link (SL) mode in which a single data link is used. It is assumed that the SMCP works in the ML mode hereafter.

C. Connection Establishment and Liberation

At the calling SMCP procedure, the connection establishment for data transfer is solicited by a higher layer entity (e.g. transport). The SMCP-h procedure requests X.25 PLP to establish a network connection to the destination and stores the corresponding information in the context. In turn, if needed, X.25 PLP first requests the SMCP-l procedure to establish data links.

The SMCP-l procedure then requests one of SLP procedures to establish a data link (using X.21 call number). After the data link is established, the SMCP-l procedure sends a multilink service NPDU (2 octets) to the destination SMCP-l. The format of this NPDU is described as below (see Fig.5):




Fig.5. The multilink service NPDU

- CMD** : this field indicates the command code (2 bits).
 00 --- multilink service request
 01 --- multilink service response
- service** : this field indicates the number of links requested in the request NPDU or indicates the number of links agreed in the response NPDU.
- SF** : this field indicates whether the packet resequencing is required.
 00 --- resequencing required
 01 --- resequencing not required

For example, this NPDU may contain information as CMD=00, service=0...01000, SF=00. After this NPDU is sent, the SMCP-l procedure waits for a

response. If the received NPDU is longer than 2 octets, it is ignored. The destination SMCP-l may respond with an NPDU containing CMD=01, service=0...0100, SF=00, that is to say, it has only 4 links available and is agreed for packet resequencing. While the calling SMCP-l procedure receives this NPDU, it is aware of the number of links that can be established for the communication (4 links instead of 8 links it requests). Then the SMCP-l procedure requests other 3 SLPs to establish data links to the destination. After the 4 data links are established, the SMCP-l procedure sends a connection indication of data links to X.25.

Then X.25 PLP sends a call request packet to the SMCP-l procedure which delivers this packet to one of the SLPs. While the call request is confirmed, the SMCP-h procedure then returns a connection confirmation to the higher layer entity. At this moment, the network connection is established. If the virtual circuit can not be established, the SMCP-h returns a disconnection indication to the higher layer entity and clears the context. At the end of a communication, if a disconnection request is received from the higher layer entity, the SMCP-h entity requests X.25 PLP to disconnect the virtual circuit.

At the called SMCP, when the SMCP-l procedure receives a 2 octet length NPDU, it verifies whether the CMD field is 00 or not. If so, it responds with a NPDU containing CMD=01, service=min(the number of links requested, its number of links available), SF=received SF. After the concerning data links are established, it sends a connection indication of data links to X.25 PLP. While X.25 PLP receives a call request packet, it sends a connection indication to the SMCP-h procedure that transfers this indication to the higher layer entity and waits for a connection response. When the connection response is received, it then requests X.25 PLP to send a connection confirmation to its correspondent. At the end of a communication, the SMCP-h procedure should generally receive a disconnection indication from X.25 PLP.

D. Resequencing and Distribution Algorithms

Since the packets are transmitted through n data links which are possibly with different propagation delays and error rates, they are likely to be received not in the order in which they are sent. The packet resequencing needs to be done by the SMCP. Two octets of sequence control information are added to the NSDUs (Network Service Data Unit) for this purpose, as shown in Fig.6. Regarding that X.25 PLP and HDLC can ensure the delivering of packets, only two fields are defined in this control information. The SI (Sequence Indication) field indicates whether this packet should be reordered or not and the SNR (Sequence Number) field indicates the sequence number of the current NSDU. SI=0 signifies that the packet sequencing is required. The SMCP-h procedure is responsible for the packet resequencing.

As to the SMCP-l procedure, it is charged of distributing the packets to be transmitted to the SLPs according to a specific algorithm. Two algorithms can be used, one of which is the cyclic algorithm (C-




Fig.6. Packet resequencing

algorithm) and the other is the adaptive algorithm (A-algorithm). By the C-algorithm, the SMCP-1 starts packet transmission from the SLP number 1 and then sends next packet to the SLP number 2. While a packet is sent to the SLP number n, the SMCP-1 returns to the SLP number 1. The major disadvantage of the C-algorithm is that the throughput degrades in the case of heterogeneous physical links used, because all data links are charged with the same traffic in spite of their real capacities⁸¹. However, in most of cases, the physical links of a CSDN seem almost homogeneous, so that using C-algorithm can offer the expected performance. In the worst case where the n data links could have very different propagation delays and error rates, the A-algorithm can be applied to improve the throughput. By the A-algorithm, before sending a packet, the SMCP-1 must find an SLP which is ready for data transmission among the SLPs and then sends the packet to the ready SLP. Therefore, better the quality of a data link, higher the traffic the data link transports. The A-algorithm is adopted in the SMCP.

E. Change of Working Mode

The working mode is set either to multilink control or to single link control in the idle time (no communication being processed by the SMCP). During a communication, the working mode can not be changed. As shown in Table-1, if the ML mode is chosen, n will be assigned to the service field of multilink service request NPDU. If the SL mode is chosen, the value of this field is one. In the latter case, n network connections can be established at the same time. Communication between the SMCPs in different working modes is possible.

Table-1. Comparison of working modes

	service request field	number of connections authorized
SMCP in the SL mode	1	n
SMCP in the ML mode	n	1

IV. INTERCONNECTION TECHNIQUES

A. Addressing

Application processes are assumed to be attached to transport service access points (TSAPs). Each TSAP is a "socket style" combination of a port number

identifying an application in a host and a unique Internet address of a host in the internets. This makes necessary to use Internet address throughout the interconnection scheme. Therefore a class A Internet address should be assigned to each host connected to a CSDN.

Mapping between an Internet address and a CSDN call number can be done in the same way as the mapping between Internet address and physical address. Using the call number, the gateway and host on CSDN can establish physical connections. The source and destination "socket" identifiers (2-octet port number and 4-octet Internet address) must be carried in X.25 packet's address fields.

B. Data Transmission and Flow Control

To cope with the very different transmission rates on the LAN and on the CSDN, buffers for temporarily stocking data are required. There are 2 buffers in the gateway, a TCP buffer is for data transmission to the internet side and a PLP buffer to the CSDN side.

To transfer data from TCP to PLP, the data received from TCP is first buffered in the PLP buffer. If the length of the received data is larger than the packet length, the data must be divided into several NSDUs. Then these NSDUs are sent to the SMCP-h procedure. After all the NSDUs are acknowledged, the relay function can receive data from the TCP entity again. Data transfer from PLP to TCP follows the same principle. In our interconnection scheme, flow control in the LAN and in the CSDN are performed separately.

C. Implementation

The gateway has been implemented on an IBM AT compatible microcomputer. The communication between the gateway and hosts on the TCP/IP-based LAN is supported by a communication card from Excelan¹⁰¹ and the access to the CSDN is done by using several ANTILOPE cards¹¹¹. Two CSDN links were used during the experiment. The results show that the end-to-end transmission rate can reach approximately to 0.80x2x64 kb/s (across Transcom)⁸¹.

V. CONCLUSIONS

The designed gateway aims to interconnect LANs (internets) and public circuit switched data networks. In the interconnection of high speed LANs and relatively slow public networks, performance degradation is a significant issue. To address this problem, we have implemented a function that can simultaneously use n (n>1) data links for data transfer. Consequently, the throughput can be increased roughly by n times higher over a single link and thus will satisfy more applications. Another purpose of the use of n independent links is that in the case where one link is sufficient for the performance, the gateway can handle n independent communications at the same time. The multilink control protocol defined in layer 3 provides the way to effectively control from one to n links for data transfer.

unique makes ut the ternet ted to a

CSDN as the hysical d host. s. The tet port carried

on rates orarily s in the 1 to the

e data uffer. If an the several MCP-b ed, the entity e same control urately.

SM AT ication -based l from one by l links s show reach across

LANs works. s and mance s this at can data n be e link other in the nance, ndent tilink ay to asfer.

REFERENCES

[1] E.K. Chew, "Interworking of Local Area Networks and Public Networks", Telecommunication Journal of Australia, Vol.34, No.3, 1984, pp221-231.

[2] W. Johannsen, W. Lamersdorf and K. Reinhardt, "Architecture and Design of an Open Systems LAN/WAN Gateway", Pro. of Computer Networking Symposium 1988, pp112-119.

[3] Internet Protocol Transition Workbook, Network Information Center SRI International, March 1982.

[4] J.P. Temime, "Une Trame Numérique à 64 kb/s pour services multimédias", Echo des Recherches, 2ème trimestre 1986.

[5] ISO 8882, "Data Communication - X.25 Packet Layer Specification for Data Terminal Equipment", October 1982.

[6] ISO 8073, "Information Processing Systems - Open Systems Interconnection - Transport Protocol Specification", 1986.

[7] T. Nishizono, K. Kanemaki and J. Yano, "Multilink Protocol for Data Transfer", Review of the Electrical Communications Labs, Vol.33, No.5, 1985.


[8] Y. Liang, "The N×64 kb/s ISDN Services: performance of the protocols and realization of a gateway", ENST Ph.D Thesis, June 1989.

[9] T. Liao, "A Realization Method of the 2×64 kbps ISDN Service : Control of Two Transcom accesses", ENST Internal Report, 1989.

[10] The LAN WorkPlace™ Network Software for PC DOS, Socket Library Application Program Interface User's Guide, Excelan.

[11] Manuel d'Utilisation ANTLOPE, CISI Ingénierie Informatique.

[12] L. Altarah and S. Motard, "Dynamic Use of ISDN B Channels", Proceedings of ISDN in Europe, The Hague, April 1989.



114

[Back to News Home](#) | [Previous News Item](#) | [Next News Item](#)

FOR IMMEDIATE RELEASE
Contact: Industry Relations Manager
Telephone: (650) 969-8000
E-Mail: pr@coyotepoint.com

Coyote Point Systems Introduces the Equalizer Load-Balancing Router for High-Volume Web Sites

Sunnyvale, CA, September 8, 1997. Coyote Point Systems, Inc. today announced the release of the Equalizer, a load balancing router which optimizes Internet server performance and reliability. Equalizer can handle more than 40,000 Web requests a minute--ideal for busy Internet sites, Internet service providers and mission critical corporate Intranets.

Unlike conventional round-robin DNS based distribution, the Equalizer ties together multiple server machines to form "logical clusters." If one server fails, the Equalizer automatically compensates, insuring that clients never receive "server down" errors.

"Equalizer easily managed the load of our multiple sites" said Dean Gaudet, Director of Infrastructure at HotWired. "This is a technology with the capability of making my job a lot easier." The Equalizer features a broad array of advantages for high load Web sites:

- Equalizer detects failed servers and automatically redirects traffic.
- Equalizer controls the distribution of requests across servers and constantly re-distributes the load. Sites using Equalizer get the most out of their server resources.
- Equalizer allows system administrators to increase maximum capacity incrementally, as needed. Equalizer allows administrators to add inexpensive servers, rather than continuously upgrading to bigger, more expensive systems.
- Equalizer provides a greater degree of control over cluster resources than other load balancing products. The browser based administration tool provides a simple, powerful means to monitor cluster performance and adjust operating controls.
- Newly added servers start sharing the load instantly without the waiting period associated with DNS based load distribution. Webmasters can bring servers down for maintenance without service interruptions.
- Security features: Enhanced support for the SSL security protocol makes Equalizer ideal for internet commerce sites. Equalizer provides complete control over which server ports are accessible. Connections to specified addresses and ports may be redirected or dropped as required. Improved resistance to denial of service attacks is another benefit.
- Equalizer supports all popular TCP/IP based protocols.

"As Web traffic becomes critical to the success of companies, the level of service provided by these sites will be determined by the effectiveness of their hardware," said M. Phillip Roth, executive vice president of Coyote Point Systems.

"The Equalizer easily handles multi-million hit-per-hour sites. It provides a means by which Web masters can manage, maintain and scale their equipment in a cost effective manner. We provide the only Web-based management solution of this type in the industry."

Three Equalizer models are available, starting at \$6,200. Higher end models can handle greater load, more servers and more virtual clusters. A redundant backup option is also available, insuring that no single point of failure can interrupt service.

Coyote Point Systems is a privately held corporation located in Sunnyvale, California, specializing in Internet performance tools and technology. Further information is available at www.coyotepoint.com or by calling 650 969-8000. Email PR@coyotepoint.com Coyote Point

<http://www.coyotepoint.com/press-release.shtml>

9/29/98 4:15:29 PM

Systems, 530 Lawrence Expressway, Suite 341 Sunnyvale, 94086

[Home](#) [About Coyote Point](#) [Contact Information](#) [News](#) [Products](#) [Partners](#) [Support](#) [Site Index](#)

Copyright © 1998 Coyote Point Systems, Inc. All Rights Reserved.

A15

Network Address Translation Technical Discussion

Not everyone wants the actual details of a NAT. But, for those that do, here they are for our products.

Technology Consulting

By the way, we are happy to work with others that wish to add NAT functionality to their products. In the nearly two years we have been providing Network Address Translation, we have learned that there is much to it beyond the RFC. Let us help you understand the finer points and not-so-obvious issues in bringing your product to market.

Network Address Translation

The standard features of Network Address Translation are detailed in RFC 1631. All of those features are supported by our NATs, including:

- Conversion of IP Addresses
- Correct Handling of FTP and ICMP Messages

In addition, the Net NAT expands upon the NAT standard with:

- Dynamic Concentration of IP Addresses
- Conversion of TCP and UDP port numbers
- BSD Authentication Server
- Creation of Virtual Servers
- Support for Great New Applications

Conversion of IP Addresses

The main feature of an RFC 1631 NAT is to enable an organization to use the free IP Networks reserved in RFC 1597 while still permitting clients or servers on this network to access, or be accessed by, the public Internet. It does this through a mechanism that substitutes a globally registered IP Address into the source IP Address part of a message leaving the private network, and restores the private IP Address into the destination part of a reply message entering the private network. For example, assume a translate table of the following nature:

Name	Private IP Address	Public IP Address
ws12	192.168.16.12	204.116.73.1
ws26	192.168.16.26	204.116.73.2
ws27	192.168.16.27	204.116.73.3
ws59	192.168.16.59	204.116.73.4

A message originating at ws12 has 192.168.16.12 in the Source IP Address part of the message header. As it passes through the NAT to the Public Internet, the NAT substitutes 204.116.73.1 into that part of the header and recalculates the various message checksums. The message is then sent to the addressed host on the "outside" as though it originated from the public address. When a message arrives at the NAT from the Public Internet addressed to 204.116.73.1, the private IP Address of ws12 is substituted into the destination part of the message header, the checksums are recalculated, and the message is delivered to ws12.

Even though the four workstations in this example are spread across a wide section of the internal RFC 1597 Class C Network (192.168.16), their Public IP Addresses have been consolidated into a very small section of the external IP Network.

We Prefer Actual and Apparent

Instead of Private IP Addresses (or "reusable addresses" from RFC 1631) we prefer to call them "Actual Addresses." In the same light, we call the Public IP Addresses (non-reusables from RFC 1631) "Apparent Addresses." This seems much more understandable, and is less limiting. Many NATs are used in large Enterprise Networks that are not connected to the Public Internet, and so have no "non-reusable" addresses at all.

Correct Handling of FTP and ICMP Messages

Some of the NAT function would be much simpler if IP Addresses were only found in headers. But, several application protocols and the Internet Control Message Protocol carry IP Addresses in the message data. If a NAT is to work with these protocols, it must identify the protocol, find the embedded IP Address and fix it there, as well as in the header. Needless to say, the creation of protocols that do this should be discouraged or standardized to make the process more robust.

Dynamic Concentration of IP Addresses

An examination of the address translation table above will explain the first enhancement made to the RFC 1631 NAT. Instead of a fixed actual-to-apparent mapping, a pool of apparent addresses is created, and then dynamically assigned to the actual users. The assignment is temporary, so that apparent addresses can be used by other actual users.

Several other vendors of NAT products do this. We chose to do it differently. Please read on!

Conversion of TCP and UDP Port Numbers

The thought of tying up a large range of apparent IP Addresses seemed very silly to us, so we built the Net NAT to perform a different kind of concentration. Instead of translating just the IP Address, we also translate the TCP or UDP Port Numbers. These are also called service numbers. In our version of the address translation table, there is only one entry, and that is a single Apparent IP Address. And, instead of a pool of apparent addresses, there is a pool of "Apparent Ports."

So, when a workstation in the private network initiates a TCP session or a UDP exchange, the Net NAT assigns a port from the apparent port pool and then substitutes the apparent address and port for the actual address and port before sending the message to the public network. In a similar way, the apparent address and port in a message from the public network are replaced with the actual address and port before the message is passed to the workstation. The actual addresses may be RFC 1597 private addresses, or someone else's public addresses that were used before a connection to the Public Internet was anticipated, and must be kept to avoid the agony and expense of reconfiguring every computer in the network.

This permits the sharing of a single Apparent IP Address between a very large number of actual users. In practice, the number of users is virtually unlimited. The true limit is in the number of simultaneous sessions. That is limited by the size of the apparent port pool and the memory available for session context blocks.

BSD Authentication Server

When you connect to some BSD-derived unix hosts, they query your system to determine the identity of the user, before ever prompting you for a login. This typically wastes ten seconds of your time for each connection. Our customers wanted to know what they could do about this, and we responded by adding a server to our NAT code. We reply with a canned user identity so that the connection will complete right away.

Creation of Virtual Servers

The Net NAT can provide support for servers in the private network that need to be "seen" from the public network. Three modes are available for this function. We call them **Fixed Mode**, **Port Mode**, and **Mux Mode**

Fixed Server Mapping Mode

Fixed mode is the equivalent of plain RFC 1631 conversion. In this mode, a table defines a fixed relationship between apparent and actual addresses. No translation of port number is performed in this mode, since that isn't required.

Port Server Mapping Mode

Port mode is very similar to Fixed Mode, except that the fixed relationships are between combinations of apparent address and port and actual address and port. Thus, the port number is added to the translation process. Consider this new table:

Name	Actual Address and Port	Apparent Address and Port
ser12	192.168.16.12 80	204.116.73.1 80
ser26	192.168.16.26 80	204.116.73.2 80
ser27	192.168.16.27 80	204.116.73.3 80
ser59	192.168.16.59 80	204.116.73.4 80

This looks a lot like our earlier table, except that only the specified services are permitted in. This example shows 4 WWW Servers on the inside that are to be offered to users on the outside. In practice, the table also includes the higher-level protocol (TCP or UDP) so that the Net NAT may discriminate to that level. For a far more interesting configuration, consider:

Name	Actual Address and Port	Apparent Address and Port
ser12	192.168.16.12 8000	204.116.73.1 80
ser12	192.168.16.12 8001	204.116.73.2 80
ser12	192.168.16.12 8002	204.116.73.3 80
ser59	192.168.16.59 23	204.116.73.1 23

This shows three publically-visible WWW Servers on apparent addresses 1-3, and an inbound telnet destination sharing the first apparent address. While the outside "sees" three separate WWW Server "machines," there is really only one "machine" on the inside. There are, however, three separate WWW Server processes running on that machine, using ports 8000 through 8002. There is probably also one on port 80 for the internal users, but that doesn't need to be in the table. It's interesting to note that WWW requests to apparent address 1 go to machine ser12, while telnet requests go to machine ser59. This allows very tight securing of the network by directing telnets to a secured telnet proxy and keeping telnet accounts off of the WWW Server.

Mux Server Mapping Mode

This is much like Port Mode, in that an external combination of IP Address and port is mapped into the inside world. Unlike Port Mode, and unlike any established product, the Mux Mode maps incoming requests to up to four internal servers, to distribute the service load across multiple platforms. We called this our "WebMUX" product in 1994, but have added it to our standard NAT line at no extra cost.

Support for Great New Applications

The Public Internet is an exciting environment for us all. What is fueling the stunning growth is the array of wonderful new client/server applications, and new browsers for older applications. Running some of these new applications through a NAT device can be a real challenge. We mentioned the kind of problem that one encounters in our discussion of embedded IP Address and Port information. We will teach the Net NAT to deal with these applications if we can. A notable example (and a wonderful product) is Realaudio from Progressive Networks. The potential of this product encouraged us to add specific modifications to the Net NAT.

This page was last modified on April 18, 1996.

Copyright © 1996 Network Safety

This information is proprietary to Network Safety. Network Safety, WebElite and NetNAT are trademarks of Network Safety. For information on our products and services, please contact our sales department. This page was prepared using WebElite, our professional editor for the Web.

**DIGITAL
Technical Journal**

Updated: 26 January 1998

Development of Router Clusters to Provide Fast Failover in IP Networks

Peter L. Higginson and Michael C. Shand

Abstract

IP networks do not normally provide fast failover mechanisms when IP routers fail or when links between hosts and routers break. In response to a customer request, a DIGITAL engineering team developed new protocols and mechanisms, as well as improvements to the DECNIS implementation, to provide a fast failover feature. The project achieved loss-of-service times below five seconds in response to any single failure while still allowing traffic to be shared between routers when there are no failures.

Introduction

A DIGITAL router engineering team has refined and extended routing protocols to guarantee a five-second maximum loss-of-service time during a single failure in an Internet Protocol (IP) network. We use the term router cluster to describe our improved implementation. A router cluster is defined as a group of routers on the same local area network (LAN), providing mutual backup. Router clusters have been in service since mid-1995.

Background

The Digital Equipment Corporation Network Integration Server (DECNIS) bridge/router is a midrange to high-end product designed and built by a DIGITAL Networks Product Business Group in Reading, U.K. [1] The DECNIS performs high-speed routing of IP, DECnet, and OSI (open system interconnection) protocols and can have the following network interfaces: Ethernet, FDDI (fiber distributed data interface), ATM (asynchronous transfer mode), HSSI (High-Speed Serial Interface), T1/E1 (digital transmission schemes), and lower-speed WAN (wide area network) interfaces. The DECNIS bridge/router is designed around a Futurebus backplane, with a number of semi-autonomous line cards, a hardware based address lookup engine, and a central control processor responsible for the control protocols and route calculation. Data packets are normally handled completely by the line cards and go to the central processor only in exception cases.

The DECNIS routers run a number of high-profile, high-availability, wide-area data networks for telephone service providers, stock exchanges, and chemical companies, as well as forming the backbone of DIGITAL's internal network.

Typically, the DECNIS routers are deployed in redundant groups with diverse interconnections, to provide very high availability. A common requirement is never to take the network down (i.e., during maintenance periods, connectivity is preserved but redundancy is reduced).

Overview

IP is the most widely used protocol for communication between hosts. Routers (or gateways) are used to link hosts that are not directly connected. When IP was originally designed, duplication of WAN links was common but duplication of gateways for hosts was rare, and no mechanisms for avoiding failed routers or broken links between hosts and routers were developed.

In 1994, we began a project to restrict loss-of-service times to below five seconds in response to any single failure; for example, failure of a router or its electrical supply, failure of a link between routers, or failure of the connection between the router and the LAN on which the host resides. In contrast, existing routing protocols have recovery times in the 30- to 45-second range, and bridging protocols are no better. Providing fast failover in IP networks required enhancements to many areas of the router's design to cover all the possible failure cases. It also required the invention of new protocols to support the host-router interaction under IP. This was achieved without requiring any changes to the host IP code.

In this paper, we start by discussing our targets and the behavior of existing routing or bridging protocols and follow this with a detailed analysis of the different failure cases. We then show how we have modified the behavior of the routing control protocols to achieve the desired failover times on links between routers or in response to the failure of intermediate routers. Finally, we describe the new IP Standby Protocol and the mechanisms we developed to achieve fast recovery from failures on the LANs local to the end hosts. This part of the problem is the most challenging because the hosts are of many types and have IP implementations that cannot realistically be changed. Thus all changes have to be made in the routers.

<http://www.asia-pacific.digital.com/DTJR02/DTJR02HM.HTM>

9/29/98 4:31:11 PM

Our secondary aims were to allow the use of router clusters in any existing network configuration, not to constrain failover to simple pairs of routers, to be able to share traffic between available routers, and to continue to use the Internet Control Message Protocol (ICMP) redirect mechanism for optimum choice of router by hosts on a per destination basis. A common problem of hosts is that they do not time out redirects. This problem is avoided by the adoption mechanism within the router cluster. Having met these aims, as well as fast failover, we can justifiably call the result router clusters.

The Customer Challenge

A particular customer, a telecommunications service provider, has an Intelligent Services Network application by which voice calls can be transferred to another operator at a different location. The data network manages the transferral and passes information about the call. The application uses User Datagram Protocol (UDP) packets in IP with retransmission from the application itself.

Because this application requires a high level of data network availability, network designers planned a duplicate network with many paired links and some mesh connections. Particular problems arise when the human initiator becomes impatient if there are delays; however, the more critical requirement was one over which the network designers had no control. The source of the calls is another system that makes a single high-level retransmission after five seconds. If that retransmission does not receive a response, the whole system at the site is assumed to have failed. This leads to new calls being routed to other service sites or suppliers, and manual intervention is required.

To resolve this issue, the customer requested a networking system that would recover from a single failure in any link, interface, or router within a five-second period. The standard test (which both the customer and we use) is to start a once-per-second ping, and to expect to drop no more than four consecutive ping packets (or their responses) upon any event. The five-second maximum break also has to apply to any disruption when the failed component recovers.

To meet the customer challenge, the router group in Reading developed the router cluster implementation on the DECNIS. In the next two sections, we discuss the bridging and routing protocols in use at the start of our project and relate our analysis of the customer's network problems.

Bridging and Routing Default Recovery Times

In a large network, a routing control protocol is essential in order to dynamically determine the topology of the network and to detect failing links. Bridging control protocols may be used similarly in smaller networks or may be used in combination with routing.

Bridging and routing control protocols often have failure recovery times in the order of a minute or more. A typical recovery consists of a detect time during which adjacent routers learn about the failure; a distribution time during which the knowledge is shared, possibly throughout the whole network; and a route recalculation time during which a new set of routes is calculated and passed to the forwarding engine.

Detection times are in the order of tens of seconds; for example, 30 seconds is a common default. The two most popular link-state routing control protocols in large IP networks are Open Shortest Path First (OSPF)² and Integrated Intermediate System-to-Intermediate System (Integrated IS-IS).³ These protocols have distribution "hold downs" (to limit the impact of route flaps) to prevent the generation of a new control message within some interval (typically 5 or 30 seconds) of a previous one. The distribution of the new information is rapid (typically less than one second), depending primarily on link speeds and network diameter; however, the distribution may be adversely affected by transmission errors which require retransmission. The default retransmission times after packet loss vary between 2 and 10 seconds. The route recalculation typically takes less than one second. These values result in total recovery times after failures (for routing protocols with default settings) in the 45- to 90-second range.

Distance vector routing protocols, such as the Routing Information Protocol (RIP),⁴ typically take even longer to recover, partly because the route computation process is inherently distributed and requires multiple protocol exchanges to reach convergence, and partly because their timer settings tend to be fixed at relatively long settings. Consequently, their use is not further considered in this paper.

Similarly, bridging protocols, as standard, use a 15-second timer; one of the worst-case recovery situations requires three time-outs, making 45 seconds in all. Another bridging recovery case requires an unsolicited data packet from a host and this results in an indeterminate time, although a timeout will cause flooding after a period.


In IP protocols, there is no simple way for a host to detect the failure of its gateway; nor is it simple for a router to detect the failure to communicate with a host. In the former case, several minutes may pass before an Address Resolution Protocol (ARP) entry times out and an alternative gateway is chosen; for some implementations, recovery may be impossible without manual intervention. Failure to communicate with a host may be the result of failure of the host itself,

which is outside the scope of this project. Alternatively, it may be due to failure of the LAN, or the router's LAN interface. In this case, there exists an alternative route to the LAN through another router, but the routing protocols will not make use of it unless the subnet(s) on the LAN are declared unreachable. This requires either manual intervention or timely detection of the LAN failure by the router.

Analysis of the Failure Cases

The first task in meeting the customer's challenge was to analyze the various failure and recovery modes and determine which existing management parameters could be tuned to improve recovery times. After that, new protocols and mechanisms could be designed to fill the remaining shortcomings.

A typical network configuration is shown in Figure 1. The target network is similar but has more sites and many more hosts on each LAN. Many of the site routers are DECNIS 500 routers with one or two WAN links and two Ethernets. The second Ethernet is used as a management rail and as a redundant local path between routers one and two (R1-R2) and between routers three and four (R3-R4).



KEY:

----- POSSIBLE MANAGEMENT LAN
ALSO PROVIDING REDUNDANT PATH

Figure 1
Typical Configuration for Router Cluster Use

In the original plans for the customer network, the router cloud consisted of groups of routers at two or three central sites and pairs of links to the host sites. In designing our solution, however, we tried to allow any number of routers on each LAN, interconnected by a general mesh network. For test purposes, both we and the customer used this set-up with direct R1-R3 and R2-R4 T1 links as the network cloud.

We have to consider what happens to packets traveling in each direction during a failure: there is little gain in delivering the data and losing the acknowledgments. Since the direction of data flow does not give rise to additional complications in the network cloud, there are just two failure cases:


1. Failure of a router in the network cloud
2. Failure of a link in the network cloud

We keep these cases distinct because the failure and recovery mechanisms are slightly different.

We also need to consider a failure local to one of the LANs on which the hosts are attached. A failure here has two consequences: (1) The packets originated by the host must be sent to a different router, and (2) The response packets from the other host through the network cloud must also be sent to a different router, so that it can send them to the host. We break down this type of failure into the following three cases:

3. Packets from the host to a failed or disconnected router
4. Packets to the host when the router fails
5. Packets to the host when the router interface fails

Note that we are using the term *router interface failure* to include cases in which the connector falls out or some failure occurs in the LAN local to the router (such that the router can detect it). In practice, failure of an interface is rare. (Removing the plug is not particularly common in real networks but is easy to test.) Figure 2 shows these failure cases; this configuration was also used for some of the testing.



FAILURE CASES

1. FAILURE OF A ROUTER IN THE NETWORK CLOUD
2. FAILURE OF A LINK IN THE NETWORK CLOUD
3. PACKETS FROM THE HOST TO A FAILED OR DISCONNECTED ROUTER
4. PACKETS TO THE HOST WHEN THE ROUTER FAILS
5. PACKETS TO THE HOST WHEN THE ROUTER INTERFACE FAILS

Figure 2
Diagram of Failure Cases Targeted for Recovery

Recovery of a link that previously failed causes no problems because the routers will not attempt to use it until after it has been detected as being available. Prior to that, they have alternate paths available. Recovery of a failed router can cause problems because the router may receive traffic before it has acquired sufficient network topology to forward the traffic correctly. Recovery of a router is discussed more fully in the section on Interface Delay.

Can Existing Bridging or Routing Protocols Achieve 5-Second Failover in a Network Cloud?

In this section, we discuss the failure of a router and the failure of a link in the network cloud (cases 1 and 2). The customer requested enhanced routing, and the existing network was a large routed WAN, so enhancing bridging was never seriously considered. Our experience has shown that the 15-second bridge timers can be reduced only in small, tightly controlled networks and not in large WANs. Consequently, bridging is unsuitable for fast failover in large networks.

For link-state routing control protocols such as OSPF and Integrated IS-IS, once a failure has been detected recovery takes place in two overlapping phases: a flood phase in which information about the failure is distributed to all routers, and a route calculation phase in which each router works out the new routes. The protocols have been designed so that only local failures have to be detected and manageable parameters control the speed of detection.

Detection of failure is achieved by exchanging Hello messages on a regular basis with neighboring routers. Since the connections are usually LAN or Point-to-Point Protocol (PPP) (i.e., with no link-layer acknowledgments), a number of messages must be missed before the adjacency to the neighbor is lost. The messages used to maintain the adjacency are independent of other traffic (and in a design like the DECNIS may be the only traffic that the control processor sees). Typical default values are messages at three-second intervals and 10 lost for a failure, but it is possible to reduce these.

Decreasing the Routing Timers

The default timer values are chosen to reduce overheads, to cover short outages, and to ensure that it is not possible for long packets to cause the adjacency to expire accidentally by blocking Hello transmission. (Note transmission of a 4,500-byte packet on a 64 kilobit-per-second link takes half a second, and queuing would normally require more than a packet time.) However, with high-quality T1 or higher link speeds in the target network and priority queuing of Hellos in the DECNIS, it is acceptable to send the Hellos at one-second intervals and count three missed as a failure. (Although we have successfully tested counts of two, we do not recommend that value for customers on WAN links because a single

link error combined with a delay due to a long data packet would cause a spurious failure to be detected.) The settings of one second and three repeats were within the existing permitted ranges for the routing protocols.

When these shorter timers are used, it is important that any LANs in the network should not be overloaded to the extent that transmissions are delayed. The network managers should monitor WAN links and disable any links that have high error rates. Given the duplication of routes, it is better to disable and initiate repairs to a bad link than to continue a poor service. Many customers, with less controlled networks and less aggressive recovery targets, have adopted the router cluster system but kept to more conservative timers (such as 1 second and 10 repeats).

Implementation and Testing Issues

In some cases, a failed link may be detected at a lower level (e.g., modem signals or FDDI station management) well before the routing protocol realizes that it has stopped getting Hellos and declares the adjacency lost. (This can lead to good results during testing, but it is essential also to test link-failure modes that are not detected by lower levels.) In the worst case, however, both the detection of a failed router or the detection of a failed link rely on the adjacency loss and so have the same timings.

Loss of an adjacency causes a router to issue a revised (set of) link-state messages reflecting its new view of the local topology. These link-state messages are flooded throughout the network and cause every router in the network to recalculate its route tables. However, because the two or more routers will normally time out the adjacency at different times, one message arrives first and causes a premature recalculation of the tables. Therefore it may require a subsequent recalculation of the route tables before a new two-way path can be utilized. We had to tune the router implementation to make sure that subsequent recalculations were done in a speedy manner.

During initial testing of these parameters, we discovered that failure of certain routers represented a more serious case. However discussion of this is deferred to the later section The Designated Router Problem.

Our target five seconds is made up of three seconds for the failure to be detected, leaving two seconds for the information about the failure to be flooded to all routers and for them to recalculate their routes. Within the segment of the network where the recovery is required, this has been achieved (with some tuning of the software).

Recovery from Failures on the LANs Local to the End Hosts

The previous section shows that we can deal with router failure and link failure in the network cloud (cases 1 and 2). Here we consider cases 3, 4, and 5, those that deal with failures on the LANs local to the end hosts.

From the point of view of other routers, a failed router on a LAN (case 4) is identical to a failed router in the network cloud (case 1): a router has died, and the other routers need to route around it. Failure case 4 therefore is remedied by the timer adjustments described in the previous section. Note that these timer adjustments are an integral part of the LAN solution, because they allow the returning traffic to be re-routed. These timer adjustments cannot work properly if the LAN parts of router clusters are using an inappropriate routing control protocol such as RIP⁴, which takes up to 90 seconds to recover from failures.

Detecting LAN Failure at the Router

A solution to case 5—packets to the host when the router interface fails—for IP requires that the router can detect a failure of its interface (for example, that the plug has been removed). If the LAN is an FDDI, this is trivial and virtually instantaneous because continuous signals on the ring indicate that it is working and the interface directly signals failure. For Ethernet, we faced a number of problems, partly due to our implementation and partly due to the nature of Ethernet itself. We formed a small team to work on this problem alone.

Because of the variety of Ethernet interfaces that might be attached, there is no direct indication of failure: only an indirect one by failure to successfully transmit a packet within a one-second interval. For maximum speed, the DECNIS implementation queues a ring of eight buffers on the transmit interface and does not check for errors until a ring slot is about to be reused. This means that an error is only detected some time after it has occurred, consuming much of our five-second budget.

The control software in the DECNIS management processor has no direct knowledge of data traffic because it passes directly between the line cards. Therefore it sends test packets at regular intervals to find out if the interface has failed. By sending large test packets occupying many buffers, it ensures that the ring circulates and errors are detected. Initially, we reduced the timers and increased the frequency of test packets to be able to detect interface failure within three seconds. (The test packets have the sender as destination so that no one receives them and, as usual, more than one failure to transmit is required before the interface is declared unusable.)

This initial solution caused several problems when it was deployed to a wider customer group; we had more complaints than previously about the bandwidth consumed by the test messages and, more seriously, a number of instances of

previously working networks being reported as unusable. These problem networks were either exceptionally busy or had some otherwise undetected hardware problem. Over time, the networks with hardware problems were fixed, and we modified the timers to avoid false triggering on very busy networks. Clearly, the three-second target required more thought.

Several enhancements have since been made. First, the timers are user configurable so that the network managers can trade off between aggressive recovery times, bandwidth used, and false detection. Second, the test packet generator takes into account other packets sent by the control processor such that they are only sent to the size and extent required for the total traffic to cause the ring to circulate. This is a significant improvement because the aggressive routing timers discussed previously cause Hello packets to be sent at one-second intervals, which is often sufficient not to require extra test packets. Third, the line card provides extra feedback to the control program about packets received and the transmission of packets not originated by the control processor. This feedback gives an indication of successful operation even if some transmits are failing.

Re-routing Host Traffic When a Router or Router Connection Fails

Case 3 was by far the most difficult problem to solve. IP does not provide a standard mechanism to re-route host traffic when a router fails, and the only method in common use (snooping RIP messages in the hosts) is both "deprecated" by the RFCs and has fixed 45-second timers that exceed our recovery target. Customers have a wide range of IP implementations on their hosts, and reliance on nonstandard features is difficult. The particular target application for this work ran on personal computer systems with a third-party IP stack, and we obtained a copy for testing. Such IP stacks sometimes do not have sophisticated recovery schemes and discussion with various experts led us to believe that we should not rely on any co-operation from the hosts.

Among other objectives, we wanted to be independent of the routing control protocol in use (if any), to permit both a mesh style of networking and more than two routers in a cluster, and to continue to route traffic by reasonably optimal routes. In addition, we wished to not confuse network management protocols about the true identity of the routers involved and, if possible, to share traffic over the WAN links where appropriate.

Electing a Primary Router

In our solution, the first requirement is for other routers on the LAN to detect that a router has failed or become disconnected, and to have a primary router elected to organize recovery. This is achieved by all routers broadcasting packets (called IP Standby Hellos) to other routers on the LAN every second. The highest priority (with the highest IP address breaking ties) router becomes the primary router, and failure to receive IP Standby Hellos from another router for *n* seconds (three is the default) causes it to be regarded as disconnected. This condition may cause the selection of a new primary router, which would then initiate recovery to take traffic on behalf of the disconnected router.

The IP Standby Hellos are sent as "all routers multicasts" and therefore do not add additional load to hosts. They are UDP datagrams⁵ to a port we registered for this purpose (digital-vrc; see the Internet Assigned Numbers Authority [IANA] on-line list). The routers are manually configured with a list of all routers in the cluster. To make configuration easier and less error prone, the list on each router includes itself, and hence an identical set of configuration parameters can be used for all the routers in a cluster. Automatic configuration was rejected because of the problem of knowing which other routers should exist.

Function of the Primary Router in ARP Mode

Our first attempt (called ARP Mode) uses a fake IP address (one per subnet for a LAN with multiple subnets), which the current primary router adopts and the hosts have configured as their default router. The primary router returns its own media access control (MAC) address when the host broadcasts an ARP request (using the standard ARP protocol⁶) for the fake IP address and thus takes the traffic from the host. After a failure, a newly elected primary router broadcasts an ARP request containing the information that the fake IP address is now associated with the new primary router's MAC address. This causes the host to update its tables and to forward all traffic to the new primary router.

The sending of ICMP redirects⁷ by the routers has to be disabled in ARP mode. Redirects sent by a router would cause hosts to send traffic to an IP address other than the fake IP address controlled by the cluster, and recovery from failure of that router would then be impossible. Disabling redirects causes an additional problem. If the primary router's WAN link fails, all the packets have to be inefficiently forwarded back over the LAN to other routers. To avoid this problem, we introduced the concept of monitored circuits, whereby the priority of a router to become the primary depends on the state of the WAN link. Thus, the primary router changes when the WAN link fails (or all the links fail if there are several), and the hosts send the packets to the new primary (whose WAN link is still intact).

ARP mode has a number of disadvantages. It does not necessarily use an optimum route when the WAN links form a mesh rather than the simple pair case, because redirects have to be disabled. The monitored circuit concept works only on the first hop from the router; more distant failures cannot change the IP Standby priority and may result in inefficient routing. Most seriously, the rules for hosts acting on information in ARP requests have only a "suggested implementation" status in the RFCs, and we found several hosts that did not change when requested or were very slow in

doing so. (Note that we did consider broadcasting an ARP response, but there is no allowance in the specifications for this message to be a broadcast packet, whereas an ARP request is normally a broadcast packet.)

MAC Mode IP Standby (to Re-route Host Traffic)

To solve these problems, we looked for a mechanism that did not rely on any host participation. The result was what we termed MAC mode. Here, each router uses its own IP address (or addresses for multiple subnets) but answers ARP requests with one of a group of special MAC addresses, configured for each router as part of the router cluster configuration. When a router fails or becomes disconnected, the primary (or the newly elected primary) router adopts the failed router. By adopt, we mean it responds to ARP requests for the failed router's IP address with the failed router's special MAC address, and it receives and forwards all packets sent to the failed router's special MAC address (in addition to traffic sent to the primary router's own special MAC address and those of any other failed routers it has adopted).

The immediate advantages of MAC mode are that ICMP redirects can continue to be used, and, providing the redirects are to routers in the cluster, the fast failover will continue to protect against further failures. The mechanism is completely transparent to the host. In a cluster with more than two routers, the primary router will use redirects to cause traffic (resulting from failure) to use other routers in the cluster if they have better routes to specific destinations. Thus multiple routers in a cluster and mesh networks are supported. This also solves the problem of hosts not timing out redirects (an omission common to many IP implementations derived from BSD), because the redirected address has been adopted.

In MAC mode, the hosts are configured with the IP address of any router in the cluster as the default gateway. (The concept that it does not matter which router is chosen is one of the hardest for users to accept.) Some load sharing can be achieved by setting different addresses in different hosts.

Since the DECNIS is a bridge router, it has the capability to receive all packets on Ethernet and many MAC addresses on FDDI; thus all packets on all the special MAC addresses are seen by all routers in the cluster, and its own and those of any adopted routers are forwarded. The special MAC addresses used are those associated with the unused DECnet area 0. They are ideal because they are part of the locally administered group and have implementation efficiencies in the DECNIS because the DECnet hi-ord (AA-00-04-00) is already decoded, and they are 16 addresses differing in one nibble only (i.e., AA-00-04-00-0x-00, where x is the hexadecimal index of the router in the cluster). Note that ARP requests sent by the router must also contain the special MAC address in the source hardware address field of the ARP packet, otherwise the hosts' ARP tables may be updated to contain the wrong MAC address.

MAC mode has minor disadvantages. Initially, it is easy to spread the load over a number of routers; however, this can be lost after redirects. In addition, a small chance of packet duplication exists during recovery because there may be a short period when both routers are receiving on the same special MAC address (which does not happen in ARP mode because the host changes the MAC address it is using). This is preferable to a period when no router is receiving on that address.

Interface Delay

Recently, we added an interface delay option to ameliorate a situation more likely to occur in large networks. In this situation, a router, rebooting after a power loss, a reboot, or a crash, reacquires its special MAC address before it has received all of the routing updates from neighboring routers and thus drops packets sent to it (and worse, returns "unreachable" to the host). Typically, the main LAN initialization would be delayed for 30 seconds while routing table updates were received over the WAN interfaces and any other LAN interfaces. The backup continues to operate during this 30 seconds. (Note that with Integrated IS-IS, we could have delayed IP on the whole router, but we did not do this because it would not have worked for OSPF, which requires IP to do the updates.) We use a fixed configurable time rather than attempting to detect the end of updating, because determining completion is difficult if the network is in a state of flux or the router's WAN links are down.

Redirects and Hosts That Ignore Them

When a router issues an ICMP redirect, the RFCs state that it must include its own IP address in the redirect packet. A host is required to ignore a redirect received from a router whose IP address is not the host's next hop address for the particular destination address. Therefore, it is necessary to ensure that the address of the failed router is correctly included when issuing a redirect on its behalf. In the DECNIS implementation, because the destination MAC address of a received packet is not available to the control processor, the primary router cannot tell whether a redirect has to be issued on behalf of itself or one of the adopted routers. The primary router therefore issues multiple redirects—one for each adopted router (in addition to its own). Since redirects are rare, this is not a problem, but they could be avoided by passing the MAC destination address of the original packet (or just five bits to flag a special MAC address and say which it is) to the control processor.

It is contrary to the basic IP rules for hosts to ignore redirects.⁸ Despite the rules, some hosts do ignore redirects and continue sending traffic which has to be sent back over the same LAN. These cause problems in all networks because of the load, and, in the DECNIS implementation, because every time the line card recognizes a redirect opportunity, it signals the control processor to consider sending a redirect. This may happen at data packet rates and is a severe load on the control processor, which slows down processing of routing updates and might then cause our five-second recovery

target to be exceeded.

To reduce the problems caused by hosts ignoring redirects, we improved the implementation to rate-limit the generation of redirect opportunity messages by the line cards. We also recommend in the documentation that, where it is known that hosts ignore redirects (or their generation is not desired), the routers be connected by a lower-cost LAN than the main service LAN (such as the management LANs shown in Figure 1). Normally, this would mean linking (just) the routers by a second Ethernet and setting its routing metric so that it is preferred to the main LAN for packets that would otherwise traverse back on the main LAN to the other router. This has two advantages. Such packets do not consume double bandwidth and cause congestion on the main LAN, and they pass only through the fast-path parts of the router, which are well able to handle full Ethernet bandwidth.

In MAC mode, it is also possible to define a router that does not actually exist (but has an IP address and a special MAC address) and is adopted by another router, depending on the state of monitored WAN circuits. Setting this as the default gateway is another way of coping with hosts that ignore redirects.

Special Considerations for Bridges

We do not recommend putting a bridge or layer 2 switch between members of a router cluster, because during failover, action would be required from the bridge in order for the primary router to receive packets that previously were not present on its side of the bridge. We cannot rely on this being the case, so we must have a way of allowing bridges to learn where the special MAC addresses currently are. More importantly, if bridges do not know where the special MAC addresses are, they often use much less efficient (flooding) mechanisms.

For greater traceability (and simpler implementation), we use the router's real MAC address as the source address in data packets that it sources or forwards. We use the special MAC address as the source address in the IP Standby Hellos. Since the Hello is sent out as an IP multicast, it is seen by all bridges or switches in the local bridged network and causes them to learn the location of the address (whereas data packets might not be seen by non-local bridges). Since we are sending the Hellos every one second anyway, there is no extra overhead.

When a primary router has adopted routers, it cycles the source MAC address used for sending its Hello between its own special MAC address and those of the adopted routers. We also send out an additional Hello immediately when we adopt a router to speed up recognition of the change.

Since the same set of special MAC addresses is used by all router clusters, we were concerned that a bridge that was set up to bridge a non-IP protocol (e.g., local area transport [LAT]) but not to bridge IP, might be confused to see the same special MAC address on more than one port. (This has been observed to happen accidentally, and the resultant meltdown has led us to avoid any risk, however slight, of this happening.) Hence we make 16 special MAC addresses available and recommend to users that they allocate them uniquely within a bridged domain, or at least use disjoint sets on either side of a bridge.

The Designated Router Problem

While testing router failures, we discovered additional delays during recovery due to the way in which link-state protocols operated on LANs. In these cases, the failure of routers not handling the data packets can also result in interruption of service due to the control mechanisms used.

For efficiency reasons in link-state routing protocols, when several routers are connected to a LAN, they elect a designated router and the routing protocols treat the LAN as having a single point-to-point connection between each real router and a pseudo router maintained by the designated router (rather than connections between all the routers). The designated router issues link-state packets on behalf of the pseudo router, showing it as having connections to each real router on the local LAN, and each router issues a link-state packet showing connection to the pseudo router. This mechanism operates in a broadly similar way in both Integrated IS-IS and OSPF; the primary difference being that the OSPF election exhibits hysteresis, thus minimizing unnecessary designated router changes.

For routing table calculations, a transit path over the LAN is taken from a router to the pseudo router and then to another router on the LAN. Hence any change in pseudo router status disrupts calculation of the network map.

When a designated router fails, a slew of updates occurs; each router on the LAN loses the adjacency to the old designated router and issues a new link-state packet. Next, the new designated router is elected (or in the case of OSPF, the backup designated router takes over), and each router issues a link-state packet showing a link to it. In parallel, the new designated router issues a set of link-state packets showing its connections. This is a new router on the network as far as the other routers are concerned; the old designated router stays, disconnected, in the tables for as long as 20 minutes to an hour. This happens at level 1 and at level 2 in Integrated IS-IS, resulting in twice as many updates. The interactions are complex; in general, they result in the sending of multiple, new link-state messages.

Apart from the pure distribution and processing problem of these updates and new link-state packets, there are deliberate delays added. A minor one is that updates in Integrated IS-IS are rate-limited on LANs (to minimize the possibility of message loss). A major one is that a particular link-state packet cannot be updated within a holding time from a previous update (to limit the number of messages actually generated). The default holding time is 30 seconds in Integrated IS-IS; it can be reduced to 1 second in the event we found that the best solution was to allow as many as 10 updates in a 10-second period. The reason for this is that the first update usually contains information about the disconnection and it is highly desirable to get the update with the connection out as fast as possible. In addition, in the wider network, an update can overtake and replace a previous one.

With OSPF, the protocol defines a minimum holding time of five seconds, which limits the recovery time when the designated router fails. The target customer's network was using Integrated IS-IS, and so we were able to achieve the five-second recovery even when the designated router failed. (Note that with two routers, one must be the designated router so it is not a rare case.) We have not, so far, felt that it is worthwhile to break the rules by allowing a shorter holding time for OSPF.

Conclusions

We successfully designed and implemented router clusters for the DECNIS router with shared workload and interruptions after failures of less than five seconds in both LAN and WAN environments. This capability has been deployed in the product since the middle of 1995. An Internet Engineering Task Force (IETF) group is currently attempting to produce a standard protocol to meet this need.⁹

Acknowledgments

Various members of the router engineering team in Reading, U.K. assisted with ideas for this work. In particular, we must mention Dave Forster who implemented the high-level IP changes, Chris Szmidt who implemented the line card forwarding, and John Rigby who implemented the bit in-between and the Ethernet cable-out detection.

References

1. D. Brash and S. Bryant, "The DECNIS 500/600 Multiprotocol Bridge Router and Gateway," *Digital Technical Journal*, vol. 5, no. 1 (Winter 1993): 84-98.
2. J. Moy, "OSPF Version 2," Internet Engineering Task Force, RFC 1583 (March 1994).
3. R. Callon, "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments," Internet Engineering Task Force, RFC 1195 (December 1990).
4. C. Hedrick, "Routing Information Protocol," Internet Engineering Task Force, RFC 1058 (June 1988).
5. J. Postel, "User Datagram Protocol," SRI Network Information Center, Menlo Park, Calif., RFC 768 (August 1980).
6. D. Plummer, "Ethernet Address Resolution Protocol," Internet Engineering Task Force, RFC 826 (November 1982).
7. J. Postel, "Internet Control Message Protocol," Internet Engineering Task Force, RFC 792 (September 1981).
8. R. Braden, "Requirements for Internet Hosts—Communication Layers," Network Information Center, RFC 1122 (October 1989).
9. R. Hinden, S. Knight, D. Weaver, D. Whipple, D. Mitzel, P. Hunt, P. Higginson, and M. Shand, "Virtual Router Redundancy Protocol," Internet Drafts <draft-ietf-vrrp-spec-03.txt> (October 1997).

Biographies



Peter L. Higginson

Peter Higginson manages the advanced development work on router products for DIGITAL's Internetworking Products

Engineering Group in Reading U.K. (IPEG-Europe). His responsibilities include improving communications on customers' large networks. Most recently, he contributed to the corporate Web gateway strategy and future router products. Peter was issued one patent on efficient ATM cell synchronization and has applied for several other patents related to networks. He has published many papers, including one on a PDP-9 for DECUS in 1971. Before joining DIGITAL in 1990, Peter was the software director for UltraNet Ltd. (now part of the Anite Group), a maker of X.25 equipment. For 12 years before that, he was a lecturer in the Department of Computer Science, University College London. He received an M. Sc. in computer science from University of London in 1970 and a B. Sc. (honours) in mathematics from University College London in 1969. Peter connected the first non-U.S. host to the Arpanet in 1973.



Michael C. Shand

Mike Shand is a consulting software engineer with DIGITAL's Network Products Business in Reading, U.K. He is currently involved in the design of IP routing algorithms and the system-level design of networking products. Formerly, Mike was a member of the NAC (Networks and Communications) Architecture Group where he designed DECnet OSI, Phase V routing architecture. Before joining DIGITAL in 1985, Mike was the assistant director (systems) of the Computing Centre at Kingston University. He earned an M.A. in the natural sciences from the University of Cambridge in 1971 and a Ph. D. in surface chemistry from Kingston University in 1974. He was awarded six patents (and has filed another) in various aspects of networking.

Trademarks

The following are trademarks of Digital Equipment Corporation: DECnet, DECNIS, and DIGITAL.

[DIGITAL home](#) [Feedback](#) [Search](#) [Map](#) [Subscribe](#) [Help](#)
[Legal](#)

-December 24, 2001-

112



Where Can I Get a Domain Name?

[Home](#) | [Support](#) | [Contact](#) | [Policy](#) | [About](#) | [Search](#) | [Referral Program](#)

Internet Access
DSL Service
Personal Accounts
Commercial Accts.
Dedicated Access
E-Commerce
Other Services
Web Design
Domain Names
Computer Systems
Subscribers Area
Tech Support
Search the Net
WebMail
Elected Postings
Press & News Point
Sign Up Now!

Navpoint Internet provides a wide range of dedicated Internet access solutions, including Frame Relay, Point-to-Point, Centrex ISDN and Dedicated Dial-Up. A dedicated account, unlike a dial-up account, is meant to be connected all the time. DSL (Navpoint DSL) is also a dedicated, "always-on" connection - please see our DSL section for complete DSL information. Frame Relay and Point-to-Point circuits are quoted based on your location and the speeds desired. Please contact sales@navpoint.com, or, fill out the information request form at the bottom of this page.

NEW - VOI ISDN

Virtual Office ISDN is a new ISDN service being offered that significantly drops the price of an ISDN connection for a business. Like Centrex ISDN, the VOI ISDN can be used as a dedicated 24/7 connection OR as a dial up ISDN connection at either 64K or 128K speeds. This document (Virtual_Office_ISDN) will explain the basics of VOI and provide general pricing information.

Frame Relay

Frame Relay is one of the most popular dedicated access services available because of the excellent price to performance ratio. The telephone company provides the circuit between the end-user and a frame relay switch, which handles the frame relay traffic for that area. Because the switch is shared between several frame relay customers, the cost of the circuit is reduced.

When you purchase a frame relay circuit, you are contracting for a specific CIR - Committed Information Rate. The CIR is a setting on the circuit that will limit your bandwidth during periods of heavy usage on the frame relay network. This can be thought of as the amount of bandwidth that you are guaranteed to receive. Frame relay also has a burst speed that is available for use when traffic on the switch is light. The burst speed is typically twice the amount of the CIR. For example, if you purchase a 768K Frame Relay, your burst speed is full T1 - 1.54mb (768K x2). Since an end user will rarely use the full bandwidth available for anything other than short bursts of time, frame relay is a more economical choice than a point-to-point connection.

Frame Relay is available in speeds from 56K to T1. Pricing is based on your location and the term of circuit commitment (month to month, 3 year or 5 year). Two pieces of hardware are required at the end user location for Frame Relay - a router and a CSU/DSU. Navpoint Internet can provide this hardware, or, advise you on whether your existing hardware will be compatible with Frame Relay.

Frame Relay Pricing

The first table is your monthly charge from Bell Atlantic for the actual Frame Relay connection. You can choose a month to month, 3 year, or 5 year commitment to Bell for the circuit (5 year gets a \$1.00 installation charge). The second table is the Navpoint charge for the bandwidth (add both monthly charges together for your total circuit monthly cost).

Bell CIR	Monthly	3 Year	5 Year
128K (256K burst)	439.00	404.00	384.00
256K (512K burst)	444.00	409.00	389.00
384K (768K burst)	447.00	412.00	392.00

Navpoint DSL

- ◆ Residential and Commercial Service
- ◆ Speeds up to 1.5Mbps
- ◆ Convenient "always on" service
- ◆ Connect multiple machines to the Net with one DSL Line

Click for More Info

Sponsored by
COVAD



512K (1024K burst)	460.00	425.00	405.00
768K (1536K burst)	463.00	428.00	408.00
Installation	1012.00		1.00

Navpoint Frame Relay Bandwidth		
CIR Speed	Install	Monthly
128K	250.00	250.00
256K	250.00	300.00
384K	250.00	350.00
512K	250.00	420.00
768K	250.00	500.00

Navpoint Frame Relay accounts include commercial domain name hosting (15mb of web space), and DNS setup so you can run your own mail server. If you will not be running your own mail server, see the Additional Services section below for POP3 email pricing.

Point-to-Point

A point-to-point connection is a direct, private, T1 (1.54Mb/sec) or T3 (28 T1s, or, 45Mb/sec) connection between the end-user and Navpoint Internet. A T1 provides enough bandwidth to service a company, or, educational institution where access is required for several hundred people at the same time. A T3 can provide access for several thousand people, or, an entire campus area network. A Point-to-Point connection offers the utmost in stability, reliability and performance for an Internet connection.

Pricing for Point-to-Point connectivity is based on exact mileage between the end-user location and Navpoint Internet. Hardware for Point-to-Point connectivity is basically the same as for Frame Relay - a router and a CSU/DSU. Navpoint Internet will provide all necessary hardware for your Point-to-Point connection. Please contact sales@navpoint.com for more information and pricing.

Centrex ISDN

Centrex ISDN is an unmetered ("always-on", no per-minute charge), 64K or 128K digital ISDN connection. Centrex ISDN may be a good choice where DSL is not available due to location or distance issues. Centrex ISDN pricing is based on exact mileage between your location and ours, so, must be quoted on an individual basis. Please contact sales@navpoint.com for more information and pricing.

Dedicated Dial-Up

Navpoint Personal and Commercial dial-up accounts are intended for unlimited interactive usage - that is, designed to be actively used. They are not intended to be left on unattended for extended periods of time. If you need a dial up connection where you don't have to worry about remaining connected all the time, then, dedicated dial up is for you. We will remove all time-outs from your dial up account, and, you will be able to stay connected at all times.

# of Modems	Monthly	or Yearly	Setup
1 modem or 64K ISDN	100.00	1000.00	50.00
2 modems or 128K ISDN	175.00	1750.00	50.00

Additional Services Available

Service	Monthly	or Yearly	Setup
Single POP3 Email	3.00	30.00	n/c*
10 POP3 Emails	25.00	250.00	n/c*
20 POP3 Emails	40.00	400.00	n/c*

* Setup fees waived when combined with any Navpoint Dedicated Access

account.


For more information about Navpoint Dedicated Access accounts, please fill out this form and one of our sales or technical staff will contact you.

Name	<input type="text"/>
Company	<input type="text"/>
Phone	<input type="text"/>
E-Mail Address	<input type="text"/>
Service	<input type="text" value="Frame Relay"/>

Questions?

© 2000, Navpoint Internet

118



The Basic Guide to **Frame Relay Networking**

Your Complete Guide

to Frame Relay

from the

Frame Relay Forum

**Frame
Relay
Forum** 
www.frforum.com

The Basic Guide to **Frame Relay** Networking

**Your Complete Guide
to Frame Relay
from the
Frame Relay Forum**



Notice

Contributors

Todd Bahner, ADC Kentrox
Skip Carlson, Cabletron Systems
Anne Exter, Bell Atlantic
Mark Kaplan, Newbridge Networks
Chris Nicoll, Current Analysis, Inc.
Cheryl Vandegriff Hyon, Sync Research
Additional contributing companies include
MCI, Visual Networks, Netrix

Editor

Jan Thibodeau
JT Communications LLC

Design

Alan Greco Design

Copyright ©1998 Frame Relay Forum

Frame Relay Forum
39355 California Street Suite 307
Fremont, California 94538

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or information storage and retrieval systems now known or to be invented, without permission in writing from the Frame Relay Forum.

TABLE OF CONTENTS

Introduction	iii
Chapter 1	
Basic Gear	7
Chapter 2	
How Frame Relay Works	17
Chapter 3	
Frame Relay Signaling Mechanisms	26
Chapter 4	
Frame Relay Standards and Interoperability	40
Chapter 5	
Where Frame Relay is Used	46
Chapter 6	
Planning Your Frame Relay Network	70
Frame Relay Glossary	75
Appendix	85

INTRODUCTION

How to Use This Guide

Understanding a new technology is a lot like taking a hike on an unfamiliar trail. It helps to have a guide. Well, since we can't hike along with you, we thought we'd do the next best thing: provide a trail guide that will help you chart your course.

We're assuming, by the way, that you're not a completely uninitiated hiker. That is, you've got some basic familiarity with data communications and basic internetworking concepts, as well as data communications devices and their functions.

Like any good trail guide, we've tried to present the information in an easy-to-read format. Stop at the **Base Camp** in each chapter. That's where you'll find out what information is in that chapter and how it's organized. Typically, the **Basic Trail** will offer the fundamental information about the chapter topic. On the **Advanced Trail**, you'll get more challenging information. Here, we'll go into a little more depth on the topic and use a few more technical terms.

The **Base camp** will also list the **View Points** in each chapter. You guessed it...the things to look at. Tables, diagrams, figures.

Finally, if you've only got 20 minutes or so, take the **Shortcut**.

We believe that most network planners will benefit by reading this guide from cover to cover. But we also know that if you're involved in networking today, your hectic schedule may force you to skim through this booklet for highlights until you can find a quiet moment to read more thoroughly. The **Shortcuts** contain summaries of the major points you should know about frame relay networking. A quick scan can tell you which sections are most relevant to your planning needs.

Here's how the content breaks out:

- In the first chapter, you'll get your basic gear: a background on why frame relay was developed and what the benefits are. The



Base Camp



Basic Trail



Advanced Trail

Advanced Trail will discuss the ways in which circuit switching and X.25 switching are less than ideal and how they measure up to the demands of certain new high speed network applications.

- Chapter 2 defines the simple way in which frame relay data is switched, based on the address at the beginning of the frame.
- Chapter 3 describes the various mechanisms that are used by the frame relay network to communicate with the user device in order to avoid congestion on the network, to recover from an overload situation, and to convey the status of various connections.
- Chapter 4 discusses frame relay standards and interoperability and the work of the Frame Relay Forum.
- Chapter 5 examines several common applications used over frame relay networks.
- Chapter 6 discusses the steps you need to take and the questions you'll want to consider if you're planning a frame relay network.
- In the back of the book, you'll find a comprehensive frame relay glossary.



View Points



Shortcut

So fill up your canteen, lace up your hiking boots, and let's set out on the trail.

CHAPTER 1

BASIC GEAR



Base Camp

In this chapter, we will introduce you to frame relay and how it works. In other words, we'll give you the basic gear you'll need to continue your exploration.

Basic Trail: The Basic Trail will give you a definition of frame relay as a technology. Then, we'll explore the networking trends that combined to create a market need for frame relay. Finally, we'll talk about the benefits of using frame relay in your network.

Advanced Trail: If you take the advanced trail, you'll find a comparison of the characteristics of frame relay and other network switching technologies, namely Time Division Multiplexing (TDM), circuit switching and X.25 packet switching. (For a more detailed comparison of frame relay processing and X.25 processing, see Chapter 2.)



View Points:

- Figure 1: Frame relay network
- Figure 2: Opens Systems Interconnection (OSI) Model
- Table 1: Comparison chart of TDM circuit switching, X.25, and frame relay

Shortcut: If you have only a few minutes, take the shortcut. It will give you a quick overview of how frame relay was developed and the benefits of frame relay.



Basic Trail

What is Frame Relay?

Frame relay is a high-speed communications technology that is used in hundreds of networks throughout the world to connect LAN, SNA, Internet and even voice applications.

Simply put, frame relay is a way of sending information over a wide area network (WAN) that divides the information into frames or packets. Each frame has an address that the network uses to determine the destination of the frame. The frames travel through a series of switches within the frame relay network and arrive at their destination.

Frame relay employs a simple form of packet switching that is well-suited to powerful PCs, workstations and servers that operate with intelligent protocols, such as SNA and TCP/IP. As a result, frame relay offers high throughput and reliability that is perfect for a variety of today's business applications.

A Quick Look at a Frame Relay Network

A frame relay network consists of endpoints (e.g., PCs, servers, host computers), frame relay access equipment (e.g., bridges, routers, hosts, frame relay access devices) and network devices (e.g., switches, network routers, T1/E1 multiplexers).

Accessing the network using a standard frame relay interface, the frame relay access equipment is responsible for delivering frames to the network in the prescribed format. The job of the network device is to switch or route the frame through the network to the proper destination user device.

(See Figure 1.)

A frame relay network will often be depicted as a network cloud, because the frame relay network is not a single physical connection between one endpoint and the other. Instead, a logical path is defined within the network. This logical path is called a virtual circuit. Bandwidth is allocated to the path until actual data needs to be transmitted. Then, the bandwidth within the network is allocated on a packet-by-packet basis. This logical path is called a virtual circuit.

We'll be talking a lot more about virtual circuits – both Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs) in the next chapter. And we'll also discuss how frames or packets are "relayed" across the network.

But, before we get too technical, let's turn our attention to how and why frame relay got its start.

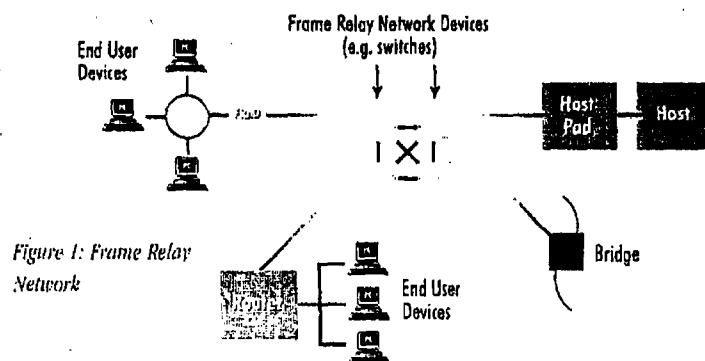


Figure 1: Frame Relay Network

Why was Frame Relay Developed?

From the beginning, frame relay was embraced enthusiastically by users because it was developed in response to a clear market need, namely the need for high speed, high performance transmission. Frame relay technology also made cost-effective use of widespread digital facilities and inexpensive processing power found in end user devices. Developed by and for data communications users, frame relay was simply the right technology at the right time. Let's explore the network trends that contributed to the development of frame relay.

As the 1980's came to a close, several trends combined to create a demand for and enable higher speed transmission across the wide area network:

- The change from primarily text to graphics interaction
- The increase in "bursty" traffic applications
- Intelligent end-user devices (PCs, workstations, X-Windows terminals) with increased computing power
- The proliferation of LANs and client/server computing
- Widespread digital networks

Need for Increased Speed

Today, rapid storage and retrieval of images for interactive applications is as common as transmitting full screens of text was in the 1970s and 1980s. Early graphics applications users who were accustomed to rapid information transfer over their LANs expected similar response times when transmitting data over the wide area. Since peak bandwidth requirements for graphics were substantially higher than for text transactions, increased bandwidth and throughput were clearly needed if response time expectations were to be met.

Dynamic Bandwidth Requirements

This type of LAN user required high bandwidth in bursts, followed by periods of idle time. "Bursty" traffic, as we call it, is well-suited for statistical sharing of bandwidth, which is a characteristic of frame relay technology.

Smarter Attached Devices

As networking requirements were changing, computing power was changing as well. Decreasing cost of processing power resulted in the proliferation of intelligent PCs and powerful workstations and servers all connected by LANs. These new end-user devices also offered the possibility of performing protocol processing, such as error detection and correction. This meant that the wide area network could be relieved of the burden of application layer protocol processing – another perfect fit for frame relay.

End-user equipment was becoming more sophisticated in its ability to recognize errors and retransmit packets at the same time as digital facilities were reducing error rates within the network. In addition, industry-standard high layer protocols, such as TCP/IP, added intelligence to end-user devices.

Without the overhead associated with error detection and correction, frame relay could offer higher throughput than other connectivity solutions, such as X.25.

Higher Performance

More LANs in general and Internet Protocol (IP) LANs in specific, fueled the need to internetwork LANs across the wide area network, another factor that drove the growth of public frame relay services.

Some users tried to solve the internetworking challenge by

simply hooking LAN bridges or routers together with dedicated lines. This approach worked for simple networks, but as complexity increased, the drawbacks became apparent: higher transmission costs, lower reliability, limited network management and diagnostics, and hidden inefficiencies.

It soon became apparent that a better approach to LAN internetworking was to connect bridges and routers into a reliable, manageable WAN backbone designed to make the best use of facilities and offer the high performance users demanded.

Frame relay technology offered distinct advantages for the wide area network. First, it was a more efficient WAN protocol than IP, using only five bytes of overhead versus 20 for IP. In addition, frame relay was easily switched. IP switching was not widely available in the WAN, and IP routing added unnecessary delays and consumed more bandwidth in the network.

Widespread Digital Facilities

As the public telecommunications infrastructure migrated from analog facilities to high quality digital facilities, bandwidth availability increased and error rates decreased. The error-correcting capabilities of X.25 and SNA, which were developed to cope with the inherent errors of analog lines, were no longer necessary in digital wide area networks.

In the Beginning

While telecommunications managers contemplated the task of how to manage growing user requirements and increased network complexity, frame relay was being conceived in Bell Labs as part of the ISDN specification. Soon, frame relay evolved into a network service in its own right.

In 1990, four companies collaborated to refine the frame relay specification. "The Gang of Four," as they were known, later formed the Frame Relay Forum, which was incorporated in 1991. Since its inception, the Frame Relay Forum has grown to more than 300 members, evidence of widespread acceptance of frame relay as the method of choice for high-speed networks.

We'll discuss the work of the Frame Relay Forum in more detail in Chapter 4. Now, let's focus on the benefits of frame relay as a technology.

Banking on Frame Relay

The success of a new technology is often dependent upon compelling economic reasons for implementation. In the years since its inception, users of frame relay have found that it provides a number of benefits over alternative technologies:

1. lower cost of ownership
2. well-established and widely-adopted standards that allow open architecture and plug-and-play service implementation
3. low overhead, combined with high reliability
4. network scalability, flexibility and disaster recovery
5. interworking with other new services and applications, such as ATM

Cost of Ownership

Frame relay provides users a lower cost of ownership than competing technologies for a number of reasons:

- It supports multiple user applications, such as TCP/IP, NetBIOS, SNA and voice, eliminating multiple private lines to support different applications at a single site.
- It allows multiple users at a location to access a single circuit and frame relay port, and it efficiently uses bandwidth, thanks to its statistical multiplexing capability.
- Because only a single access circuit and port are required for each site, tremendous savings are often realized in recurring costs of transmission facilities.
- Customers realize a significant reduction in hardware, such as the number of router cards and DSU/CSUs required, reducing up-front costs and on-going maintenance costs when compared with point-to-point technologies.

Standards

Well established, widely-adopted standards are key to equipment interoperability and efficient use of capital acquisition funds.

With frame relay, users can relax knowing that frame relay standards are in place both in the United States and around the world. This ensures that equipment and services provided today will be functional for the long run, with constantly evolving standards to support new applications and to meet dynamic market place needs.

Low Overhead and High Reliability

By using only two to five bytes of overhead, frame relay makes efficient use of each frame. This means that more of the frame relay bandwidth is used for sending user data and less for overhead. Bandwidth utilization of frame relay is nearly equivalent to leased lines and better than numerous other technologies, such as X.25 or IP switching.

When the effects are spread over a large network with numerous sites, results improve exponentially:

- Simplified switching means less delay.
- Statistical multiplexing leads to more efficient bandwidth use.
- Low overhead means bandwidth is used only for user data, not for data transport.

Network Scalability, Flexibility and Disaster Recovery

To the end user, a frame relay network appears straightforward: one user simply connects directly to the frame relay cloud. A frame relay network is based on virtual circuits which may be meshed or point-to-point, and these links may be permanent or switched. (See Chapter 2 for more details.)

Because of this structure, frame relay is more easily scalable than a fixed point-to-point network. This means that additions and changes in a network are transparent to end users, giving telecommunications managers the flexibility to modify network topologies easily and scale networks as applications grow and sites are added.

This inherent flexibility lends itself equally well to the provision of alternate routes to disaster recovery sites, which are, in many cases, transparent to the end user.

Interoperability with New Applications and Services

Compared with using point-to-point leased lines, frame relay suits meshed networks and hub and spoke networks equally well. This means that frame relay easily accommodates new applications and new directions of existing networks, for example, SNA migration to APPN.

In addition, frame relay standards have been developed to interwork with newly evolving services such as ATM. As new applications emerge and/or bandwidth requirements increase, networks can gracefully migrate to the appropriate technology without stranding existing network equipment.



Advanced Trail

If you're an experienced hiker, we've provided a little more of a challenging trail in this section. Here, you'll read why frame relay offers advantages over other technologies.

Frame Relay: The Right Mix of Technology

Frame relay combines the statistical multiplexing and port sharing features of X.25 with the high speed and low delay characteristics of TDM circuit switching. Defined as a "packet mode" service, frame relay organizes data into individually addressed units known as frames rather than placing it into fixed time slots. This gives frame relay statistical multiplexing and port sharing characteristics.

Unlike X.25, frame relay completely eliminates all Layer 3 processing. (See Figure 2.)

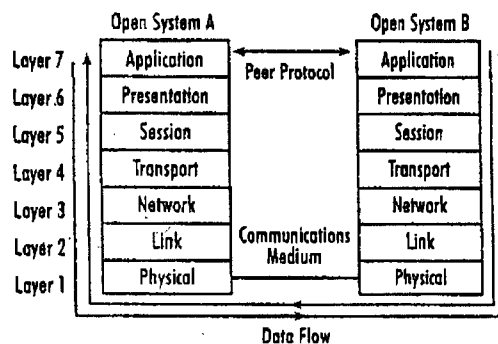


Figure 2: Open Systems Interconnection (OSI) Model

Only a few Layer 2 functions, the so-called "core aspects," are used, such as checking for a valid, error-free frame but not requesting retransmission if an error is found. Thus, many protocol functions already performed at higher levels, such as sequence numbers, window rotation, acknowledgments and supervisory frames, are not duplicated within the frame relay network.

Stripping these functions out of frame relay dramatically increases throughput (i.e., the number of frames processed per second for a given hardware cost), since each frame requires much less processing. For the same reason, frame relay delay is lower than X.25 delay, although it is higher than TDM delay, which does no processing.

In order to remove this functionality from the frame relay network, end devices must ensure error-free end-to-end transmission of data. Fortunately, most devices, especially those attached to LANs, have the intelligence and processing power to perform this function.

Table 1 summarizes the characteristics of TDM circuit switching, packet switching and frame relay

	TDM Circuit Switching	X.25 Packet Switching	Frame Relay
Time-slot multiplexing	yes	no	no
Statistical (virtual circuit) multiplexing	no	yes	yes
Port sharing	no	yes	yes
High speed (per \$)	yes	no	yes
Delay	very low	high	low

Table 1: Comparison of TDM circuit switching, packet switching, and frame relay

Frame relay uses a variable length framing structure, which, depending on user data, ranges from a few to more than a thousand characters. This feature, similar to X.25, is essential for interoperability with LANs and other synchronous data traffic, which requires variable frame size. It also means that traffic delays (although always lower than X.25) will vary, depending on frame size. Some traffic types do not tolerate delay well, especially variable delay. However, frame relay technology has been adapted to carry even delay-sensitive traffic, such as voice.



Shortcut

As the 1980's came to a close, several network trends combined to create a need for a new form of wide area network switching:

- Growth in high speed, high throughput applications
- Proliferation of end-user devices
- Increased availability of error-free, high-speed transmission lines.

This new wide area switching technology required high speed, low delay, port sharing, and bandwidth sharing on a virtual circuit basis. While existing TDM circuit switching and X.25 packet switching had some of these characteristics, only frame relay offered a full complement. These characteristics make frame relay an ideal solution for the bursty traffic sources found in LAN-WAN internetworking.

Frame relay provides a number of benefits over alternative technologies:

- lower cost of ownership
- well-established and widely-adopted standards that allow open architecture and plug-and-play service implementation
- low overhead, combined with high reliability
- network scalability, flexibility and disaster recovery
- interworking with other new services and applications, such as ATM

Frame relay offers users the ability to improve performance (response time) and reduce transmission costs dramatically for a number of important network applications. In order to be effective, frame relay requires that two conditions be met:

1. the end devices must be running an intelligent higher-layer protocol
2. the transmission lines must be virtually error-free

Other wide area network switching technologies, such as X.25 packet switching and TDM circuit switching, will remain important where line quality is not as good, when the network itself must guarantee error-free delivery, or when the traffic is intolerant of delay.

CHAPTER 2

HOW FRAME RELAY WORKS



Base Camp

In this chapter, we will discuss in more detail how frame relay works. We'll concentrate on the basic flow of data within a frame relay network.

Basic Trail: The basic trail will give beginners an overview of virtual circuits. Next, we'll show you the frame relay frame, how it's constructed and how it moves across the frame relay network. Finally, on the Basic Trail, we'll introduce the concept of discarding frames.

Advanced Trail: If you take the advanced trail, you'll find a comparison of X.25 and frame relay processing and a more detailed discussion of error recovery by higher layer protocols.



View Points:

- Figure 3: Basic frame structure of popular synchronous protocols
- Figure 4: Frame structure and header format of the frame relay frame
- Figure 5: DLCI path through the network
- Figure 6: X.25 versus frame relay processing flow chart

Shortcut: The shortcut summarizes the basic flow of data in a frame relay network.



Basic Trail

Virtual Circuits in Frame Relay

Frame relay technology is based on the concept of using virtual circuits (VCs). VCs are two-way, software-defined data paths between two ports that act as private line replacements in the network. While today there are two types of frame relay connections, switched virtual circuits (SVCs) and permanent virtual circuits (PVCs), PVCs were the original service offering. As a result, PVCs were more commonly used, but SVC products and services are growing in popularity. A more detailed discussion of SVCs and their benefits occurs in Chapter 3. For now, let's discuss the basic differences between PVCs and SVCs.

Using PVCs

PVCs are set up by a network operator – whether a private network or a service provider – via a network management system. PVCs are initially defined as a connection between two sites or endpoints. New PVCs may be added when there is a demand for new sites, additional bandwidth, alternate routing, or when new applications require existing ports to talk to one another.

PVCs are fixed paths, not available on demand or on a call-by-call basis. Although the actual path taken through the network may vary from time to time, such as when automatic rerouting takes place, the beginning and end of the circuit will not change. In this way, the PVC is like a dedicated point-to-point circuit.

PVCs are popular because they provide a cost-effective alternative to leased lines. Provisioning PVCs requires thorough planning, a knowledge of traffic patterns, and bandwidth utilization. There are fixed lead times for installation which limit the flexibility of adding service when required for short usage periods.

Using SVCs

Switched virtual circuits are available on a call-by-call basis. Establishing a call by using the SVC signaling protocol (Q.933) is comparable to normal telephone use. Users specify a destination address similar to a phone number.

Implementing SVCs in the network is more complex than

using PVCs, but is transparent to end users. First, the network must dynamically establish connections based on requests by many users (as opposed to PVCs where a central network operator configures the network). The network must quickly establish the connection and allocate bandwidth based on the user's requests. Finally, the network must track the calls and bill according to the amount of service provided.

Although SVCs were defined in the initial frame relay specifications, they were not implemented by the first carriers or vendors of frame relay. Today, applications well-suited to SVCs are driving its deployment. While PVCs offer the statistical bandwidth gain of frame relay, SVCs deliver the any-to-any connectivity that can result in network savings and flexibility. (See Chapter 3 for a more complete discussion of SVC applications and benefits.)

The Frame Relay Header and DLCI

Now that we know about virtual circuits, and the fundamental differences between PVCs and SVCs, let's take a look at the basic structure of a frame relay frame and how it accommodates other technologies.

In the most popular synchronous protocols, data is carried across a communications line in frames which are similar in structure, as shown in Figure 3.




Figure 3:
Basic frame structure for several popular synchronous communications protocols

In a frame relay frame, user data packets are not changed in any way. Frame relay simply adds a two-byte header to the frame. Figure 4 shows the frame relay frame structure and its header in more detail.




Figure 4: Frame structure and header format for frame relay.

For now, let's look at the largest portion of the header, the DLCI. The remaining six bits of the frame relay header are discussed in the next chapter.

The frame relay header contains a 10-bit number, called the Data Link Connection Identifier (DLCI). The DLCI is the frame relay virtual circuit number (with local significance) which corresponds to a particular destination. (In the case of LAN-WAN internetworking, the DLCI denotes the port to which the destination LAN is attached.) As shown in Figure 5, the routing tables at each intervening frame relay switch in the private or carrier frame relay network route the frames to the proper destination.

Note: In the figures illustrating frame relay networks, the user devices are often shown as LAN routers, since this is a common frame relay application. They could also be LAN bridges, hosts, front end processors, FRADs or any other device with a frame relay interface.




Figure 5: The DLCI denotes the port on which the destination is located

The DLCI allows data coming into a frame relay switch (often called a node) to be sent across the network using a simple, three-step process, which is shown as a flow chart in Figure 6 in this chapter.

1. Check the integrity of the frame using the Frame Check Sequence (FCS) — if it indicates an error, discard the frame.
2. Look up the DLCI in a table — if the DLCI is not defined for this link, discard the frame.
3. Relay the frame toward its destination by sending it out the port or trunk specified in the table.

Simple Rule: If there is a problem, discard the data

In order to simplify frame relay as much as possible, one simple rule exists: if there is any problem with a frame, simply discard it. There are two principal reasons why frame relay data might be discarded:

- detection of errors in the data
- congestion (the network is overloaded)

But how can the network discard frames without destroying the integrity of the communications? The answer lies in the existence of intelligence in the endpoint devices, such as PCs, workstations, and hosts. These endpoint devices operate with multilevel protocols which detect and recover from loss of data

in the network. Incidentally, this concept of using intelligent upper layer protocols to make up for a backbone network is not a new idea. The Internet relies on this method to ensure reliable communication across the network.

If you're interested in a more detailed discussion of how the upper layer protocols recover from the loss of a frame and the causes of discarded frames, continue on to the advanced trail.

If you prefer to go right to Chapter 3, you'll find a discussion of how the frame relay network handles congestion and frame discards.



Advanced Trail

Processing: Frame Relay Versus X.25

The frame relay node processes data in a relatively simple manner compared to more fully-featured protocols like X.25. Figure 6 contrasts the simplicity of frame relay with the more complex processing of X.25. (For the sake of simplicity, the diagram reflects the path of a valid data packet. Showing the steps for error recovery and non-information frame processing for X.25 would make it much more complicated.)

Recovery by Higher Layer Protocol

As you can see in Figure 6, frame relay technology simplifies the processing task, and it relies on the endpoint devices to compensate for frame loss.

How does an upper layer protocol recover from the loss of a frame? It keeps track of the sequence numbers of the various frames it sends and receives. Acknowledgments are sent to let the sending end know which frame numbers have been successfully received. If a sequence number is missing, the receiving end will request a retransmission after waiting for a "time-out" period.

In this manner, the two end devices ensure that all of the frames eventually are received without errors. This function occurs at Layer 4, the Transport Layer, in protocols like TCP/IP and OSI Transport Class 4. By contrast, X.25 networks perform this function at Layers 2 and 3, and the endpoints need not duplicate the function in Layer 4.

While the higher layers will reliably recover from frame discards, end-to-end recovery is costly. A single lost frame will result in retransmitting all unacknowledged frames. Such recovery takes extra cycles and memory in the endpoint computers, and it uses extra network bandwidth to retransmit multiple frames.

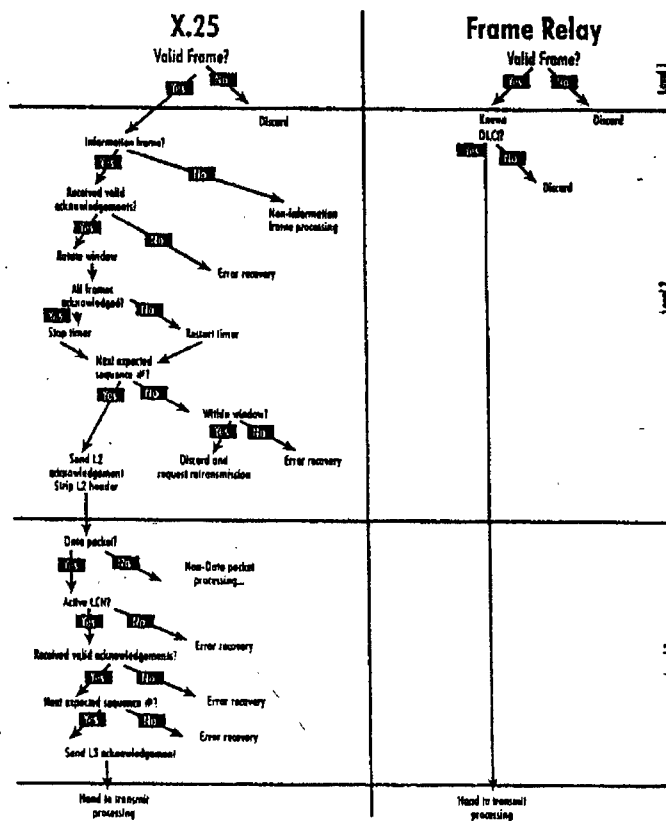


Figure 6: Simplified model of X.25 and frame relay processing

Worst of all – it causes large delays due to the higher layer time-outs (the time spent waiting for the frame to arrive before declaring it lost) and the time spent retransmitting.

Even though the higher layers can recover when discards occur, a major factor in the overall performance of a network is the ability of the network to minimize frame discards.

Two causes of frame discards are bit errors and congestion.

Frame Discards Caused by Bit Errors

When an error occurs in a frame, typically caused by noise on the line, it is detected upon receipt of the frame using the Frame Check Sequence (FCS). (See Figure 4.)

Unlike X.25, the frame relay node detecting the error does not request the sender to correct the error by retransmitting the frame. The node simply throws the frame away and moves on to receive the next frame. It relies on the intelligence of the PC or workstation that originated the data to recognize that an error has occurred and to resend the frame. Because the cost of having the higher layers recover is great, this approach would have a disastrous effect on network efficiency if the lines are noisy, generating many errors.

Fortunately, most backbone lines are based on fiber optics and experience extremely low error rates. This lowers the frequency of error-induced endpoint data recovery on lines and effectively eliminates the problem. Thus, frame relay is useful with clean, digital lines that have low error rates, while X.25 may be required for good performance on lines with higher error rates.

Frame Discards Caused by Congestion

Network congestion occurs for two reasons. First, a network node receives more frames than it can process. This is called receiver congestion. Second, a network node needs to send more frames across a given line than the speed of the line permits, which is called line congestion.

In either case, the node's buffers (temporary memory for incoming frames awaiting processing or outgoing frames lining up to be sent) are filled and the node must discard frames until the buffers have room.

Since LAN traffic is extremely bursty, the probability of congestion occurring occasionally is high unless, of course, the user excessively overconfigures both the lines and the switches – and thereby overpays on network costs. As a result, it is very important that the frame relay network have excellent congestion management features both to minimize the occurrence and severity of congestion and to minimize the effect of the discards when they are required. Congestion management features are discussed in more detail in the following chapter.



Shortcut

The basic flow of data in a frame relay network can best be described in a series of key points:

- Data is sent through a frame relay network using a data link connection identifier (DLCI), which specifies the frame's destination.
- If the network has a problem handling a frame due to line errors or congestion, it simply discards the frame.
- The frame relay network does no error correction; instead, it relies on the higher layer protocols in the intelligent user devices to recover by retransmitting the lost frames.
- Error recovery by the higher layer protocols, although automatic and reliable, is costly in terms of delay, processing and bandwidth; thus, it is imperative that the network minimize the occurrence of discards.
- Frame relay requires lines with low error rates to achieve good performance.
- On clean lines, congestion is by far the most frequent cause of discards; thus, the network's ability to avoid and react to congestion is extremely important in determining network performance.

CHAPTER 3

FRAME RELAY SIGNALING MECHANISMS



Base Camp

In this chapter, we'll discuss how frame relay technology handles interface signaling for control. If that sounds too complicated, think of it this way: interface signaling provides information about what is happening on the network so that users can get the response time they expect and the network will have the greatest efficiency possible. Signaling mechanisms can also provide options for building different types of frame relay networks to match your applications and the performance they demand.

Basic Trail: The basic trail will acquaint you with three types of signaling mechanisms used in frame relay:

- congestion notification mechanisms
- status of the connections
- SVC signaling

Advanced Trail: If you take the advanced trail, you'll find more information about the Local Management Interface (LMI) specification, which is a connection status mechanism.



View Points:

- Figure 7: The importance of congestion management
- Figure 8: Frame relay frame showing the FECN, BECN and DE bits
- Figure 9: The use of FECN and BECN in explicit congestion notification Table 2: LMI Specifications
- Figure 10: PVC Signaling using the LMI specification
- Table 2: LMI specification

Shortcut: The shortcut will quickly cover the three types of congestion management. It will also give highlights of the two other types of interface signaling discussed in the chapter, PVC status and SVC signaling.



Basic Trail **The Need for Signaling Mechanisms**

When frame relay was first proposed, it was based on a simple rule: keep the network protocol simple and let the higher layer protocols of the end devices worry about the other problems. But upon further study, it became apparent to the standards organizations that practical implementation of frame relay in real-world networks would need to specify signaling mechanisms to address three important issues:

- Allowing the network to signal that congestion exists
- Telling the status of connections (PVCs)
- Setting up new calls (SVCs)

Although these mechanisms add complexity to frame relay, the standards have an important provision which allows basic frame relay to remain simple: the use of signaling mechanisms is optional. That is, a vendor is not required to implement these features.

Without the signaling mechanisms, the resulting frame relay interface would still be compliant with the standard and data will still flow. With the signaling mechanisms, however, the throughput of the network, the response time to users, and the efficiency of line and host usage are improved.

Let's look at how these frame relay signaling mechanisms work.

Congestion Notification Mechanisms

Congestion management mechanisms, like the other signaling mechanisms, are optional for compliance, but they will affect performance. The importance of congestion management is illustrated in Figure 7.

The traffic entering the network is called the "offered load." As the offered load increases, the actual network throughput increases linearly. The beginning of congestion is represented by Point A, when the network cannot keep up with the entering traffic and begins flow control.

If the entering traffic continues to increase, it reaches a state of severe congestion at Point B, where the actual effective throughput of the network starts to decrease due to the number of re-




Figure 7: The importance of congestion management

transmissions. This causes a given frame to be sent into the network multiple times before successfully making it through.

In severe congestion, the overall network throughput can diminish, and the only way to recover is for the user devices to reduce their traffic. For that reason, several mechanisms have been developed to notify the user devices that congestion is occurring and that they should reduce their offered load.

The network should be able to detect when it is approaching congestion (Point A) rather than waiting until Point B is reached before notifying the end devices to reduce traffic. Early notification can avoid severe congestion altogether.

The ANSI specifications are very clear about the mechanisms used to indicate the existence of congestion in the network. There are two types of mechanisms to minimize, detect and recover from congestion situations, in effect providing flow control:

- Explicit Congestion Notification
- Discard Eligibility

Another mechanism that may be employed by end user devices is implicit congestion notification.

These mechanisms use specific bits contained within the header of each frame. The location of these specific bits (FECN,

BECN and DE) are shown in Figure 8.

Let's look at how each of these mechanisms function.




Figure 8: Frame relay frame showing the FECN, BECN and DE bits

Explicit Congestion Notification (ECN) Bits

The first mechanism uses two Explicit Congestion Notification (ECN) bits in the frame relay header. They are called the Forward Explicit Congestion Notification (FECN) and the Backward Explicit Congestion Notification (BECN) bits. Figure 9 depicts the use of these bits.




Figure 9: The use of FECN and BECN in explicit congestion notification

Let's suppose Node B is approaching a congestion condition. This could be caused by a temporary peak in traffic coming into the node from various sources or by a peak in the amount of traffic on the link between B and C. Here is how forward congestion notification would occur:

- Node B would detect the onset of congestion based on internal measures such as memory buffer usage or queue length.

- Node B would signal Node C (the downstream node, toward the destination) of the congestion by changing the forward ECN (FECN) contained within the frames destined for Node C from 0 to 1.
- All interim downstream nodes, as well as the attached user device, would thus learn that congestion is occurring on the DLCI(s) affected.

Depending upon the protocols used and the capabilities of the CPE device and the network switches, it is sometimes more useful to notify the source of the traffic that there is congestion, so the source can slow down until congestion subsides. (This assumes that the source is capable of responding to receipt of the congestion notification signals.) This is called Backward Congestion notification.

This is how backward congestion notification occurs:

- Node B watches for frames coming in the other direction on the connection.
- Node B sets the backward ECN bit within those frames to signal the upstream node(s) and the attached user device.

The FECN and BECN process can take place simultaneously on multiple DLCIs in response to congestion on a given line or node, thus notifying multiple sources and destinations. The ECN bits represent an important tool for minimizing serious congestion conditions.

Implicit Congestion Notification

Some upper layer protocols, such as Transport Control Protocol (TCP), operating in the end devices have an implicit form of congestion detection. These protocols can infer that congestion is occurring by an increase in round trip delay or by detection of the loss of a frame, for example. Reliance on network traffic characteristics to indicate congestion is known as implicit congestion notification.

These upper layer protocols were developed to run effectively over networks whose capacity was undetermined. Such protocols limit the rate at which they send traffic onto the network by means of a "window," which allows only a limited number of frames to be sent before an acknowledgment is received.

When it appears that congestion is occurring, the protocol can reduce its window size, which reduces the load on the network. As congestion abates, the window size is gradually increased.

The same window-size adjustment is also the normal way for the end-user devices to respond to explicit congestion notification - FECN and BECN. The ANSI standards state that implicit and explicit congestion notification are complementary and can be used together for best results.

Discard Eligibility

Frame relay standards state that the user device should reduce its traffic in response to congestion notification. Implementation of the recommended actions by the user device will result in a decrease in the traffic into the network, thereby reducing congestion. If the user device is incapable of responding to the signaling mechanisms, it might simply ignore the congestion signal and continue to transmit data at the same rate as before. This would lead to continued or increased congestion.

In this case, how does the network protect itself? The answer is found in the basic rule of frame relay: if there is a problem, discard the data. Therefore, if congestion causes an overload, more frames will be discarded. This will lengthen response times and reduce overall network throughput, but the network will not fail.

When congestion does occur, the nodes must decide which frames to discard. The simplest approach is to select frames at random. The drawback of this approach is that it maximizes the number of endpoints which must initiate error recovery due to missing frames.

A better method is to predetermine which frames can be discarded. This approach is accomplished through the use of the Committed Information Rate (CIR). The CIR is the average information capacity of the virtual circuit. When you subscribe to or buy a frame relay service from a carrier, you specify a CIR depending on how much information capacity you think your network will need.

In each frame header, there is a bit called the Discard Eligibility (DE) bit (see Figure 8). A DE bit is set to one (1) by the CPE device or the network switch when the frame is above the CIR. When the DE bit is set to 1, it makes the frame eligible for discard in response to situations of congestion. A frame with a DE bit of 1 is discarded in advance of non-discard-eligible data (those frames with a DE bit set to zero (0)). When the discard of

DE-eligible data, by itself, is not sufficient to relieve severe congestion, additional incoming frames are discarded without regard to the setting of the DE bit.

Status of Connections (PVCs and SVCs)

The next type of optional signaling mechanism defines how the two sides of a frame relay interface (e.g., the network and the router) can communicate with each other about the status of the interface and the various PVCs on that interface.

Again, these are optional parameters. It is possible to implement a frame relay interface and pass data without implementing these parameters. This signaling mechanism simply enables you to retrieve more information about the status of your network connection.

This status information is accomplished through the use of special management frames with a unique DLCI address which may be passed between the network and the access device. These frames monitor the status of the connection and provide the following information:

- Whether the interface is still active — this is called a "keep alive" or "heartbeat" signal
- The valid DLCIs defined for that interface
- The status of each virtual circuit; for example, if it is congested or not

The connection status mechanism is termed the Local Management Interface (LMI) specification. There are currently three versions of the LMI specification:

Protocol	Specification
LMI	Frame Relay Forum Implementation Agreement (IA) FRF.1 superseded by FRF.1.1
Annex D	ANSI T1.617
Annex A	ITU Q.933 referenced in FRF.1.1

Table 2: LMI Specifications

While LMI was used colloquially for the FRF.1 IA, it may also be used as a generic term to refer to any and all of the protocols.

The revised Frame Relay Forum IA FRF.1.1 calls for the mandatory implementation of Annex A of ITU Q.933.

Each version includes a slightly different use of the management protocol. Virtually all equipment vendors support LMI and most support Annex D, while Annex A is supported by fewer vendors. To ensure interoperability when your network consists of equipment from different vendors, the same version of management protocol must be at each end of the frame relay link.

For a little history of LMI and more detail on the functions of the different versions, see the advanced trail later in this chapter.

Switched Virtual Circuits

The final signaling mechanism we will discuss is SVC signaling. Unlike the previous two signaling mechanisms discussed – congestion status and connection status – SVC signaling does not offer the network operator information about the network. Rather, SVC signaling specifications allow an alternative to permanent virtual circuits. In turn, SVC signaling must provide call setup and call disconnect. Call setup includes information about the call, such as measuring data sent, acceptance, addresses and bandwidth parameters.

SVCs can also provide opportunities for new applications and new network usage. This section will discuss those alternatives and opportunities.

SVC Implementation Agreement

Implementation Agreement FRF.4 defines the needed messages and procedures to establish an SVC. Basically, the network alerts the requested destination of the incoming call and the destination chooses whether or not to accept it. If the destination accepts, the network builds the SVC across the network switches. Once the network establishes the SVC, the two endpoints can transfer information. When the endpoints no longer need the connection, either one notifies the network to terminate the call.

SVC Benefits

While current provisions for PVCs are adequate for the vast majority of near-term applications, SVC capability is beginning to gain momentum for use in public frame relay networks and for very large private networks. With SVCs, users can request set up of virtual connections only when needed and negotiate through-

put rate and burst size depending on the application.

SVC Network Applications

Some of the benefits become clearer as we look at the various applications where SVC technology is well-suited.

Remote Connectivity

At the fringes of the network or in sites where there is little need to contact other locations, SVCs are an excellent way to provide basic connectivity cost effectively. The customer pays only for the use of the network when needed, without requiring PVCs at the user-to-network interface. This application holds a great deal of promise for remote locations accessing high-speed frame relay implementations.

Overflow Traffic

There may be times of the day or night when using the burst capability of the main PVC alone cannot satisfy the need for excess capacity. Since SVCs can be set up on an adhoc basis, they can fulfill the demands of seasonal, sporadic or finite-use traffic and offer true bandwidth on demand.

Intranets and Extranets

These two applications are compelling because they allow frame relay (with SVCs) to access the Internet territory. For customers uncomfortable with the variations in quality of the Internet, building an intranet or extranet using frame relay may be a good alternative. This represents a whole new set of services carriers could offer.

Dial Access

To access a frame relay service from a carrier, a "local loop" connects the user premises to the carrier's nearest point-of-presence (POP). This local loop can be either a leased line or a dial line. Users dialing into the frame relay network can connect to either a PVC network or an SVC network.

Disaster Recovery or Alternate Network Paths

For networks using back-up or recovery sites and alternate network paths, SVCs can provide an economical alternative to leased lines, switched services or PVCs. They provide the net-

work flexibility required when leased lines are not available or there is no time to provision PVCs.



Advanced Trail

The advanced trail will discuss two topics in more depth: the LMI specification and the SVC Implementation Agreement.

In this section, we will refer to the two major standards-setting organizations, the American National Standards Institute (ANSI) and the International Telecommunications Union - Telecommunications Services Sector (ITU-T). For more information on standards, please read Chapter 4.

LMI Specification

As you may recall if you went through the basic trail, there are three versions of the LMI specification:

FRF.1 superceded by FRF.1.1, ANSI T1.617 and ITU Q.933 referenced in FRF.1.1.

The first definition for PVC status signaling was in the LMI specification. The protocol defined for the LMI provides for a "status inquiry" message which the user device (e.g., router) can send, either simply as a "keep alive" message to inform the network that the connection to the router is still up, or as a request for a report on the status of the PVCs on that port.

The network then responds with a "status" message, either in the form of a "keep alive" response or in the form of a full report on the PVCs. (See Figure 10.) An additional optional message, "status update," is also defined which enables the network to provide an unsolicited report of a change in PVC status.

Notice that the LMI status query only provides for one-way querying and one-way response, meaning that only the user device (e.g., router) can send a "status inquiry" message, and only the network can respond with a "status" message. While this approach was simple to implement, it resulted in some limitations in functions. Using status inquiries in this manner, both sides of the interface are unable to provide the same commands and responses. Most notably, it addressed only the User Network Inter-

face (UNI) and would not work in a Network-to-Network Interface (NNI) due to the one-way communications of the interface. UNI provides the end device interface to the network.

NNI provides the ability for networks to query and respond to one another. When only UNIs are available, this could lead to problems within hybrid private/public networks, where a private network node would have a frame relay NNI interface to a public frame relay service.

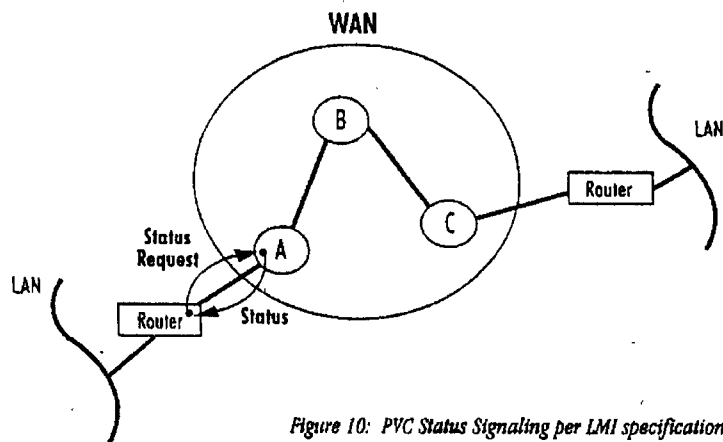


Figure 10: PVC Status Signaling per LMI specification

Therefore, just before final approval of the standard for frame relay signaling, ANSI extended the standard to provide a bi-directional mechanism for PVC status signaling that is symmetric. The bi-directional mechanism provides the ability for both sides of the interface to issue the same queries and responses. This mechanism is contained in Annex D of T1.617, known simply as Annex D. Annex D works in both the UNI and NNI interfaces.

In contrast to the LMI (which uses DLCI 1023), Annex D reserves DLCI 0 for PVC status signaling. The current requirement in FRF.1.1, Annex A signaling, is similar to Annex D and also uses DLCI 0.

To insure interoperability in a multi-vendor network environment, the same version of link management protocol must be at each end of the frame relay link.

SVC Implementation Agreement

The SVC Implementation Agreement is based on existing SVC

standards in ANSI and ITU-T. The current SVC standards are T1.617 in ANSI and Q.933 in ITU-T. These two documents lay the basis for Q.2931, the standard for access signaling for ATM (Asynchronous Transfer Mode), as well as for the PVC management procedures for frame relay.

The SVC Implementation Agreement can enable expanded service in frame relay networks. Use in internal networks involves implementing SVCs that are internal to a public or private network. The SVCs would remain transparent to the users who maintain their user-to-network interface PVCs, for example, in the case of disaster recovery. In wide area networks SVCs may be used over large geographic areas such as transatlantic applications, which have been traditionally cost-prohibitive.

ISDN and Switched Access for PVCs and SVCs

Access on demand for PVCs and SVCs, whether via Integrated Services Digital Network (ISDN) or switched access is another method to reach the frame relay network. Access on demand holds a great deal of promise for remote locations accessing high-speed frame relay implementations.

In switched access, a circuit-switched connection to the frame relay switch can be established using the existing voice network. An indication is then sent to the switch that a frame relay call is being established; the switch makes the connection and bills the call appropriately. The customer pays only for the use of the local loop when needed, without requiring PVCs at the user-to-network interface. The same benefits are true for ISDN access and E.164 addressing and lead to true, any-to-any connectivity through ISDN or switched access.



Shortcut

Interface signaling mechanisms provide information about the frame relay network so that network operators can improve efficiency.

Signaling mechanisms also provide optional ways of configuring your frame relay network to match applications usage.

There are three types of signaling mechanisms used in frame relay:

- congestion notification mechanisms
- status of the connections
- SVC signaling

The ANSI standard defines a method for the network to signal the existence of congestion called the Explicit Congestion Notification (ECN) bits.

Frame relay uses FECN (Forward ECN) and BECN (Backward ECN) bits to notify end user devices about network congestion.

Although the frame relay protocol does not respond to congestion, some higher layer protocols for end-user devices may respond to Implicit Congestion Notification by recognizing that end-to-end delays have increased or that frames have been dropped.

The use of the "discard eligibility" (DE) bit can be a powerful tool for managing throughput, including the ability to meter traffic and to guarantee a level of service.

The ANSI and ITU standards define a mechanism for communicating the status of PVCs on a frame relay interface based on a modification of the method in the LMI specification.

SVC signaling allows an alternative to permanent virtual circuits which can improve the efficiency of the network. SVCs can also provide opportunities for new applications and new network usage.

CHAPTER 4

FRAME RELAY STANDARDS AND INTEROPERABILITY



Base Camp

No discussion of frame relay would be complete without mentioning standards.

This chapter will discuss frame relay standards – those that currently exist and how they developed. We'll also acquaint you with the Frame Relay Forum and the Implementation Agreements the Forum develops to ensure frame relay interoperability.

Basic Trail: The basic trail will discuss the development of ANSI and ITU standards for frame relay. This section also gives an overview of the Frame Relay Forum and lists the current Implementation Agreements.

There is no advanced trail in this chapter.



View Points:

- Table 3: Frame relay standards
- Table 4: List of frame relay Implementation Agreements (IAs)

Shortcut: The shortcut summarizes current frame relay standards and the work of the Frame Relay Forum.



Basic Trail

How Did Frame Relay Standards Develop?

The remarkable degree of industry consensus about the need for frame relay to supplement existing switching technologies resulted in rapid development of industry standards. There are two major standards organizations which are active in this area:

- American National Standards Institute (ANSI)
- International Telecommunications Union - Telecommunications Services Sector (ITU-T), which was formerly called the Consultative Committee for International Telephone and Telegraph known as CCITT

To understand how frame relay standards developed, we need to go back to 1988. That year, ITU-T (then called CCITT) approved Recommendation I.122, "Framework for additional packet mode bearer services."

I.122 was part of a series of ISDN-related specifications. ISDN developers had been using a protocol known as Link Access Protocol - D channel (LAPD) to carry the signaling information on the "D channel" of ISDN. (LAPD is defined in ITU Recommendation Q.921.)

Developers recognized that LAPD had characteristics that could be very useful in other applications. One of these characteristics is that it has provisions for multiplexing virtual circuits at level 2, the frame level (instead of level 3, the packet level as in X.25). Therefore, I.122 was written to provide a general framework outlining how such a protocol might be used in applications other than ISDN signaling.

At that point, rapid progress began, led by an ANSI committee known as T1S1, under the auspices of the Exchange Carrier Standards Association (ECSA). This work resulted in a set of standards defining frame relay very clearly and completely. The principal frame relay standards are shown in Table 3.

Description	ANSI Standard	Status	ITU Standard	Status
Service Description	T1.606	Standard	I.233	Approved
Core Aspects	T1.618 (previously known as T1.6ca)	Standard	Q.922 Annex A	Approved
Access Signaling	T1.617 (previously known as T1.6fr)	Standard	Q.933	Approved

Table 3: Frame relay standards

T1.606 was approved early in 1990. Thanks to the hard work of the ANSI committee, coupled with a clear mandate from the market, the remaining ANSI standards sped through the stages of the standards process to receive complete approval in 1991.

Frame Relay Standards

The fast pace of frame relay standards work at ANSI was matched by an outstanding degree of cooperation and consensus in the international arena. As a result, the ITU-T recommendations for frame relay are in alignment with the ANSI standards and have also moved rapidly through the approval process. (Authors' note: although we refer to ANSI standards throughout this book, most of the discussion applies equally to the ITU-T standards.)

Interoperability and Standards Compliance

With the number of options in the standards and the range of design choices faced by vendors, what does it mean to a customer interested in interoperability?

Minimum requirements: Basic Data Handling

In order to achieve interoperability, frame relay network equipment must comply with the basic data transport method specified in the ANSI standard, which states that frame relay takes place using the DLCI in the two-byte frame relay header. This subject is covered in Chapter 2. With that relatively simple re-

quirement met, there is interoperability. The remaining requirements determine how well the network performs and whether it can be managed.

Required for real-world networks: Interface signaling

The interface control mechanisms described in Chapter 3 are optional. Data flows without them, and ignoring them is not a violation of the standard.

In real networks, however, you may find interface signaling essential to ensure that the network operates with adequate performance. Otherwise, there is no way for a network to control congestion. This means that as the traffic increases, network throughput may decrease. And throughput may continue to decrease as congestion is further exacerbated by more discards and retransmissions.

The Frame Relay Forum

The Frame Relay Forum is a non-profit organization dedicated to promoting the acceptance and implementation of frame relay based on national and international standards. Established in 1991, the Forum now has more than 300 member companies worldwide.

The Forum develops and approves Implementation Agreements (IAs) to ensure frame relay interoperability and facilitates the development of standard protocol conformance tests for various protocols. Since the earliest frame relay IAs, additional features, such as multicast, multiprotocol encapsulation and switched virtual circuit signaling, have been defined in subsequent IAs to increase the capabilities of frame relay.

Work by the Frame Relay Forum has resulted in the completion of several implementation agreements, which are listed in Table 4. Work on implementation agreements and standards is ongoing to add enhancements and broaden the applications for frame relay.

An updated listing of IAs can be found on the Frame Relay Forum web site at <<www.frforum.com>> and in the Forum's quarterly newsletters.

FRF.1.1	User-to-Network (UNI) Implementation Agreement
FRF.2.1	Frame Relay Network-to-Network (NNI) Implementation Agreement
FRF.3.1	Multiprotocol Encapsulation Implementation Agreement (MEI)
FRF.4	Switched Virtual Circuit Implementation Agreement
FRF.5	Frame Relay/ATM PVC Network Interworking Implementation Agreement
FRF.6	Frame Relay Service Customer Network Management Implementation Agreement (MIB)
FRF.7	Frame Relay PVC Multicast Service and Protocol Description Implementation Agreement
FRF.8	Frame Relay/ATM PVC Service Interworking Implementation Agreement
FRF.9	Data Compression over Frame Relay Implementation Agreement
FRF.10	Frame Relay Network-to-Network Interface SVC Implementation Agreement
FRF.11	Voice over Frame Relay Implementation Agreement
FRF.12	Frame Relay Fragmentation Implementation Agreement

Table 4: Frame Relay Forum Implementation Agreements (IAs)



Shortcut

There are two major standards organizations:

- American National Standards Institute (ANSI)
- International Telecommunications Union- Telecommunications Services Sector (ITU-T)

The initial frame relay standard was approved in 1990 by ANSI, and the remaining standards were approved by 1991. ITU recommendations for frame relay are in alignment with the ANSI standards.

In order to achieve interoperability, frame relay network equipment must comply with the basic data transport method specified in the ANSI standard, which states that frame relay takes place using the DLCI in the two-byte frame relay header.

Although interface control mechanisms are optional, they are essential to ensure that the network operates with adequate performance.

The Frame Relay Forum is a non-profit organization dedicated to promoting the acceptance and implementation of frame relay based on national and international standards.

The Forum develops and approves Implementation Agreements (IAs) to ensure frame relay interoperability. Since the earli-

est frame relay IAs, additional features, such as multicast, multiprotocol encapsulation and switched virtual circuit signaling, have been defined to increase the capabilities of frame relay.

CHAPTER 5

WHERE FRAME RELAY IS USED



Base Camp

Previous chapters covered the what and how of frame relay. By focusing on applications, this chapter provides insight into the practical benefits of frame relay. We'll discuss frame relay applications that are widely deployed and others that are emerging.

Basic Trail: The basic trail will give you an overview of four popular and growing applications: LAN peer-to-peer networking over frame relay, SNA over frame relay, voice over frame relay (VoFR) and frame relay-to-ATM interworking.

Advanced Trail: On the advanced trail, we'll go into more detail about three of the applications covered on the basic trail. Specifically, you'll also read how PRF.3.1 provides interoperability in SNA networks and how traffic is managed in SNA over frame relay applications. We'll also talk about how voice over frame relay works and the associated trade-offs. Finally, you'll find a discussion of the Interworking Function (IWF) and FUNI or ATM frame-based UNI.



View Points:

- Figure 11: Traditional Solution for LAN or Client/Server Networking
- Figure 12: Frame Relay Solution for LAN or Client/Server Networking
- Figure 13: Parallel bank branch networks
- Figure 14: Consolidated bank network

- Figure 15: Integrated voice and data network
- Figure 16: Frame/ATM Network Interworking
- Figure 17: Frame/ATM Service Interworking
- Figure 18: Typical Multidrop SNA Network
- Figure 19: SNA Network Migrated to Frame Relay
- Figure 20: NCP direct FRF.3.1 Network
- Figure 21: FRAD FRF.3.1 Network
- Figure 22: Normal Speech Components
- Figure 23: Frame/ATM Network Interworking (Encapsulation)
- Figure 24: Frame/ATM Service Interworking (Transparent)
- Figure 25: Frame/ATM Service Interworking (Translation)
- Figure 26: ATM DXI and ATM FUNI
- Table 5: Comparison of Frame Relay/ATM Interworking with FUNI and ATM DXI

Shortcut: The shortcut will give the highlights of the four applications discussed and the associated benefits.



Basic Trail

Initially, frame relay gained acceptance as a means to provide end users with a solution for LAN-to-LAN connections and to meet other data connectivity requirements. Frame relay's compelling benefit is that it lowers the cost of ownership compared to competing technologies:

- Frame relay supports multiple user applications, such as TCP/IP, NetBIOS, SNA and voice and thus eliminates the need for multiple private line facilities supporting different applications at a single site.
- Because it statistically multiplexes, frame relay allows multiple users at a location to access a single circuit and frame relay port, making efficient use of the bandwidth.
- Since only a single access circuit and port are required for each user site, users often realize tremendous savings in the cost of transmission facilities.
- Customers realize a significant reduction in the number of router cards and DSU/CSUs required, reducing up-front costs as well as ongoing maintenance compared with point-to-point technologies.

Application #1: Meshed LAN Peer-to-Peer Networking

In a traditional solution for LAN or client/server networking across a WAN, meshed network implementations can be costly. Since private line pricing is distance sensitive, the price of the network increases as geographic dispersion increases. Changes in network design normally require physical reconfigurations in addition to software changes, which increases the time to administer the changes. (See Figure 11.)

Frame Relay for LAN or Client/Server

By moving to frame relay for LAN or client/server applications, additional VCs between locations can be provisioned for minimal incremental cost. Most public frame relay pricing is distance insensitive. Virtual connections are software configurable. Changes to VCs can be done relatively quickly. This makes frame relay ideal for meshed configurations. (See figure 12.)




Figure 11: Traditional Solution for LAN or Client/Server Networking




Figure 12: Frame Relay Solution for LAN or Client/Server Networking

Application #2: SNA Over Frame Relay

Over the past few years there has been a migration of legacy traffic, such as BSC (binary synchronous communications) and Systems Network Architecture (SNA), from low speed leased lines onto frame relay services. Ratified standards from the Internet Engineering Task Force (IETF) and the Frame Relay Forum enable the encapsulation of multiple protocols, including SNA, over frame relay networks. Together, they provide a standard method of combining SNA and LAN traffic on a single frame relay link. This enables FRADs and routers, which provide network connectivity, to handle time-sensitive SNA and bursty LAN traffic simultaneously.

The integration of legacy and LAN-to-LAN traffic provides network administrators with a more efficient, flexible and cost-effective network as well as a number of other benefits:

- Simplify the network
- Leverage investment in capital equipment
- Move in SNA's stated direction, with migration strategies to distributed and peer-to-peer enterprise networks
- Dramatically lower line costs – a potential of 30 to 40 percent compared to dedicated links
- Provide up to a 40 percent increase in network utilization through frame relay's multiprotocol support
- Experience no disruption of operations – integrity and control of the network are sustained with NetView and SNMP management
- Offer high performance networking for Advanced Peer-to-Peer Networking (APPN)

Let's explore a few of these benefits. SNA installations have expanded their use of frame relay because it is a mature, proven and stable technology. Moreover, users can leverage existing capital equipment and maintain existing network management practices.

Frame relay can be integrated into SNA networks with little or no disruption. Users may migrate to frame relay without any changes to Front End Processor (FEP) hardware or software, SNA naming or network topologies.

Frame relay allows the familiar NetView tools and practices to be maintained, so there is no need to retool network operations or retrain operations staff. This allows users to migrate at their own pace to multi-vendor enterprise network management such as SNMP.

Because frame relay is consistent with SNA's stated direction, end users have the comfort of adopting a migration strategy to distributed and peer-to-peer enterprise networks to advance and optimize their SNA network. Users benefit from improved response times and session availability with higher performance than the original multidrop network.

A typical frame relay SNA application will illustrate these benefits.

Frame Relay in a Banking Application

In a large bank with many branch locations, SNA and BSC devices are co-located with LANs, resulting in parallel branch networks and high monthly WAN costs. (See Figure 13.)




Figure 13: Parallel SNA, BSC, Alarm, and LAN Branch networks

Consolidating traffic on frame relay customer premise equipment (CPE) combines serial protocol networks and LAN networks in each branch. The CPE provides an integration of legacy devices typically found at a branch with emerging devices that support client/server applications.

How does it work? The frame relay CPE consolidates SNA/SDLC (Synchronous Data Link Control) and BSC data and LAN data onto the frame relay-based WAN. This eliminates multiple single protocol leased lines connecting the branch to its host resources. It also exploits the advantages of the higher performance LAN internetwork to consolidate serial and LAN traffic, compared with low speed analog leased-line networks.

The result is better performance, greater reliability and lower costs. Because one frame relay access line can be used to reach the same number of sites as multiple leased lines, the amount of networking equipment may be reduced. Monthly telecommuni-

cations charges are also reduced and the network is simplified. Other benefits include efficient bandwidth utilization, predictable and consistent response times, enhanced session reliability and simplified network topologies and troubleshooting.

For example, a BSC-to-frame relay conversion can be performed by the frame relay CPE for BSC Automated Teller Machine locations that are supported by an SNA host at the data center. This leverages the bank's investment in capital equipment. (See Figure 14.)




Figure 14: Consolidated Bank Network

SNA Over Frame Relay Pays

Frame relay enables mission-critical SNA networks to improve performance and reduce operating costs. These savings are available because frame relay meets the response time, availability and management requirements of mission-critical applications.

If you're interested in the details of how frame relay can be used as a replacement for SDLC point-to-point networks and how FRF.3.1 provides interoperability, take a look at the SNA section in the advanced trail.

Application #3: Voice Over Frame Relay (VoFR)

Today, non-traditional uses for frame relay are beginning to emerge. One new application, voice over frame relay (VoFR), offers telecommunication and network managers the opportunity to consolidate voice and voice-band data (e.g., fax and analog modems) with data services over the frame relay network. (See Figure 15.)




Figure 15: Integrated Voice and Data Network

In migrating leased line networks to frame relay, many network administrators found a cost-effective solution for their data needs. However, since many leased-line networks also carried voice, a solution was needed to address corporate voice requirements.

With the ratification of FRF.11, a standard was established for frame relay voice transport. Among other things, FRF.11 defines standards for how vendor equipment interoperates for the transport of voice across a carrier's public frame relay network.

Maximizing Frame Relay Networks

FRF.11 enables vendors to develop standards-based equipment and services that interoperate. It also enables network managers seeking to reduce communications costs and maximize their frame relay network to consider VoFR as an option to standard voice services.

In some cases, users may find they have excess bandwidth in

their frame relay network that could efficiently support voice traffic. Other telecommunications managers may find that the incremental cost of additional frame relay bandwidth for voice traffic may be more cost-effective than standard voice services offered by local or long distance carriers.

VoFR can provide end users with a cost effective option for voice traffic transport needs between company locations. For instance, the network manager may integrate some voice channels and serial data over a frame relay connection between a branch office and corporate headquarters. By combining the voice and data traffic on a frame relay connection already in place, the user has the potential to obtain cost-effective intracompany calling and efficient use of the network bandwidth.

Because it does not significantly increase network architecture, link speeds or CIR, the integration of voice, fax and data traffic over a single access link provides a viable option for network managers and adds to the growing list of new and non-traditional applications for frame relay.

For a discussion of how VoFR works, refer to the advanced trail in this chapter.

Users are finding that frame relay offers another significant advantage: the ability to interwork with other advanced services, such as ATM.

Application #4: Frame Relay-to-ATM Interworking

Frame Relay/ATM Interworking is a viable solution that provides users with low cost access to high speed networks. Ratified by both the ATM and Frame Relay Forums, the Frame Relay/ATM PVC Interworking Implementation Agreements (IAs) provide a standards-based solution for interworking between existing or new frame relay networks and ATM networks without any changes to end user or network devices.

Why do users want to interwork frame relay and ATM? While frame relay is well suited for many applications including LAN internetworking, SNA migration and remote access, other applications, such as broadcast video and server farm support, may be better suited for ATM networks.

Users are also interested in interworking frame relay and ATM networks to protect their capital investment in existing frame relay networks and to support planned migrations from frame relay to ATM.

Frame Relay/ATM SVC Interworking IAs are currently being developed.

Frame Relay-to-ATM Interworking Standards

There are two Frame Relay/ATM Interworking IAs for PVCs, each encompassing two different types of interworking. The first one, Frame Relay/ATM Network Interworking for PVCs (FRF.5) allows network administrators to scale the backbone beyond the 45 Mbps trunks supported by frame relay. In other words, it provides the standards for ATM to become a high speed backbone for frame relay PVC users. The second one, Frame Relay/ATM Service Interworking for PVCs (FRF.8) defines the standard for frame relay PVC and ATM PVC end users or systems to communicate seamlessly.

Frame Relay/ATM Network Interworking for PVCs can be thought of as encapsulation while Frame Relay/ATM Service Interworking for PVCs is translational between the two protocols. Let's take a closer look at each standard.

Frame Relay/ATM Network Interworking for PVCs

Frame Relay/ATM Network Interworking allows Frame Relay end-user or networking devices such as FRADs or routers to communicate with each other via an ATM network employed as the backbone.

For example, SNA terminal users connected to FRADs in branch offices communicate with frame relay-attached IBM 3745 Communications Controllers located in corporate headquarters locations using a high speed ATM network as the backbone.




Figure 16: Frame Relay/ATM Network Interworking

The frame relay devices interact as if they are using frame relay for the entire connection without knowing that an ATM network is in the middle. An ATM backbone connecting multiple frame relay networks can provide scalability and high speed support

for a large number of locations and end-user devices, without requiring changes to the devices themselves.

Frame Relay/ATM Service Interworking for PVCs

Frame Relay/ATM Service Interworking enables communication between an ATM and frame relay network or end user devices. Frame Relay/ATM Service Interworking allows existing frame relay devices in the remote branch offices to communicate with end users at headquarters who are using ATM-based applications.

By enabling existing devices to access new ATM-based applications, Frame Relay/ATM Service Interworking protects the investment in existing equipment. This promotes the decoupling of client and server sides of the network, allowing each to use the resources that best meet bandwidth requirements and budget constraints.




Figure 17: Frame Relay/ATM Service Interworking

A Quick Look at IWF and FUNI

Before we leave frame relay to ATM interworking, there are two more topics we need to touch upon briefly. Then, if you're interested in more details on these topics, you can proceed to the advanced trail. The topics are Interworking Function (IWF) and Frame-based User-to-Network Interface (FUNI).

An important advantage of Frame Relay/ATM Interworking is that it provides solutions to support communications between frame relay and ATM environments without modifications to end-user devices. However, successful support of end-to-end communications in a Frame Relay/ATM Interworking environment requires performing technical functions to compensate for the differences between frame relay and ATM. These functions are defined within the Frame Relay/ATM Service and Network Interworking IAs and are provided by the IWF generally located on the switch at the boundaries of the frame relay and ATM

services. The advanced trail will discuss the responsibilities of the IWF and how it works.

FUNI was defined by the ATM Forum to provide frame-based access to ATM networks. It is an alternative to Frame Relay/ATM Service Interworking and it is most viable where the wide area infrastructure uses ATM. FUNI enables ATM quality of service levels for network throughput and delay to be maintained end-to-end, despite the fact that the access method is frame-based, rather than native or cell-based ATM.

Approved by the ATM Forum in 1995, the FUNI specification provides improved efficiency of access line bandwidth. FUNI enables users to transmit variable length (low overhead) frames rather than fixed length cells to the ATM network. The advanced trail will discuss how FUNI differs from ATM DXI (Data Exchange Interface) and the benefits of FUNI.



Advanced Trail

On the advanced trail, we'll go into more detail about three of the applications covered on the basic trail: SNA over frame relay, Voice over Frame Relay and Frame Relay/ATM Interworking.

Frame Relay as an SNA SDLC Point-to-Point Line Replacement

Traditional SNA networks are based on leased lines which connect multiple controllers to the Front End Processor (FEP). These are typically low-speed analog lines, which represent a single point of failure between user and host, as in Figure 18.




Figure 18: Typical multidrop SNA network

Multidrop leased-line networks are a familiar evil, subjecting network managers to the complexities of a multitude of leased lines. Despite advances in networking technology, many organizations continue to base their SNA mission-critical applications on multidrop private lines.

SNA networks using point-to-point, non-switched lines can be migrated from SDLC to frame relay without any changes to the existing applications or hardware. Frequently, all that is needed is an upgrade to the communications software in the controllers.

Controllers that cannot be upgraded to support frame relay may be connected to a FRAD or router for frame relay connectivity. Frame relay uses the same hardware framing as SDLC, so all SDLC line interface couplers, modems and DSU/CSUs can be used with frame relay networks.

Another item to be considered when connecting controllers to a frame relay network is how the FRADs or routers connect to the remote and host sites on the SNA network. SNA has a form of maintenance communications called polling. An SNA device responsible for a sub-area network regularly polls each downstream controller for status, inquiring if it has data to send. At the remote site, each local controller responds.

Frame relay access devices can provide local polling or a decoupled polling capability in order to provide optimal networking conditions. A frame relay access device can poll its downstream devices or respond on their behalf. This process, called spoofing, eliminates polls on the network because only data is passed end-to-end. Extracting polls increases usable network bandwidth, directly impacting network performance.

Frame relay, as a virtual private line replacement, offers straightforward migration from the complexities of multidrop leased lines to a higher performance and more cost effective network.

As shown in Figure 19, migrating SNA networks to frame relay can occur without any change to FEP hardware or software. Users can realize significantly lower monthly WAN costs, which can pay for a frame relay migration within months.




Figure 19: SNA Network Migrated to Frame Relay

Upgrading to frame relay allows fully meshed topologies for redundancy and backup without managing a large number of dedicated lines. Adding and deleting virtual connections is done via network management and service subscription versus adding and deleting hardware. For high traffic volumes at a data center, frame relay supports access speeds up to 45 Mbps (e.g., T3/E3).

Frame relay supports "one-to-many" and "many-to-many" connections over a single line, where SDLC requires a multidrop line for a "one-to-many" configuration. SNA multipoint hardware configurations must be changed to point-to-point hardware configurations to use frame relay. The changes can be chosen to provide the best economic solution by combining the positioning of frame relay switches and frame relay terminal equipment. For example, frame relay switches may be used to provide the best use of point-to-point line tariffs, and FRADs may be used to provide frame relay to SDLC interworking, where SDLC multipoint lines are less expensive.

If users want additional cost savings, other migration paths are possible. For example, the FEP can be upgraded to allow direct connections to NCP (Network Control Program) from an FRF.3.1-compatible FRAD, as shown in Figure 20. This eliminates a FRAD or router at the host, which reduces hardware costs and complexity.




Figure 20: NCP-direct FRF.3.1 Network

Alternatively, the FEP may be upgraded to a token ring connection from an SDLC line, and a FRAD provides connectivity to the host, as shown in Figure 21, eliminating hardware on the host.



Figure 21: FRAD FRF.3.1 Network

With the additional bandwidth available from frame relay, overall performance including session availability and user response time is improved as users migrate from multidrop lines which are typically 4.8/9.6 Kbps to frame relay connections of 56/64 Kbps to T1/E1.

How FRF.3.1 Provides Interoperability

Recognizing the ability of frame relay networks to carry multiple protocols, members of the Internet Engineering Task Force (IETF) developed a standardized method to encapsulate various protocols in frame relay. This multiprotocol encapsulation technique is called RFC 1490 after its IETF designation. ANSI and the Frame Relay Forum enhanced the multiprotocol encapsulation method to include support of the SNA protocols (FRF.3.1). FRF.3.1 was adopted and implemented by numerous vendors and is invaluable in multi-vendor environments. FRF.3.1 is used to carry SNA traffic across a frame relay network and may also be used to transport IP.

Protocol Encapsulation and Practical Implementation

Typically, SNA controllers, routers and FRADs encapsulate SNA as multiprotocol data as described in the Frame Relay Forum FRF.3.1 IA. SNA topologies supported across a frame relay network include:

- Intermediate Network Node (INN)
- Boundary Network Node (BNN)
- SNA Network Interconnect (SNI)
- Advanced Peer-to-Peer Networking (APPN), including High Performance Routing (HPR)
- Boundary Access Node (BAN)

FRF.3.1 specifies how to encapsulate SNA Subarea, SNA/APPN with and without HPR within the RFC 1490 multiprotocol framework. Because data is transparent to the frame relay network, it allows multiple distinct protocols to be multiplexed across a single frame relay interface. Frame relay network access nodes are responsible for converting the user data into an appropriate FRF.3.1 format for SNA and LAN traffic.

There are other alternatives to FRF.3.1 for transporting SNA over frame relay. One method uses routers to encapsulate SNA data within TCP/IP using a standard such as Data Link Switching

(DLSw) for link layer transport. The transport method a user selects depends on the application involved and the type of network equipment used.

Traffic Management Considerations

The mission-critical nature of SNA applications requires prioritization and bandwidth allocation mechanisms to avoid poor response times and SNA session failures caused by large bursts of other data traffic. One solution is to assign a higher priority to SNA data than LAN IP/IPX data if both are multiplexed over the same virtual connection. Another alternative is to send the data streams over two separate virtual connections and use the frame relay CIR mechanism to allocate bandwidth dynamically to each virtual connection.

Bandwidth allocation by percentage of CIR is a feature supported by some FRADs and routers, and it may not require separate PVCs. A smaller amount of bandwidth may be allocated (e.g., 20 percent) to the LAN traffic connection, giving the SNA traffic more frequent transmission opportunity. Further, it is recommended that both the frame relay service provider and the frame relay equipment support explicit congestion management indicators such as FECN/BECN and Discard Eligibility (DE). If these mechanisms are supported, SNA traffic flow is adjusted properly and packet discards are minimized. As with other applications carried over frame relay, congestion management plays an important role in supporting SNA applications.

Please visit the Frame Relay Forum's Web site (www.frforum.com) for a white paper on SNA over Frame Relay. This paper goes into more detail on the many options available for SNA over Frame Relay.

Voice over Frame Relay (VoFR)

Unlike most data which can tolerate delay, voice must be handled in near real time. This means that transmission and network delays must be kept small enough to remain imperceptible to the user. Until recently, packetized voice transmission was unattainable due to the voice bandwidth requirements and transmission delays associated with packet based networks.

Human speech is burdened with a tremendous amount of redundant information that is necessary for communications to occur in the natural environment, but which is not needed for a

conversation to occur. Analysis of a representative voice sample shows that only 22 percent of a typical dialog contains essential speech components that must be transmitted for complete voice clarity (see Figure 22). The balance is made up of pauses, background noise, and repetitive patterns.

Packetized voice is possible and low-bit rates are attained by analyzing and processing only the essential components of the voice sample, rather than attempting to digitize the entire voice




Figure 22: Normal speech components

sample with all its associated pauses and repetitive patterns. Current speech processing technology takes the voice digitizing process several steps further than conventional encoding methods.

VoFR Trade-offs

There are potential trade-offs when implementing VoFR. These include:

- loss of the quality commonly associated with toll traffic due to VoFR's use of voice compression
- loss of management and administrative benefits associated with carrier voice services (i.e., the loss of consolidated voice billing and invoice itemization, end user charge back capabilities, and other advanced features such as ID and accounting codes)
- lack of equipment interoperability between customer premise equipment vendors
- lack of standards defining the acceptable levels of quality for voice transport over a carrier's frame relay network

These trade-offs do not necessarily negate the value and promise of VoFR. Significant advances in digital signal processors and compression algorithms often provide voice at a level approaching

toll quality, for a fraction of the cost of public service. VoFR vendors continue to add advanced capabilities in management and administration capabilities. In addition, future industry work will also seek to develop standards which define acceptable levels of quality and performance metrics for voice transport through carriers' frame relay networks.

Please visit the Frame Relay Forum's Web site (www.frforum.com) for a white paper on Voice over Frame Relay. This paper goes into more detail on how VoFR works and the mechanics of voice compression.

More on the IWF

As we discussed on the basic trail, support of end-to-end communications in a Frame Relay/ATM network requires performing technical functions to compensate for the differences between frame relay and ATM. These functions are provided by the IWF generally located on the switch at the boundaries of the frame relay and ATM services.

Primary responsibilities for services provided by the IWF include mapping various parameters or functions between frame relay and ATM networks. These include:

- Frames or cells are formatted and delimited as appropriate.
- Discard eligibility and cell loss priority are mapped.
- Congestion indications are sent or received appropriately (frame relay's FECN is mapped to ATM's EFCI (Explicit Forward Congestion Indicator)).
- DLCI to VPI/VCI (Virtual Path Identifier/Virtual Circuit Identifier) mapping is performed.

The IWF also supports traffic management by converting ATM and frame relay traffic conformance parameters, supporting PVC management interworking via status indicators and providing upper layer user protocol encapsulation.

Frame Relay to ATM Network Interworking for PVCs may be thought of as encapsulating frame relay in ATM, since the ATM transport is transparent to the two frame relay users. The end user protocol suite remains intact. The IWF provides all mapping and encapsulation functions necessary to ensure that the service provided to the frame relay CPE is unchanged by the presence of an ATM transport. This is also sometimes referred to as frame relay transport over ATM. (See figure 23.)

Figure 24 and 25 illustrate the Interworking Function in a

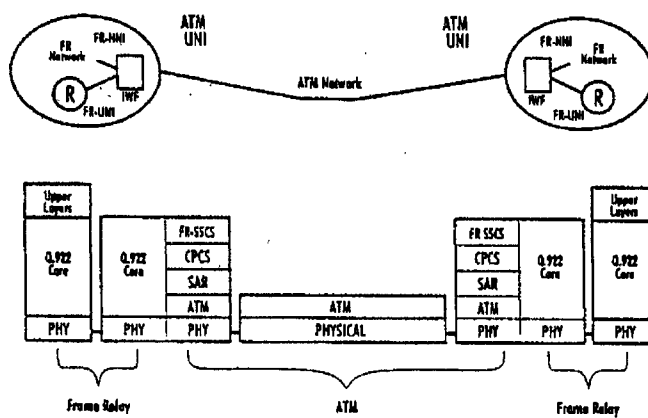


Figure 23: Frame/ATM Network Interworking (Encapsulation)

Frame Relay/ATM Service Interworking environment. To enable communications between a frame relay desktop device and the ATM based application, the IWF performs all tasks associated with mapping the frame relay User-to-Network Interface (UNI) Q.922 core-based message in the frame relay network to the ATM UNI adaptation layer in the ATM network.

Figure 24 shows Transparent Protocol Support. For encapsulation methods other than FRF.3.1 (RFC 1490) and 1483 or when a single protocol is used, the IWF forwards the data unaltered.

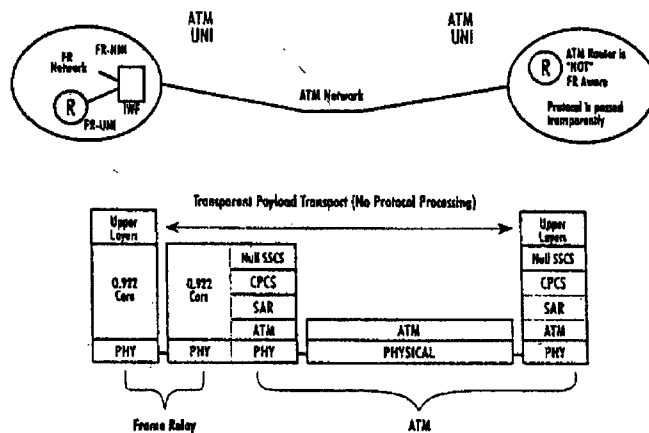


Figure 24: Frame/ATM Service Interworking (Transparent)

Transparent mode can be used when the terminal equipment on one side of the IWF uses the encapsulation method of the terminal equipment on the other side. For example, an ATM CPE may use the RFC-1490 encapsulation method which is directly compatible with the frame relay equipment and no translation is required by the IWF.

Figure 25 shows Protocol Translation Mode. Encapsulation methods for carrying multiple upper layer user protocols (e.g. LAN-to-LAN) over a frame relay PVC and an ATM PVC conform to the standard FRF.3.1 and RFC 1483, respectively. The IWF performs mapping between the two encapsulation methods. Translation Mode supports the interworking of routed and/or bridged protocols (e.g., ARP translation). For more details, please refer to the Frame Relay and Frame-Based ATM white paper on the Frame Relay Forum web site (www.frforum.com).




Figure 25: Frame/ATM Service Interworking (Translation)

How Does FUNI Work?

FUNI requires FUNI-compatible software in the user equipment and a complementary frame-based interface, as well as FUNI software in the switch to which user equipment connects. Within the switch interface, the frames are segmented into cells and sent into the network. Cells coming from the network are re-assembled into frames and sent to the user. Thus, hardware costs

of segmentation and reassembly are moved from the user equipment to the switch where it can be shared across a large number of users.

How Does FUNI Differ from ATM DXI?

Both the ATM DXI (Data Exchange Interface) and the ATM FUNI specifications translate frames of up to 2000 bytes into 53-byte ATM cells, but they differ as to where the translation takes place. The DXI standard requires a DXI-enhanced DSU to convert frames sent over an access line into cells and DXI software in the user equipment as well.

In contrast, the FUNI specification allows frames to be sent directly to the ATM switch where they are divided into cells, an approach which reduces processing and memory overhead in the remote server or workstation and makes more efficient use of the access line bandwidth. Figure 26 illustrates the two approaches.

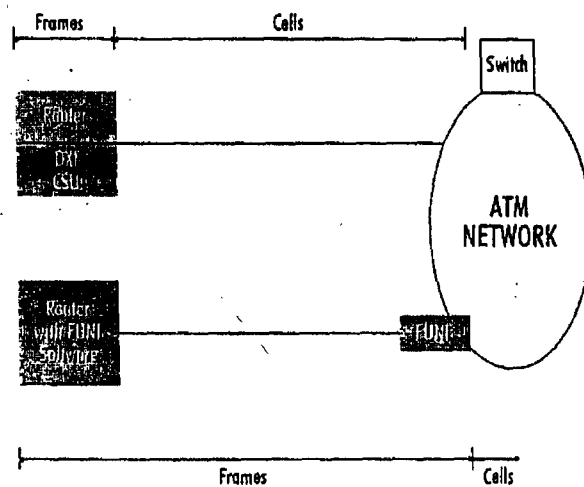


Figure 26: ATM DXI and ATM FUNI

The major benefit of ATM FUNI is that it is "ATM ready." Although limited to VBR (variable bit rate) services, it uses the same schemes as cell-based ATM UNI in the following areas:

- upper layer multiprotocol encapsulation and address resolution
- traffic parameters
- ILMI (Interim Length Management Interface)

- OAM (Operations, Administration and Maintenance) cells (future requirement)
- ATM SVC signaling (future requirement)

By contrast, frame relay requires Frame Relay/ATM Interworking functionality to achieve the same interoperability.

The current FUNI specifications address T1/E1 and Fractional T1/E1 (256 VCs per interface), whereas DXI supports full T1/E1, but not fractional T1/E1. Like frame relay and DXI, FUNI support for CPE routers consists of a software option.

Comparing Frame Relay/ATM Interworking with FUNI and ATM DXI

The following table helps to compare frame relay/ATM interworking with frame-based ATM.

	Frame Relay ATM Service Interworking	FUNI	ATM DXI
Access Transport	Frame based	Frame based	Cells
Software Requirements	IWF in frame relay or cell relay switch	FUNI software in user device and ATM network switch	DXI software in user device and DSU
Hardware Requirements	None	None	May require enhanced DSU

Table 5: Comparison of Frame Relay/ATM Interworking with FUNI and ATM DXI

When should you consider FUNI over Frame Relay-ATM service interworking? Clearly, one technology is not superior to the other. Rather, it amounts to selecting the solution which best addresses the network requirements, current topology and future network needs.

Since frame relay dominates the wide area architecture and remote site connectivity, Frame Relay/ATM Service Interworking is usually the most logical solution for most applications today. As the deployment of ATM approaches 50 percent of the wide area infrastructure, the case for deploying FUNI-based access becomes more compelling.



Shortcut

This chapter discussed four popular applications for frame relay: meshed LANs over frame relay, SNA over frame relay, voice over frame relay (VoFR) and Frame Relay/ATM Interworking.

- Frame relay enables networks to improve performance and provide cost reductions. These savings are available because frame relay meets the response time, availability and management requirements of business applications.
- Frame relay enables peer-to-peer, meshed LAN internetworking without the expense of a fully meshed leased line networking.
- Frame relay enables mission-critical SNA networks to improve performance and reduce costs. These savings are available because frame relay meets the response time, availability and management requirements of mission-critical applications.
- With compatible frame relay network access devices, most branch office equipment can connect to frame relay without hardware or software changes. With configuration changes or upgrades, additional savings are possible.
- The hidden costs of operation and maintenance are reduced by allowing network management staff to use the tools they are familiar with while providing a migration path to enterprise network management using SNMP.
- Voice over frame relay (VoFR) technology consolidates voice and voice-band data (e.g., fax and analog modems) with data services over the frame relay network. It has the potential to provide end users with greater efficiencies in the use of access bandwidth and cost-effective voice traffic transport for intra-company communications.
- Wide area network savings are possible because of the interworking of LAN, ATM, SNA and other legacy protocol, and voice traffic over frame-relay using industry standards. This may reduce branch office CIRs and port access speeds while always improving CIRs and port access speeds at the data center. This improves user response times and reduces WAN costs.
- Frame Relay/ATM Network Interworking allows frame relay end-user or networking devices such as FRADs or routers to communicate with each other via an ATM network. Thus,

frame relay devices interact as if they are using frame relay for the entire connection without knowing that an ATM network is in the middle.

- **Frame Relay/ATM Service Interworking** enables communication between ATM and frame relay network or end user devices. Enabling existing devices to access new ATM based applications allows low cost access to high speed networks while protecting the investment in existing equipment.

CHAPTER 6

PLANNING YOUR FRAME RELAY NETWORK




Base Camp

In this chapter, we will discuss the steps you need to take and the questions you should consider if you're planning a new frame relay network or planning to deploy new applications over an existing frame relay network.

Basic Trail: The basic trail will discuss four steps and several considerations to help you plan your frame relay network.

Advanced Trail: The advanced trail presents more detailed information to consider in planning a frame relay network.



Basic Trail

It doesn't matter where in the world you're located or how large your telecommunications network is – you're probably trying to do more with less and leverage your telecommunications devices and services to extend your reach.

Frame relay offers compelling advantages over today's leased line networks, including flexibility, reduced WAN costs and scalability. Migrating your leased-line network, however, requires careful planning to assess network requirements because WAN savings at the cost of network performance and reliability is not a savings at all.

Let's look at four steps and several considerations to help you plan your frame relay network.

Assess Your Network Requirements

Before you migrate your leased-line applications to a frame relay network or add new applications to your existing frame relay network, consider these questions:

1. What is the average and peak bandwidth required by your target applications?
2. What is the maximum network latency your applications can tolerate before having an impact on users?
3. What are your objectives for application and network availability?

Addressing these questions can help you assess the proper access trunk speed, Committed Information Rate (CIR), Excess Information Rate (EIR) or burst capabilities, and the optimal Service Level Agreement (SLA) from your carrier or service provider.

There are many service level management products available to enable your network to gather this type of information. These systems help you to establish your current network baselines, and they also report on ongoing network performance which can be compared against your SLAs.

Assess the Impact on Your Management Procedures

Because you are "outsourcing" a large part of your network, your network management procedures will change. Understanding these changes and how they affect application availability is critical to the success of your frame relay network. Consider these questions:

1. How will problem identification, tracking and resolution procedures change?
2. What are the responsibilities of your organization and your carrier or service provider?
3. Do you or your service provider have tools to isolate and diagnose frame relay related problems?

Examine Your Service Level Agreement

The Service Level Agreement (SLA) between you and your service provider states network performance and availability commitments. These are some guidelines:

1. Maximum network transit latency (delay): influences application response times.
2. Network availability: defining measurement of service reliability.
3. Mean time to restoral: how fast service is restored after an outage.
4. Measurement intervals: how often the service provider measures these metrics.
5. Reporting: in what form the SLA metrics are reported and how often.
6. Data delivery rate (throughput): what percentage of your data actually is delivered at the destination side.

Conduct Ongoing Capacity and Performance Planning (Service Level Management)

Launching your frame relay network is only the beginning. Changes in the organization, applications and user population necessitate a continual assessment of your networking needs. Plan ahead for network changes. Consider how these factors can impact your network:

1. Enhancements to existing applications
2. Deployment of new applications
3. Increase in user populations
4. Acquisitions or reorganization
5. Service provider and switch loading factors



Advanced Trail

Designing a Frame Relay Network

When designing a frame relay network, you should ask yourself:

Is Frame Relay Service Right for this Network?

Here are some general rules that help identify applications that are suited for frame relay service.

- **Connecting Multiple Sites:** frame relay service will most likely be beneficial when multiple sites must be connected, not just a pair of sites.
 - **High Speed:** if a network is using X.25 or a large number of analog private lines and is approaching the limits of existing bandwidth, frame relay may prove to be a cost-effective way to gain speed and efficiency.
 - **Multi-Vendor, Multi-Protocol Environment:** If the network has a multi-vendor, multi-protocol environment, frame relay service may be a good choice because of its network transparency.
 - **A Goal to Reduce Networking Costs:** If a network has been over-engineered to meet connectivity requirements, frame relay may offer a cost-effective solution.
 - **Interactive or Bursty Traffic:** frame relay is a better choice when the traffic pattern between sites is interactive or bursty.
 - **Widely Separated Locations:** frame relay is a better choice to link fairly widely separated locations, because its pricing structure is usually insensitive to distance; that is, it does not have the "mileage rates" usually attached to the private line tariffs.
- If the network passes the preliminary qualification audit, the next step is to diagram the proposed frame relay network. This should include your diagramming current configuration, labeling the locations, listing the CPE, and noting the WAN connections. This will give you a better sense of the benefits of frame relay.

Specific Network Design

Frame relay network design consists of two steps:

- Diagram your existing network. For example, is it hubbed or meshed? What are the speeds of existing connections?
- Identify traffic patterns and flow characteristics. This will help determine the bandwidth and logical port connectivity requirements.

When you have completed a first draft frame relay network design, consult with your technical support team to ensure that the diagram and applications are well matched. Keep in mind that frame relay is an interface, not an architecture, and the applications to be run on interconnected LANs or via legacy protocols must conform to a distributable architecture.



Shortcut

The major challenge in migrating your leased-line network to a frame relay service is achieving the reliability, performance and network availability your users and applications require while maximizing your WAN networking budget. With careful planning, you can achieve these goals. Four steps are helpful in planning a frame relay network:

- Assess your network requirements
 - Assess the impact on your management procedures
 - Examine your Service Level Agreement
 - Conduct ongoing capacity and performance planning
- To help identify applications that are and are not well suited for frame relay service are, examine your network for these characteristics:
- Connecting multiple sites
 - High speed
 - Multi-vendor, multi-protocol environment
 - A goal to reduce networking costs
 - Interactive or bursty traffic
 - Widely separated locations

FRAME RELAY GLOSSARY

Access Line. A communications line (e.g. circuit) interconnecting a frame-relay-compatible device (DTE) to a frame-relay switch (DCE). See also Trunk Line.

Access Rate (AR) The data rate of the user access channel. The speed of the access channel determines how rapidly (maximum rate) the end user can inject data into a frame relay network.

American National Standards Institute (ANSI) Devises and proposes recommendations for international communications standards. See also Comite Consultatif International Telegraphique et Telephonique (CCITT) and International Telecommunications Union-the Telecommunications Services Sector (ITU-T).

Asynchronous Transfer Mode (ATM) A high-bandwidth, low-delay, connection-oriented packet-like switching and multiplexing technique. Usable capacity is segmented into 53-byte fixed-size cells, consisting of header and information fields, allocated to services on demand. Also referred to as cell relay.

ATM Forum An industry organization which focuses on speeding the development, standardization and deployment of Asynchronous Transfer Mode (ATM).

Backward Explicit Congestion Notification (BECN) A bit set by a frame relay network to notify an interface device (DTE) that congestion avoidance procedures should be initiated by the sending device.

Bandwidth The range of frequencies, expressed in Kilobits per second, that can pass over a given data transmission channel within a frame relay network. The bandwidth determines the rate at which information can be sent through a channel - the greater the bandwidth, the more information that can be sent in a given amount of time.

Bridge A device that supports LAN-to-LAN communications. Bridges may be equipped to provide frame relay support to the LAN devices they serve. A frame-relay-capable bridge encapsulates LAN frames in frame relay frames and feeds those frame relay frames to a frame relay switch for transmission across the network. A frame-relay-capable bridge also receives frame relay frames from the network, strips the frame relay frame off each LAN frame, and passes the LAN frame on to the end device. Bridges are generally used to connect local area network (LAN) segments to other LAN segments or to a wide area network (WAN). They route traffic on the Level 2 LAN protocol (e.g., the Media Access Control address), which occupies the lower sub layer of the LAN OSI data link layer. See also Router.

Burstiness In the context of a frame relay network, data that uses bandwidth only sporadically; that is, information that does not use the total bandwidth of a circuit 100 percent of the time. During pauses, channels are idle; and no traffic flows across them in either direction. Interactive and LAN-to-LAN data is bursty in nature, because it is sent intermittently, and in between data transmissions the channel experiences idle time waiting for the DTEs to respond to the transmitted data user's input of waiting for the user to send more data.

Channel Generically refers to the user access channel across which frame relay data travels. Within a given T1 or E1 physical line, a channel can be one of the following, depending on how the line is configured.

Unchannelized:

The entire T1/E1 line is considered a channel, where:

- The T1 line operates at speeds of 1.536 Mbps and is a single channel consisting of 24 T1 time slots.
- The E1 line operates at speeds of 1.984 Mbps and is a single channel consisting of 20 E1 time slots.

Channelized:

The channel is any one of N time slots within a given line, where:

- The T1 line consists of any one or more channels. Each channel is any one of 24 time slots. The T1 line operates at speeds in multiples of 56/64 Kbps to 1.536 Mbps, with aggregate speed not exceeding 1.536 Mbps.

- The E1 line consists of one or more channels. Each channel is any one of 31 time slots. The E1 line operates at speeds in multiples of 64 Kbps to 1.984 Mbps, with aggregate speed not exceeding 1.984 Mbps.

Fractional:

The T1/E1 channel is one of the following groupings of consecutively or nonconsecutively assigned time slots:

- N T1 time slots (NX56/64Kbps where N = 1 to 23 T1 time slots per FT1 channel).
- + N E1 time slots (NX64Kbps, where N = 1 to 30 time slots per E1 channel).

Channel Service Unit (CSU)

An ancillary device needed to adapt the V35 interface on a frame relay DTE to the T1 (or E1) interface on a frame relay switch. The T1 (or E1) signal format on the frame relay switch is not compatible with the V35 interface on the DTE; therefore, a CSU or similar device, placed between the DTE and the frame relay switch, is needed to perform the required conversion.

Committed Burst Size (Bc) The maximum amount of data (in bits) that the network agrees to transfer, under normal conditions, during a time interval T_c . See also Excess Burst Size (Be).

Comite Consultatif International Telegraphique et Telephonique (CCITT) International Consultative Committee for Telegraphy and Telephony, a standards organization that devises and proposes recommendations for international communications. The CCITT is now known as the ITU-T, the International Telecommunications Union-the Telecommunications Services Sector. See also American National Standards Institute (ANSI) and the International Telecommunications Union (ITU-T).

Committed Information Rate (CIR) The committed rate (in bits per second) at which the ingress access interface trunk interfaces, and egress access interface of a frame relay network transfer information to the destination frame relay end system under normal conditions. The rate is averaged over a minimum time interval T_c .

Committed Rate Measurement Interval (Tc) The time interval during which the user can send only Bc-committed amount of data and Be excess amount of data. In general, the duration of Tc is proportional to the "burstiness" of the traffic. Tc is computed (from the subscription parameters of CIR and Bc) as $Tc = Bc/CIR$. Tc is not a periodic time interval. Instead, it is used only to measure incoming data, during which it acts like a sliding window. Incoming data triggers the Tc interval, which continues until it completes its committed duration. See also Committed Information Rate (CIR) and committed Burst Size (Bc).

Cyclic Redundancy Check (CRC) A computational means to ensure the accuracy of frames transmitted between devices in a frame relay network. The mathematical function is computed, before the frame is transmitted, at the originating device. Its numerical value is computed based on the content of the frame. This value is compared with a recomputed value of the function at the destination device. See also Frame Check Sequence (FCS).

Data Communications Equipment (DCE) Term defined by both frame relay and X.25 committees, that applies to switching equipment and is distinguished from the devices that attach to the network (DTE). Also see DTE.

Data Link Connection Identifier (DLCI) A unique number assigned to a PVC end point in a frame relay network. Identifies a particular PVC endpoint within a user's access channel in a frame relay network and has local significance only to that port.

Discard Eligibility (DE) A bit indicating that a frame may be discarded in preference to other frames if congestion occurs to maintain the committed information rate. See also Excess burst Size (Be) and CIR.

Egress Frame relay frames leaving a frame relay network in the direction toward the destination device. Contrast with Ingress.

End Device The ultimate source or destination of data flowing through a frame relay network sometime referred to as a Data Terminal Equipment (DTE). As a source device, it sends data to an interface device for encapsulation in a frame relay frame. As a destination device, it receives de-encapsulated data (i.e., the frame relay frame is stripped off, leaving only the user's data) from the interface device. Also see DCE

NOTE: An end device can be an application program or some operator-controlled device (e.g., workstation). In a LAN environment, the end device could be a file server or host.

Encapsulation A process by which an interface device places an end device's protocol-specific frames inside a frame relay frame. The network accepts only frames formatted specifically for frame relay; hence, interface devices acting as interfaces to a frame relay network must perform encapsulation. See also Interface device or Frame-Relay-Capable Interface Device.

Excess Burst Size (Be) The maximum amount of uncommitted data (in bits) in excess of Bc that a frame relay network can attempt to deliver during a time interval Tc. This data (Be) generally is delivered with a lower probability than Bc. The network treats Be data as discard eligible. See also Committed burst Size (Bc).

E1 Transmission rate of 2.048 Mbps on E1 communications lines. An E1 facility carries a 2.048 Mbps digital signal. See also T1 and channel.

File Server In the context of frame relay network supporting LAN-to-LAN communications, a device servicing a series of workstations within a given LAN.

Forward Explicit Congestion Notification (FECN) A bit set by a frame relay network to notify an interface device (DTE) that congestion avoidance procedures should be initiated by the receiving device. See also BECN.

Frame Check Sequence (FCS) The standard 16-bit cyclic redundancy check used for HDLC and frame relay frames. The FCS detects bit errors occurring in the bits of the frame between the opening flag and the FCS, and is only effective in detecting errors in frames no larger than 4096 octets. See also Cyclic Redundancy Check (CRC).

Frame Relay Access Device (FRAD) A device that is responsible for framing data with header and trailer information (control information) prior to presentation of the frame to the frame relay switch. On the receiving end, the FRAD strips away the frame relay control information so that the target device is presented with the data in its original form. A FRAD may be a standalone device or it may be embedded in a router, switch, multiplexer or similar device.

Frame-Relay-Capable Interface Device A communications device that performs encapsulation or a device with an integral FRAD. Frame-relay-capable routers and bridges are examples of interface devices used to interface the customer's equipment to a frame relay network. See also Interface Device and Encapsulation.

Frame Relay Forum Worldwide organization of frame relay equipment vendors, service providers, end users and consultants working to speed the development and deployment of frame relay. Web site address: www.frforum.com.

Frame Relay Frame A variable-length unit of data, in frame-relay format that is transmitted through a frame relay network as pure data. Contrast with Packet. See also Q.922A.

Frame Relay Network A telecommunications network based on frame relay technology. Data is multiplexed. Contrast with Packet-Switching Network.

High Level Data Link control (HDLC) A generic link-level communications protocol developed by the International Organization for Standardization (ISO). HDLC manages synchronous, code-transparent, serial information transfer over a link connection. See also Synchronous Data Link Control (SDLC).

Hop A single trunk line between two switches in a frame relay network. An established PVC consists of a certain number of hops, spanning the distance from the ingress access interface to the egress access interface within the network.

Host Computer A communications device that enables users to run applications programs to perform such functions as text editing, program execution, access to data bases, etc.

Ingress Frame relay frames from an access device toward the frame relay network. Contrast with Egress.

Interface Device Provides the interface between the end device(s) and a frame relay network by encapsulating the user's native protocol in frame relay frames and sending the frames across the frame relay backbone. See also Encapsulation and Frame-Relay-Capable Interface Device.

International Telecommunications Union-the Telecommunications Services Sector (ITU-T) Formerly known as the Comite Consultatif International Telegraphique et Telephonique (CCITT), the ITU-T is a standards organization that devises and proposes recommendations for international communications. See also Comite Consultatif International Telegraphique et Telephonique (CCITT).

Latency The time it takes for information to get through a network, sometimes referred to as delay.

Link Access Procedure Balanced (LAPB) The balanced-mode, enhanced, version of HDLC. Used in X.25 packet-switching networks. Contrast with LAPD.

Link Access Procedure on the D-channel (LAPD) A protocol that operates at the data link layer (layer 2) of the OSI architecture. LAPD is used to convey information between layer 3 entities across the frame relay network. The D-channel carries signaling information for circuit switching. Contrast with LAPB.

Local Area Network (LAN) A privately owned network that offers high-speed communications channels to connect information processing equipment in a limited geographic area.

LAN Protocols A range of LAN protocols supported by a frame relay network, including Transmission Control Protocol/Internet Protocol (TCP/IP), Apple Talk, Xerox Network System (XNS), Internetwork Packet Exchange (IPX), and Common Operating System used by DOS-based PCs.

LAN Segment In the context of a frame relay network supporting LAN-to-LAN communications, a LAN linked to another LAN by a bridge. Bridges enable two LANs to function like a single, large LAN by passing data from one LAN segment to another. To communicate with each other, the bridged LAN segments must use the same native protocol. See also Bridge.

Local Loop The physical wires that run from the subscriber's telephone set or PBX to the telephone company central office.

Open Systems Interconnection (OSI) Model The only internationally accepted framework of standards for communication between different systems made by different vendors. Developed by the International Standards Organization (ISO).

Packet A group of fixed-length binary digits, including the data and call control signals, that are transmitted through an X.25 packet-switching network as a composite whole. The data, call control signals, and possible error control information are arranged in a predetermined format. Packets do not always travel the same pathway but are arranged in proper sequence at the destination side before forwarding the complete message to an addressee. Contrast with Frame Relay Frame.

Packet-Switching Network A telecommunications network based on packet-switching technology, wherein a transmission channel is occupied only for the duration of the transmission of the packet. Contrast with Frame Relay Network. Typically refers to an X.25 packet network.

Parameter A numerical code that controls an aspect of terminal and/or network operation. Parameters control such aspects as packet size, data transmission speed, and timing options.

Permanent Virtual Circuit (PVC) A frame relay logical link, whose endpoints and class of service are defined by network management. Analogous to an X.25 permanent virtual circuit, a PVC consists of the originating frame relay network element address, originating data link control identifier, terminating frame relay network element address, and termination data link control identifier. Originating refers to the access interface from which the PVC is initiated. Terminating refers to the access interface at which the PVC stops. Many data network customers require a PVC between two points. Data terminating equipment with a need for continuous communication use PVCs. See also Data Link Connection Identifier (DLCI).

Point-of-Presence (POP) Physical place where a long distance carrier interfaces with the network of the local exchange carrier (LEC).

Q.922 Annex A (Q.922A) The international draft standard that defines the structure of frame relay frames. Based on the Q.922A frame format developed by the CCITT. All frame relay frames entering a frame relay network automatically conform to this structure. Contrast with Link Access Procedure Balanced (LAPB).

Q.922A Frame A variable-length unit of data, formatted in frame-relay (Q.922A) format, that is transmitted through a frame relay network as pure data (i.e., it contains no flow control information). Contrast with Packet. See also Frame Relay Frame.

Router A device that supports LAN-to-LAN communications by connecting multiple LAN segments to each other or to a WAN. Routers route traffic on the Level 3 LAN protocol (e.g., the Internet Protocol (IP) address). Routers may be equipped to provide frame relay support to the LAN devices they serve. A frame-relay-capable router encapsulates LAN frames in frame relay frames and feeds the frame relay frames to a frame relay switch for transmission across the network. See also Bridge.

Statistical Multiplexing Interleaving the data input of two or more devices on a single channel or access line for transmission through a frame relay network. Interleaving of data is accomplished using the DLCI.

Synchronous Data Link Control (SDLC) A link-level communications protocol used in an International Business Machines (IBM) Systems Network Architecture (SNA) network that manages synchronous, code-transparent, serial information transfer over a link connection. SDLC is a subset of the more generic High-Level Data Link Control (HDLC) protocol developed by the International Organization for Standardization (ISO).

Switched Virtual Circuit (SVC) A virtual circuit connection established across a network on an as-needed basis and lasting only for the duration of the transfer.

T1 Transmission rate of 1.544 Mbps on T1 communications lines. A T1 facility carries a 1.544 Mbps digital signal. Also referred to as digital signal level 1 (DS-1). See also E1 and channel.

Time Division Multiplexing (TDM) A method of transmission which relies on providing bandwidth based on fixed time slots or channels, also called circuit switching. See also channel.

Trunk Line A communications line connecting two frame relay switches to each other.

Virtual Circuit A communications link - voice or data - that appears to the user to be a dedicated point-to-point circuit. A virtual circuit is referred to as a logical, rather than a physical path, for a call.

APPENDIX

RELEVANT DOCUMENTS

- IBM Frame Relay Guide (IBM GG24-4463-00)x
- Systems Network Architecture - Format and Protocol Reference Manual: Architectural Logic (IBM SC30-3112-2)
- System Network Architecture - Advanced Peer to Peer Networking: Architecture Reference (IBM SC30-3422-03)

F. 1. 2.



**Frame
Relay
Forum** 
www.frforum.com

Frame Relay Forum
39355 California Street, Suite 307
Fremont, CA 94538

119



About Us | Contact Us |

Click Above to Support Our Sponsor!

Member Login

HANDLE

GO!

PASSWORD

Remember Me
Forgot Password?

Join Us!

[Keyword Search](#)

[Product Ratings](#)

[Browse Forums](#)

[Tell A Friend](#)

Come Join Us!

Be Notified Of Responses To Your Posts

Keyword Search

Turn Off Ad Banners

One-Click Access To Your Favorite Forums

Automated Signatures On Your Posts

Tell Others About Yourself In Your Personal Profile

Best Of All, It's Free!

E-mail*

Handle

Select A Type

Password

Verify P'word

GO!

*Opt in e-mail system. Click Personal Profile after login to opt out.

Partner With Us!

"Best Of Breed" Forums Add Stickiness To Your Site

Search **GO!** Find A Forum Find An Expert

Home > Forums > Data Transmission > Data Transmissions > Frame Relay

NNI & UNI
thread588-165402

Reply	E-mail It	Back To
-----------------------	---------------------------	-------------------------

norsyam21 (ISP) Nov 16, 2001

Hi,
I have questions regarding NNI & UNI. What is the different between NNI & UNI, their advantages & disadvantages and so on.

Thank in advance.

143101 (Visitor) Nov 21, 2001

NNI stands for Network to Network Interface. This is a frame relay connection that is used to link different carrier's networks to each other. For example, Carrier A might need to point some PVC's to routers on Carrier B's network. Carrier A would need to be connected to Carrier B in some form (T-1, DS-3) in order to pass traffic between each other. An NNI is a pipe between networks.

UNI stands for User to Network Interface. This is the port on a frame relay network that a customer hooks a location into. A UNI port is the customer port(s) in a frame cloud.

[E-mail This Thread To A Friend](#)

Your Reply

Before you start Have you answered this question before?
Write your own Frame Relay FAQ and point norsyam21 to it!
(Members please login at left before posting.)

Step 1 Choose Handle The name by which you'll be known in

Step 2 Message



(Download This Button Today!)

Member Feedback...

"...This was the ONLY place that I could use to resolve the problem. So thanks once again to member TomSark and the SQL forum!..."

[More...](#)

Partners

- [TopXML](#)
- [W3Schools](#)
- [DevelopersDex](#)
- [DevGuru](#)
- [Programmers Heaven](#)
- [VisualBuilder.com](#)
- [XMLPitstop](#)
- [Code Project](#)
- [Zvon - Guide to XML](#)
- [Tek-Tips Forums](#)
- [Search Us!](#)

[Empty text area for post content]

Step 3 Options	<input checked="" type="checkbox"/> E-Mail Notification	<input checked="" type="checkbox"/> Emoticons/Smileys	<input checked="" type="checkbox"/> Proc
Step 4 Submit Post	<input type="button" value="Preview Post"/>	<input type="button" value="Submit Post"/>	

Promoting, selling and recruiting are not allowed in the forums. [Click Here to find out why.](#)

LINK TO THIS FORUM!

(Add Stickiness To Your Site By Linking To This Professionally Managed Technical

TITLE: Frame Relay Forum at Tek-Tips

URL: <http://www.tek-tips.com/gthreadminder.cfm/lev2/5/lev3/34/pid/588>

DESCRIPTION: Frame Relay technical support forum and mutual help system for computer pi
Selling and recruiting forbidden.

[Add a Site Here!](#)

[Use Us!](#)

[Use Us!](#)

[Use Us!](#)

[Use Us!](#)

Copyright © 1998-2001 Tecumseh Group, Inc. All rights reserved. Unauthorized reproduction forbidden.

A20

Disaster Recovery for Frame Relay Networks

1. The Case for Frame Relay Disaster Recovery	1
1.1. Potential losses.....	2
1.2. Notable Frame Relay failures.....	2
1.3. Steps to successful disaster recovery planning.....	2
1.4. Disaster Recovery Options.....	3
2. How Leased Line Backup Differs	3
3. Disaster Recovery Options	4
3.1. Carrier-Based Dial Backup Options.....	5
3.1.1. Multiple Frame Carriers (Duplicate Networks)	5
3.1.2. Alternate PVC Solution	7
3.1.3. Alternate Site Solution	8
3.1.4. Carrier Dial Backup Into the Cloud	9
3.2. Enterprise-Based Dial Backup Options	9
3.2.1. Router-Based	10
3.2.2. DSU/CSU-Based	12
4. Summary	11
5. ADTRAN Options for Frame Relay	13

1. The Case for Frame Relay Disaster Recovery

Everyday, more and more mission critical data is being carried over public Frame Relay networks. Even though Frame Relay is a virtual network with integrated redundancy, the network is still subject to outages from events such as switch failures, network mishaps, or simply backhoe fade (cable breakage). Any outage translates into lost revenue and productivity, which is why Wide Area Networks (WAN) downtime has become such an important issue. According to a recent GartnerGroup report* on WAN total cost of ownership, backup charges are only seven percent of annual costs_a wise investment for most companies.

1.1. Potential losses

According to a study from Infonetics Research**, the average corporation (1,000 employees with revenues between \$150 million to \$2.8 billion) has the potential to lose up to \$7.8 million to WAN downtime each year. The study also states that the major reasons for the downtime are providers' equipment, enterprise hardware problems (mainly routers), bandwidth consumption (certain segments at certain times) and cabling problems.




Figure 1: Major Reasons for Downtime

*GartnerGroup Data, WAN Total Cost of Ownership: NSP Costs for a Typical U.S. Enterprise, 31 July 1998.

**Infonetics Research, WAN Downtime and SLAs, December 1998.

1.2. Notable Frame Relay failures

One of the most memorable Frame Relay failures in recent history occurred on April 13, 1998 on the AT&T network. Even though the outage lasted less than two days, Wal-Mart lost millions of dollars in revenue because they were unable to process credit card transactions. Another notable Frame Relay failure occurred on November 8, 1998 involving UUnet Technologies. Routing problems on UUnet's backbone interrupted Internet access for 70,000 customers over several hours. While these occurrences do not happen everyday, the fact is that they do happen.

1.3. Steps to successful disaster recovery planning

Good disaster planning doesn't have to cost a fortune. A properly thought-out and executed plan could end up being the cheapest insurance your company has ever bought. Here are some general guidelines.

- 1) Understand the network architecture. Perform a full site evaluation to differentiate between sites that require 100 percent uptime and sites that do not require backup.
- 2) Determine the type of backup service needed. This is a function of site requirements. For example, sites that require full-time connection with no reduction in throughput will require a completely redundant network, while others with reduced throughput requirements may temporarily utilize a lower bandwidth.
- 3) Analyze the numbers. How much does the dial backup solution actually cost? What productivity losses would be incurred in the event of significant downtime? What is the trade-off between the price of various services and throughput?

1.4. Disaster Recovery Options

There are two basic approaches to Frame Relay backup: Carrier-based and enterprise-based (build-your-own).

Carrier Options

- Dual carriers
- Back-up PVCs
- Back-up sites
- Dial back-up into the Frame Relay network (or cloud)

Enterprise-based Options

- Dial around the cloud from a router
- Dial around the cloud from a Frame Relay-aware DSU/CSU

Each of these options will be discussed, with pros and cons listed for each. But first, here is a brief explanation of why the traditional leased line dial-around-the-cloud method will not work for Frame Relay networks.

2. How Leased Line Backup Differs

When recovery is performed on the enterprise side, the Data Service Unit (DSU) at the customer premises is equipped with the ability to perform a "smart" switching function when the primary leased circuit goes down. This same method will not work on a Frame Relay line. Here's why.




Figure 2: Point-to-Point Example

In the typical point-to-point network (see Figure 2), traffic flows directly from one location to the other. If the point-to-point network should experience a failure, the DSU will go into alarm and initiate a dial backup call through the switched network; allowing traffic to flow around the failure through an alternate path. This example illustrates an Integrated Switched Digital Network (ISDN) backup network, but the Public Switched Telephone Network (PSTN) or a Switched56 network would work as well. This methodology will not work for Frame Relay networks because it cannot accommodate the signaling embedded in the Frame Relay protocol.

In Frame Relay networks, the customer node device, typically a router or Frame Relay Access Device (FRAD), communicates with the switch via LMI (Local Management Interface) signaling. If the FRAD/router no longer senses LMI, it determines that the Frame Relay network is inaccessible and that the path is no longer valid, and shuts down the interface.

In order to properly route Frame Relay frames, data is tagged with a DLCI (Data Link Control Identifier) number or address to be read by the local switch. Once the data enters the network, DLCI numbers change depending on the route the data traffic takes. Therefore, any attempt by a DSU to dial around the Frame Relay network will fail because Frame Relay traffic will be delivered to the FRAD/router with unrecognized DLCIs, and the data cannot be routed.

So, the two major reasons why the leased line dial backup method will not work for Frame Relay are 1) DLCI routing and 2) the need to maintain LMI signaling. There are, however, numerous disaster recovery solutions designed specifically for Frame Relay networks, including carrier-based and enterprise-based options.

3. Disaster Recovery Options

3.1. Carrier-Based Dial Backup Options

Today, there are a number of options available from Frame Relay carriers to provide a disaster recovery solution. These include duplicate networks, alternate PVC (Permanent Virtual Circuit), alternate site, or switched access into the Frame Relay network.

3.1.1. Multiple Frame Carriers (Duplicate Networks)

This first approach requires the customer either 1) completely construct two identical Frame Relay networks using two different carriers or, 2) mix two carriers in a single network. This method reduces the risk of a single-carrier failure bringing down the entire network. Options to consider when using the multiple carrier approach for network redundancy include:

- Complete redundancy requires separate local loops for the different carriers.
- Multiple carriers may require Network-to-Network Interfaces (NNI) to support complete network connectivity. (Can be difficult to manage.)
- Duplicate networks may require router reconfiguration to switch over when the primary network fails.
- Mixed networks may require an additional form of dial backup to provide complete redundancy.

Real-World Example

In designing a disaster recovery solution, Pierl Imports selected the multiple frame carrier strategy. The company used AT&T's Frame Relay network to connect one distribution center, three zone offices, and 35 regional offices; and used Sprint to link all retail outlets. During the AT&T outage (see Figure 3), Pierl was able to continue accessing applications on the Sprint network which included processing credit card transactions, generating inventory reports, and accessing bridal registry data. The company had implemented ISDN dial backup for their distribution center and zone offices, but the 35 regional offices on the AT&T network were unprotected, and therefore "cut-off" for the duration of the outage.




Figure 3: Pier 1 Disaster Recovery Solution

Very resilient	Expensive, multiple lines, duplicate PVC and CPE equipment
Spreads risk across multiple carriers	Lost multi-service discount opportunities
A single failure won't disrupt the entire network	Must deal with multiple carriers and multiple bills
	Opportunities for finger pointing "It's not our network"

3.1.2. Alternate PVC Solution

A second choice for network redundancy from a carrier is to provide an alternate or secondary PVC. Some carriers refer to this method as *Growable PVCs* or *Backup PVCs*. This method ensures redundancy in your Frame Relay network by defining multiple, diverse logical circuits to each site.

As shown Figure 4, the customer's primary and backup traffic travel along the same local loop at both the host and remote sites. The redundancy is within the Frame Relay network itself. When the primary PVC is lost, a second PVC is defined through an alternate route through the network.




Figure 4: Alternate PVC Example

There are a number of options to consider when using the alternate PVC method:

- The size of the backup PVC may be less than or equal to the size of the primary PVC, depending on application bandwidth demands and circuit costs.
- The backup PVC may be used for load sharing even when the primary PVC is healthy; or, it may be inactive until the primary PVC fails.

- If the alternate PVC is inactive until required, the cut-over may not occur automatically. It may require user intervention (e.g. customer must call carrier).

This method can be implemented using a single carrier or multiple carriers. If multiple carriers are used, the opportunity for multi-service discounts is lost, and multiple local loops may be required (requiring more equipment).

If a single carrier is used, the subscriber becomes vulnerable to a carrier network failure. In developing their disaster recovery solution AAA Travel implemented this type of structure throughout their network. When AT&T's network failed last year, AAA's single carrier backup solution left their several hundred travel agencies incapacitated for nearly an entire day.

No additional equipment	Potential points of failure: network, switch, port, local loop
No monthly costs for switched services (ISDN)	Additional monthly PVC charge
Takes advantage of inherent redundancy of Frame Relay network	Customer must ensure a separate path is used for redundant PVC, possibly requiring additional CPE equipment
	Requires additional router configuration
	May require user intervention to activate PVC or to allocate enough CIR (Committed Information Rate) to be effective

3.1.3. Alternate Site Solution

Another carrier provided backup scheme is to provide secondary PVCs to a backup site. This method ensures redundancy in your Frame Relay network by redirecting traffic to a backup location if the primary location goes down as seen in Figure 5. This method is using a secondary PVC just like the previous example, but the second PVC is mapped to a completely different local loop at a different location.




Figure 5: Alternate Site Example

There are a number of options to consider when using this approach:

- Redundant equipment is obviously required at the backup site.
- Depending on the application, redundant equipment may have to be updated frequently so that it contains the same information as the primary site.
- This will probably require dedicated circuits between the main site and backup site for things like disk mirroring/replication.

This method can be implemented using a single carrier or multiple carriers. If multiple carriers are used, multi-service discounts are unavailable. If a single carrier is used, the network becomes vulnerable to a carrier network failure. A multiple carrier approach would probably require multiple local loops at each site.

PROS	CONS
Eliminates monthly switched line costs	Network subject to a single point of failure
Takes advantage of inherent redundancy of Frame Relay network	Additional monthly PVC charge
	Customer must ensure a separate path is used for redundant PVC
	Equipment cost

3.1.4. Carrier Dial Backup Into the Cloud

A final carrier approach for end-to-end network redundancy is to provide dial-up access into the Frame Relay network. This method is only available if the carrier offers switched access into the Frame Relay network in addition to dedicated access.

If the subscriber has a dedicated connection to the Frame Relay network (either DDS, FT1, or T1), and if the local loop, the Frame Relay switch, or the switch port fails, the equipment at the customer's premises initiates a dial-up call into the carrier's Frame Relay switch. Figure 6 shows ISDN as the dial backup service, but Switched 56 or PSTN (analog modem) would work as well. Typically, the router or the DSU will monitor the local loop while in dial backup mode, so that the dial-up call can be terminated when the dedicated service is restored.




Figure 6: Carrier Dial Back Example

Can provide 100% coverage	Potential points of failure: network
Does not affect remote sites	Monthly switched service charge
	Normally requires additional router configuration for second PVC map

3.2. Enterprise-Based Dial Backup Options

All carrier-based options are viable, but extremely expensive and still subject to a single point of failure. As an alternative to carrier-based solutions, the enterprise can elect to implement a build-your-own solution for disaster recovery. The most commonly implemented options are dial around the cloud through the router or DSU.

3.2.1. Router-Based

Router-based dial backup scenarios operate independently of the network. Figure 7 shows a simple Frame Relay network with routers at each site linked using a single PVC provided by the carrier. If the local loop or the Frame Relay network experiences a failure, the router must detect the failure on its Frame Relay port, and initiate a call from a second port interface.




Figure 7: Router-Based Example

Depending on the routing protocol used, it could take some time for the newly established routes to be discovered. The rediscovery time is also a function of the protocol used. Legacy protocols, such as SNA, are prone to timeouts, so most routers feature a programmed delay that prevents lapse into backup during a protocol timeout. This would further delay the rediscovery of a new route. If the Frame Relay network is no longer in use, of a Frame Relay encapsulation also has no meaning. Therefore, the router's backup interface must use a protocol other than Frame Relay, such as point-to-point protocol (PPP).

This solution requires a very intelligent router/FRAD with multiple interfaces, along with an experienced user to program it correctly. Most of today's popular routers have this capability; but older legacy networks typically do not.

One advantage of the router-based approach is that the backup network is completely separate from the primary network. So, in the case of a network outage (like AT&T's), this is a good solution. ABN-AMRO Services Co. of Chicago, a global banking services company, was able to continue doing business during the AT&T outage because of a successful ISDN dial backup plan at each of 50 offices. Each remote office had a basic rate ISDN line, while headquarters had multiple primary rate ISDN lines to accept calls from the remote offices.

When dealing with Frame Relay networks, the issue of Service Level metrics arises. Service Level Agreements (SLAs) are a series of metrics that define the quality of service for the Frame Relay network by measuring parameters such as delay, availability, and throughput. These metrics can be monitored by the customer or the carrier to guarantee Quality of Service (QoS), troubleshoot the network, or provide valuable network trend information.

In router-based applications, metrics collection is limited. The router typically has no ability to capture interval information and other statistics. When the router switches to dial backup functionality, the Frame Relay network is marked "unavailable". This leaves the carrier and customer with no information except that the network is down.


Frame Relay network independent	Complex FRAD/Router configuration
Normally an integrated solution	Secondary interface not available on all FRADs
	Limited switched digital service availability_only ISDN backup
	Minimal performance monitoring
	Lack of carrier insight into customer's network status

3.2.2. DSU/CSU-Based

The issues of router complexity, router upgradability, and the fact that the router can only collect minimal Frame Relay statistics help make the case for a DSU/CSU-based dial backup solution. This DSU/CSU can be either a FRAD or a frame-aware DSU/CSU.

A frame-aware DSU/CSU is able to monitor the performance of the Frame Relay network. These devices combine the monitoring capability of a network probe and the functionality of a DSU/CSU in a single unit. Aware of LMI and of the PVCs (or DLCIs) traveling through it, the frame aware DSU/CSU adds great value to a dial backup solution.

In the network shown in Figure 8, each site is equipped with a Frame Relay-aware DSU/CSU with integral dial backup capability. If a failure occurs either in the Frame Relay network or on the local loop, the DSU detects the failure and initiates the backup call through the public switched network. The network failure is never



Primary PVC ———

apparent to the router/FRAD. Therefore, the router is unaware that a different route is being used during dial backup conditions.

Figure 8: DSU-Based Example

Since the solution is transparent to the router, the Frame Relay-aware DSU handles the LMI and DLCI number corrections required to bypass the original PVC. These devices recognize a failure at either the physical layer or the Frame Relay layer. Some Frame Relay-aware DSU/CSUs are available with optional dial backup interfaces, allowing the DSU to make the decision to initiate the backup call upon failure and to hang up the call upon service restoral. These features are vendor dependent thus require the customer to rely on a single vendor solution.

Frame Relay network independent	Vendor dependent solution
Router/FRAD independent	Monthly switched service charge
Easily configured	
Supports ISDN, SW56, or analog backup	
Complete SLA monitoring	
Providers have link to DSU for network evaluation	

4. Summary

In summary, no network is 100 percent failure proof, and no carrier is immune from disaster. A disaster recovery plan, when well thought-out, can be inexpensive insurance to protect from potential disaster. As this paper has tried to address, there are many choices available for implementation of a Frame Relay disaster recovery solution. Whether a carrier-based or enterprise-based solution is chosen, it should be carefully determined what is expected and required. When making this determination, a user should consider every possible failure point in the network, and how to recover from it. To completely satisfy a particular need, it may be necessary to implement one or more of the disaster recovery methods presented in this paper.

5. ADTRAN Options for Frame Relay

A customer needs a Disaster Recovery solution which can recover from both physical and virtual network failures, keep down time needs to be minimized, is independent of the router and the Frame Relay network, and is cost effective. This is exactly what Safe-T-Net does. Safe-T-Net is a customer-deployed solution to dial around the frame relay network. This feature provides a complete host to remote solution for disaster recovery over Frame Relay while still providing complete network statistics, SLA

verification, and voice/data integration. Safe-T-Net is featured in ATLAS 800^{PLUS}, and the IQ Family.

ADTRAN's ATLAS integrated access system and IQ performance monitoring devices provide a complete host-to-remote solution for disaster recovery over Frame Relay, while providing complete network statistics and SLA verification. These products provide a system for automatic dial backup upon interruption of Frame Relay services.

ADTRAN IQ units monitor both the physical link and the Frame Relay signal to determine if an interruption has occurred. Once detected, the IQ unit automatically initiates a dial-up call around the Frame Relay network (see Figure 9). The host ATLAS device accepts calls from the remote IQ unit. Once connected, ATLAS merges backup traffic with the primary traffic still being received from unaffected remote sites. The router connected to the IQ unit still receives all data as Frame Relay traffic over the primary connection, allowing a virtually transparent transition. Once the failed condition has been cleared, and the Frame Relay interruption is over, the IQ unit automatically restores traffic to the primary link.




Figure 9: Example of an ADTRAN-based disaster recovery solution for monitored Frame Relay networks.

For more information about ADTRAN's Frame Relay disaster recovery solutions, call 800 615-1176 to speak to an ADTRAN applications engineer.

421

Comm

Site Search

Advanced Search

- Current Issue
- Past Issues
- Online Partners
- The Pulse
- Tutorials

Edit Services

- Editorial Beats
- Editorial Calendar
- Cyberguide
- Reader Service Info
- FAQ

Ad Services

- Media Kit
- Ad Specs
- Sales Info
- Contact Us
- Subscribe
- About Data.com

The CommWeb Magazine Network

- Call Center
- Communications
- Convergence
- Network Magazine
- tele.com
- Teleconnect

Tech Centers

- New Public Network
- Enterprise Network
- Business Telecom
- Convergence
- Wireless Solutions
- Call Centers & CRM

Tell us what you think
Take the Networking
Quick Poll

Solutions Center

Buyer's Guide

NetworkMagazine.com



**you ride
the waves**



Visit these other CommWeb channels

- [Buyer's Guide](#)
- [Lab Tests](#)
- [Case Studies](#)
- [Events](#)
- [Product Reviews](#)
- [Tutorials](#)
- [Chats/Forums](#)
- [Subscriptions](#)

- FREE N
- CTN
 - CallK
 - Cam
 - Data
 - Con
 - Wire
 - Inter
 - VAR
- Your E
- Sign**
- Get descri
- eNe

Watching Your Back

Affordable frame relay backups can keep you covered during an outage.

by Tom Nolle

Network Magazine

11/01/99, 3:00 a.m. ET

- Utilities**
- [print this article](#)
 - [e-mail this article](#)

In the past year, both MCI WorldCom and AT&T experienced major frame relay network outages. The causes were similar: a new switch software update was installed under conditions the network hardware vendor didn't anticipate. Maybe the software should have been more tolerant of operational conditions. Maybe the operators should have been more prudent. Either stance would have prevented both problems, but nevertheless the problems occurred. Now we have to learn from them.

These network failures were significant because they were systemwide. With leased-line private networks, either a line fails or a node fails. Users of leased-line private networks apply backup on a resource basis to solve the problem. With frame relay networks, or with any public data service network, we have learned a terrible lesson: while public data networks fix their own resource problems—their own line or node failures—you, the user, have to fix the total network failures.

How do you back up an entire frame relay or IP network? Bypass it. The key to effective public data network backup is to assume that the entire network will fail—and to route traffic around it altogether. That means careful selection of equipment and backup architecture.

All public data networks fit into an application protocol's network architecture. Frame relay may support an SNA network or an IP network, for example. A bypass-backup strategy relies



[Product Reviews](#)
[Lab Tests](#)
[Tutorials](#)
[Case Studies](#)

Resource Center

[In the Classroom](#)
[Career Center](#)
[TrekMail](#)
[Auctions](#)
[Polls](#)
[Chats/Forums](#)
[TechEncyclopedia](#)
[Subscriptions](#)
[Market Research](#)

Visitors Center

[Contact Us](#)
[About Us](#)
[Privacy Statement](#)
[License Agreement](#)

Other CMPNet Sites

[TechWeb](#)
[ChannelWeb](#)
[PlanetIT](#)
[EDTN](#)

[Home](#)

TechEncyclopedia

default

This page uses styles sheets!

on the fact that most of these protocol-based architectures will recognize multiple options for network connection. You can create a backup strategy by using this connection multiplicity to support both a primary and an alternate network connection.

The most obvious, but not necessarily easiest, way to create an alternate network connection is to employ multiple frame relay carriers. If critical sites have enough traffic to justify it, connect each site via two carriers and share the load among them. For most users, however, the best option will be ISDN, and ISDN backup must be tuned to the type of protocol your network uses.

ISDN TO THE RESCUE

The easiest type of network for most users to back up is a routed IP network. Nearly all routers will support a dial-up ISDN port. Give each branch or secondary-site router an ISDN basic rate dial-up line with a Terminal Adapter (TA) and equip branch routers with an ISDN port connected to the TA. At the main site, you can provide either multiple ISDN basic rate connections or one or more of the faster ISDN primary rate connections. One or two 64Kbit/sec digital channels are provided by basic rate ISDN, and 23 channels are provided by primary rate. You may want or need to have inverse multiplexing capabilities in the TA or router to use more than one 64Kbit/sec digital channel per site for backup.

When your data network fails, the ISDN connections can be manually dialed. When the connections are made, the routers will "discover" the new paths, and network service will be restored. There may be a delay during the discovery process, and the new ISDN lines may be slower than the original frame relay or IP service, but with luck you will only use the facility for a short time.

This strategy works well for IP, for Novell SPX/IPX, and for networks built on LAN bridges. It may or may not work with IBM's SNA, because the SNA devices may not recognize an alternate ISDN route. If that's the case, you may have to rethink the way you build your primary network so that you can use ISDN backup. However, most LAN-attached SNA devices can be routed via source-route bridging, so this approach should work, except for native SNA terminals like 3270s.

EPISODIC BACKUP

Technical challenges aren't the only reason users often don't

provide backup; it's also expensive. But the recent public frame relay failures may point toward a cost-containing solution: episodic backup.

To date, frame relay failures tend to be system-wide events associated with major network hardware or software upgrades. Therefore, it follows that if you can predict those upgrades and prepare a backup strategy, you can address your highest risk at a fraction of the cost of a year-round approach. To do this, you need the cooperation of your carrier.


Episodic backup requires your frame relay carrier to provide you with 60 to 90 days' notice of major software upgrades or hardware changes. The notice should be sufficient for you to "turn on" ISDN or other backup services that will protect you during the carrier's changeover. If you maintain the backup for 30 days beyond the change, you will probably manage the longest period of risk.

It's important not to overreact to the frame relay network failures this year. Users still rate their frame relay networks as the most "suitable to mission" of all the public data services—far higher than they rate the Internet. Frame relay outages are rare, which is why they get so much attention. Many network applications can tolerate outages of several days! Yes, it's inconvenient, but it's often possible to route paper transactions or defer tasks.

When choosing your backup strategy, weigh its cost against the real harm the failure could produce. That must be done objectively, not in the heat of publicity. Overreaction often makes good press but bad backup policy.

Tom Nolle is president of CIMI (www.cimicorp.com), a consulting firm for strategic technology planning. He also lectures internationally on advanced information networking and infrastructure issues. He can be reached at tnolle@cimicorp.com.

[Buyer's Guide](#) | [Product Reviews](#) | [Lab Tests](#) | [Tutorials](#)
[Case Studies](#) | [Chats/Forums](#) | [Tech Events](#) | [Subscriptions](#) | [Contact Us](#)



you ride
the waves



A22



SEARCH FOR

- business
- technology
- shopping
- entertainment
- news
- sport
- education
- what's new

join
search
products & services
help
webmail
remote mail
manage a/c
preferences
contact us

documentation

- Introduction
- Simple Access
- Connect
- CLEAR Frame
- Multi Access
 - Failure Points
 - Prerequisites
 - Multi-attach
 - Multi-home
- Roadmap
- Configuration
- Faults
- Reports

Multi-Attached and Multi-Homed Dedicated Access

Multi-Homing your connection is to connect to more than one service provider, so that your network is insulated against a fault in one of the service providers' networks.

Multi-Attaching your connection is to connect to a single service provider more than once, so that your network is protected against a fault in a single access network.

This document discusses the relative merits of both approaches, and highlights some of the issues surrounding redundancy and fault-tolerance in your Dedicated Access connection.

- Where are the possible points of failure?
- Prerequisites for multi-attaching or multi-homing
- Multi-attaching to CLIX
- Multi-homing with CLIX and a second independent service provider

Where are the possible points of failure?

CLEAR operates a highly-resilient network, and neither of these measures are really required unless you have especially high uptime requirements on a particular remote service (or a local service which you provide to others on the Internet). If you fit into this category, you may decide that one of these measures is worthwhile for you. However, before resorting to these measures, there are several questions you should ask yourself to gauge exactly what you are protecting yourself against.

Consider the business-critical application you are trying to protect. Perhaps you operate a service accessible via the web for which downtime represents lost revenue; users of this service may connect from anywhere in the world. The following is a list of elements between the client and your server which might represent points of failure, leaving your service down.

- the database which serves your web server or application
- the web or application server
- the LAN which connects the web server to your router

These are all areas in which you have direct control. You may decide to introduce mirrored disk arrays, server clusters, redundant ethernet switches

and physical diversity into your hardware to reduce the impact of a critical fault. Your software processes may include multi-tiered testing phases, with automatic roll-back in the event of software failures.

- the router which connects to your LAN and to CLIX
- the CLIX access network (e.g. CLEAR Frame, Citylink/WIX)
- the CLIX access router to which your access network connects

These areas may be protected by connecting to CLIX multiple times, especially using multiple customer-side routers and different access network types. For example, you might use two different routers, one of which connects to one CLIX access router via CLEAR Frame, the other being connected to a different CLIX access router via Citylink/WIX.

- the CLIX backbone
- connections between CLIX and external networks

CLEAR's IP backbone has been built from scratch with reliability in mind, with multiple highly-available routers deployed at every network access point, connected by a redundant mesh of ATM circuits carried over CLEAR's protected SDH/fibre-optic national infrastructure. While access routers occasionally experience scheduled downtime for maintenance reasons, the integrity of the backbone should never be compromised by planned outages or single-point equipment failures.

CLEAR also maintains multiple connections to national and international networks, configured in such a way that if a single connecting circuit is lost, connectivity is not jeopardised.

Nevertheless, if these areas are of concern, they may be protected by connecting to multiple providers, since any internal problems within a single providers' network should be transparent to the others. The added administrative and operational cost of using multiple providers is discussed below.

- external transit networks' backbones
- connections between different external transit networks
- the remote clients' service providers, network connections and computers

These are largely beyond your control to protect against. However, if your application involves many simultaneous transactions with many clients all over the world, it may be that the failure of a single client is not of particular concern to you - while the loss of *all* clients will hurt your business. Major transit networks have usually been built with a high degree of reliability, and so the chance of a single- or double-point failure in a large transit network preventing large numbers of clients from accessing your service should be low.

Prerequisites for multi-attaching or multi-homing

The selection of alternate routing paths to or from a customer network is accomplished automatically by the use of a routing protocol. The routing protocol used between CLIX access routers and the routers of customers who are multi-attached or multi-homed is

the Border Gateway Protocol, version 4 (BGP4).

Independently-operated networks which need to use BGP4 are each assigned a globally-unique "autonomous system" number (ASN) by **IANA**, or by a regional authority acting on behalf of IANA. The regional authority for the Asia-Pacific region (including New Zealand) is **APNIC**. Customers are required to obtain such an ASN directly from the regional authority.

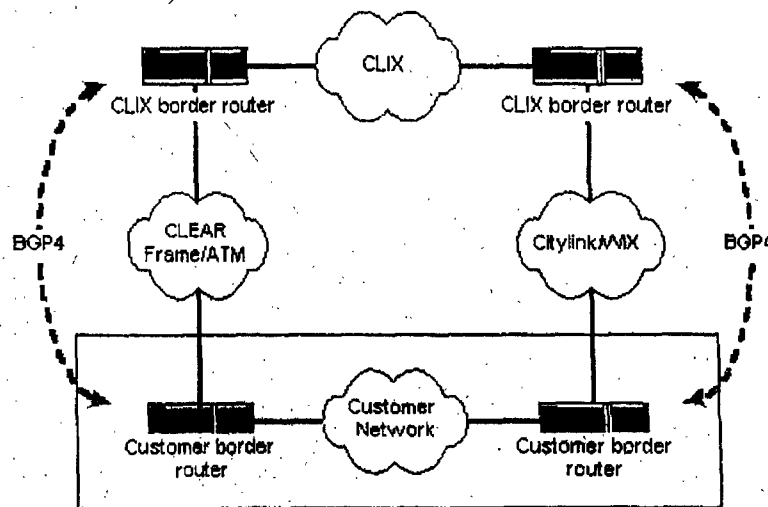
CLEAR maintains a strict policy of route filtering, based on the routing policy stored in the **Internet Routing Registry (IRR)**. This is conformant with current best practice in the network operators' community.

In summary, the following are all prerequisites for multi-attaching or multi-homing with CLIX.

- A "border" router capable of running BGP4, which will interoperate with the Cisco 7500 series border routers used by CLIX
- A globally unique ASN, obtained from an appropriate regional authority
- All customer-operated route objects should be stored in a routing registry which is part of the IRR

Multi-attaching to CLIX

Using two parallel circuits between a customer's network and different CLIX routers will satisfy most customers high-availability requirements. For optimum resilience, you should ensure that the two CLIX access circuits do not share any common elements (e.g. a single unprotected tail circuit, a single CLEAR Frame AXIS shelf, or a single mux card), and use separate routers for each access circuit, powered from separate protected power sources if possible.



Your border routers will advertise routes for your networks to both CLIX border routers using BGP4; these advertisements will be carried through the CLIX backbone and will be advertised to all external IP networks, national and international, to which CLEAR maintains connectivity.

The two IP-layer circuits between your routers and CLIX can be arranged in

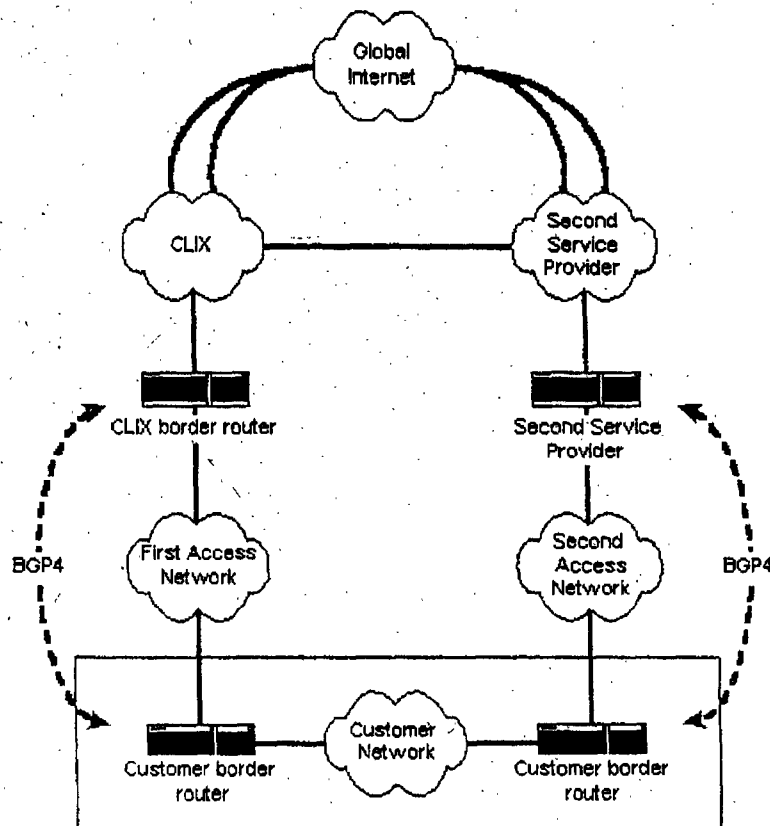
a live/backup configuration, or in a load-sharing mode - either scenario is arranged with your routers' BGP configuration, and you should be able to change between them at any time without involving CLEAR.

When multi-attached to CLIX using two routers and protected access circuits, your applications and service are insulated from planned or unplanned outages in the following elements:

- your routers
- access circuits
- frame-relay or ATM switching elements
- CLIX access routers

Multi-homing with CLIX and a second independent service provider

The arrangements described **above** can be made with different access circuits to CLIX and to a border router operated by a second, independent service provider.



This provides additional protection against a *major* fault in the CLIX backbone, or between CLIX and external networks, which may be required to meet very high service uptime requirements.

Although an arrangement such as this can be made very effectively, it is important to understand the provisioning and operational overheads which can often become more than twice as large when dealing with two

independent service providers. Load-balancing traffic over circuits to different service providers is often challenging, since both providers will often have very different connectivity to the global network. You must be confident of your in-house skills in global IP routing before deciding to multi-home.

**Business | Technology | Shopping | Entertainment | Jobs | Motoring | News | Sport | Education | What's New | Travel
Join | Search | Products & Services | Help | Webmail | Remote Access | Manage Account | Preferences | Contact
Home | Copyright**

423

Novell's[®] Complete Encyclopedia of Networking

Werner Feibel



Novell Press, San Jose

PUBLISHER: Rosalie Kearsley
EDITOR-IN-CHIEF: Dr. R. S. Langer
EXECUTIVE EDITOR, NOVELL PRESS: David Kolodney
ACQUISITIONS EDITOR: Dianne King
INSTANT INDEX CONCEPT AND DEVELOPMENTAL EDITOR: David Kolodney
EDITOR: Marilyn Smith
PROJECT EDITOR: Kristen Vanberg-Wolff
TECHNICAL EDITORS: Kelley Lindberg and Ken Neff
NOVELL TECHNICAL ADVISOR: Kelley Lindberg
BOOK DESIGNER: Helen Bruno
TECHNICAL ILLUSTRATORS: John Corrigan, Alissa Feinberg, Cuong Le, and Rick Van Genderen
PAGE LAYOUT AND TYPESETTING: Deborah Maizels
PROOFREADER/PRODUCTION ASSISTANT: Dave Nash
INDEXER: Ted Laux
COVER DESIGNER: Archer Design
LOGO DESIGN: Jennifer Gill
COVER PHOTOGRAPHER: Michael Kenna

SYBEX is a registered trademark of SYBEX Inc.

Novell Press and the Novell Press logo are trademarks of Novell, Inc.

TRADEMARKS: SYBEX and Novell have attempted throughout this book to distinguish proprietary trademarks from descriptive terms by following the capitalization style used by the manufacturer.

Every effort has been made to supply complete and accurate information. However, neither SYBEX nor Novell assumes any responsibility for its use, nor for any infringement of the intellectual property rights of third parties which would result from such use.

Copyright ©1995 SYBEX Inc., 2021 Challenger Drive, Alameda, CA 94501. Copyright ©1995, Novell Inc, 2180 Fortune Drive, San Jose, CA 95131 for DynaText encoded electronic version. World rights reserved. No part of this publication may be stored in a retrieval system, transmitted, or reproduced in any way, including but not limited to photocopy, photograph, magnetic or other record, without the prior agreement and written permission of SYBEX and Novell.

Library of Congress Card Number: 94-66403
ISBN: 0-7821-1290-0

Manufactured in the United States of America
10 9 8 7 6 5 4 3 2



internetworking unit (IWU), 466

Internetwork Link

An internetwork link serves to connect two or more networks. The networks may be identical, similar, or dissimilar. They may be located near each other or far apart. The figure "Context of internetwork links" summarizes these types of connections.

Identical networks use the same PC and network architectures and the same or comparable cabling. For example, a bridge may link two token-ring networks or a thin (10Base2) Ethernet network to a twisted-pair (10BaseT) network. These types of networks are often created for convenience. For example, an internetwork may be created to turn a large network into two smaller ones, in order to reduce network traffic.


Similar networks use the same PC architecture (for example, Intel-based) but may use different network architectures, such as Ethernet and Token Ring. Dissimilar networks use different hardware and software, such as Ethernet and an IBM mainframe.

Internetwork links differ in the level at which they operate. This difference also affects the kinds of networks they can link. The following links may be used:

- A *bridge* provides connections at the data-link layer, and it is often used to connect networks that use the same architecture. A bridge serves both as a link and as a

filter: passing messages from one network to the other, but discarding messages that are intended only for the local network. This filtering helps reduce traffic in each network.

- A *router* determines a path to a destination for a packet, and then starts the packet on its way. The destination may be in a network removed from the router by one or more intermediate networks. To determine a path, a router communicates with other routers in the larger (inter)network. Routers operate at the network layer, and most are protocol-dependent; that is, each router generally can handle only a single network-layer protocol. Special multiprotocol routers, such as Novell's Multiprotocol Router, are available. Because they need to do much more work to get a packet to its destination, routers tend to be slower than bridges.
- A *brouter* combines the features of a bridge and a router. It has the forwarding capabilities of a router, and the protocol independence of a bridge. Brouters can process packets at either the data-link or network level.
- A *gateway* moves packets between two different computer environments, such as between a local-area network and a mainframe environment or between Macintosh and PC networks. Gateways operate at the session layer and above. Because they connect dissimilar networks, gateways may need to do data translation (for example, between ASCII and EBCDIC), compression or expansion, encryption or decryption, and so on.



Context of internetwork links

- A *switch* (in this context) is a multiport bridge or gateway. Whereas a gateway connects two environments (for example, two electronic-mail systems), a mail switch can connect several such sys-

tems. Similarly, an Ethernet switch can direct packets to any of several Ethernet subnetworks to which the switch is attached.

A24

Computer Networks

Third Edition

Andrew S. Tanenbaum

*Vrije Universiteit
Amsterdam, The Netherlands*

For book and bookstore information



<http://www.prenhall.com>



Prentice Hall PTR
Upper Saddle River, New Jersey 07458

Library of Congress Cataloging in Publication Data

Tanenbaum, Andrew S. 1944-

Computer networks / Andrew S. Tanenbaum. -- 3rd ed.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-349945-6

1. Computer networks. I. Title.

TK5105.5.T36 1996

96-4121

004.6--dc20

CIP

Editorial/production manager: *Camille Trentacoste*

Interior design and composition: *Andrew S. Tanenbaum*

Cover design director: *Jerry Votta*

Cover designer: *Don Martinetti, DM Graphics, Inc.*

Cover concept: *Andrew S. Tanenbaum, from an idea by Marilyn Tremaine*

Interior graphics: *Hadel Studio*

Manufacturing manager: *Alexis R. Heydt*

Acquisitions editor: *Mary Franz*

Editorial Assistant: *Noreen Regina*



© 1996 by Prentice Hall PTR
Prentice-Hall, Inc.
A Simon & Schuster Company
Upper Saddle River, New Jersey 07458

The publisher offers discounts on this book when ordered in bulk quantities. For more information, contact:

Corporate Sales Department, Prentice Hall PTR, One Lake Street, Upper Saddle River, NJ 07458.
Phone: (800) 382-3419; Fax: (201) 236-7141. E-mail: corpsales@prehall.com

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

All product names mentioned herein are the trademarks of their respective owners.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-349945-6

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Simon & Schuster Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

5.4. INTERNETWORKING

Up until now, we have implicitly assumed that there is a single homogeneous network, with each machine using the same protocol in each layer. Unfortunately, this assumption is wildly optimistic. Many different networks exist, including LANs, MANs, and WANs. Numerous protocols are in widespread use in every layer. In the following sections we will take a careful look at the issues that arise when two or more networks are together to form an **internet**.

Considerable controversy exists about the question of whether today's abundance of network types is a temporary condition that will go away as soon as everyone realizes how wonderful [fill in your favorite network] is, or whether it is an inevitable, but permanent feature of the world that is here to stay. Having different networks invariably means having different protocols.

We believe that a variety of different networks (and thus protocols) will always be around, for the following reasons. First of all, the installed base of different networks is large and growing. Nearly all UNIX shops run TCP/IP. Many large businesses still have mainframes running SNA. DEC is still developing DECnet. Personal computer LANs often use Novell NCP/IPX or AppleTalk. ATM systems are starting to be widespread. Finally, specialized protocols are often used on satellite, cellular, and infrared networks. This trend will continue for years due to the large number of existing networks and because not all vendors perceive it in their interest for their customers to be able to easily migrate to another vendor's system.

Second, as computers and networks get cheaper, the place where decisions get made moves downward. Many companies have a policy to the effect that purchases costing over a million dollars have to be approved by top management, purchases costing over 100,000 dollars have to be approved by middle management, but purchases under 100,000 dollars can be made by department heads without any higher approval. This can easily lead to the accounting department installing an Ethernet, the engineering department installing a token bus, and the personnel department installing a token ring.

Third, different networks (e.g., ATM and wireless) have radically different technology, so it should not be surprising that as new hardware developments occur, new software will be created to fit the new hardware. For example, the average home now is like the average office ten years ago: it is full of computers that do not talk to one another. In the future, it may be commonplace for the telephone, the television set, and other appliances all to be networked together, so they can be controlled remotely. This new technology will undoubtedly bring new protocols.

As an example of how different networks interact, consider the following example. At most universities, the computer science and electrical engineering departments have their own LANs, often different. In addition, the university computer center often has a mainframe and supercomputer, the former for faculty

members in the humanities who do not wish to get into the computer maintenance business, and the latter for physicists who want to crunch numbers. As a consequence of these various networks and facilities, the following scenarios are easy to imagine:

1. LAN-LAN: A computer scientist downloading a file to engineering.
2. LAN-WAN: A computer scientist sending mail to a distant physicist.
3. WAN-WAN: Two poets exchanging sonnets.
4. LAN-WAN-LAN: Engineers at different universities communicating.

Figure 5-33 illustrates these four types of connections as dotted lines. In each case, it is necessary to insert a "black box" at the junction between two networks, to handle the necessary conversions as packets move from one network to the other.

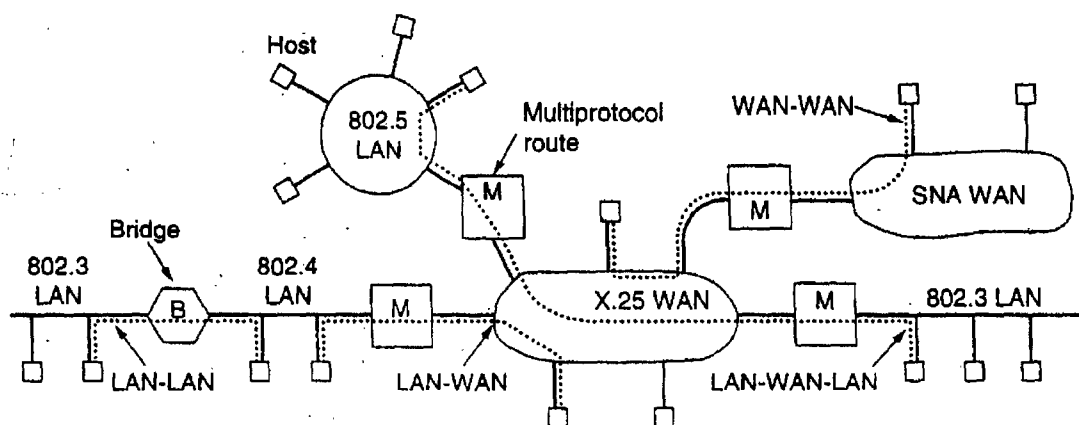


Fig. 5-33. Network interconnection.

The name used for the black box connecting two networks depends on the layer that does the work. Some common names are given below (although there is not much agreement on terminology in this area).

- Layer 1: Repeaters copy individual bits between cable segments.
- Layer 2: Bridges store and forward data link frames between LANs.
- Layer 3: Multiprotocol routers forward packets between dissimilar networks.
- Layer 4: Transport gateways connect byte streams in the transport layer.
- Above 4: Application gateways allow interworking above layer 4.

For convenience, we will sometimes use the term "gateway" to mean any device that connects two or more dissimilar networks.

Repeaters are low-level devices that just amplify or regenerate weak signals. They are needed to provide current to drive long cables. In 802.3, for example, the timing properties of the MAC protocol (the value of τ chosen) allow cables up to 2.5 km, but the transceiver chips can only provide enough power to drive 500 meters. The solution is to use repeaters to extend the cable length where that is desired.

Unlike repeaters, which copy the bits as they arrive, **bridges** are store-and-forward devices. A bridge accepts an entire frame and passes it up to the data link layer where the checksum is verified. Then the frame is sent down to the physical layer for forwarding on a different network. Bridges can make minor changes to the frame before forwarding it, such as adding or deleting some fields from the frame header. Since they are data link layer devices, they do not deal with headers at layer 3 and above and cannot make changes or decisions that depend on them.

Multiprotocol routers are conceptually similar to bridges, except that they are found in the network layer. They just take incoming packets from one line and forward them on another, just as all routers do, but the lines may belong to different networks and use different protocols (e.g., IP, IPX, and the OSI connectionless packet protocol, CLNP). Like all routers, multiprotocol routers operate at the level of the network layer.

Transport gateways make a connection between two networks at the transport layer. We will discuss this possibility later when we come to concatenated virtual circuits.

Finally, **application gateways** connect two parts of an application in the application layer. For example, to send mail from an Internet machine using the Internet mail format to an ISO MOTIS mailbox, one could send the message to a mail gateway. The mail gateway would unpack the message, convert it to MOTIS format, and then forward it on the second network using the network and transport protocols used there.

When a gateway is between two WANs run by different organizations, possibly in different countries, the joint operation of one workstation-class machine can lead to a lot of finger pointing. To eliminate these problems, a slightly different approach can be taken. The gateway is effectively ripped apart in the middle and the two parts are connected with a wire. Each of the halves is called a **half-gateway** and each one is owned and operated by one of the network operators. The whole problem of gatewaying then reduces to agreeing to a common protocol to use on the wire, one that is neutral and does not favor either party. Figure 5-34 shows both full and half-gateways. Either kind can be used in any layer (e.g., half-bridges also exist).

That all said, the situation is murkier in practice than it is in theory. Many devices on the market combine bridge and router functionality. The key property of a pure bridge is that it examines data link layer frame headers and does not inspect or modify the network layer packets inside the frames. A bridge cannot




Fig. 5-34. (a) A full gateway between two WANs. (b) A full gateway between a LAN and a WAN. (c) Two half-gateways.

tell and does not care whether the frame it is forwarding from an 802.x LAN to an 802.y contains an IP, IPX, or CLNP packet in the payload field.

A router, in contrast, knows very well whether it is an IP router, an IPX router, a CLNP router, or all three combined. It examines these headers and makes decisions based on the addresses found there. On the other hand, when a pure router hands off a packet to the data link layer, it does not know or care whether it will be carried in an Ethernet frame or a token ring frame. That is the data link layer's responsibility.

The confusion in the industry comes from two sources. First, functionally, bridges and routers are not all that different. They each accept incoming PDUs (Protocol Data Units), examine some header fields, and make decisions about where to send the PDUs based on header information and internal tables.

Second, many commercial products are sold under the wrong label or combine the functionality of both bridges and routers. For example, source routing bridges are not really bridges at all, since they involve a protocol layer above the data link layer to do their job. For an illuminating discussion of bridges versus routers, see Chap. 12 of (Perlman, 1992).

5.4.1. How Networks Differ

Networks can differ in many ways. In Fig. 5-35 we list some of the differences that can occur in the network layer. It is papering over these differences that make internetworking more difficult than operating within a single network.

Item	Some Possibilities
Service offered	Connection-oriented versus connectionless
Protocols	IP, IPX, CLNP, AppleTalk, DECnet, etc.
Addressing	Flat (802) versus hierarchical (IP)
Multicasting	Present or absent (also broadcasting)
Packet size	Every network has its own maximum
Quality of service	May be present or absent; many different kinds
Error handling	Reliable, ordered, and unordered delivery
Flow control	Sliding window, rate control, other, or none
Congestion control	Leaky bucket, choke packets, etc.
Security	Privacy rules, encryption; etc.
Parameters	Different timeouts, flow specifications, etc.
Accounting	By connect time, by packet, by byte, or not at all

Fig. 5-35. Some of the many ways networks can differ.

When packets sent by a source on one network must transit one or more foreign networks before reaching the destination network (which also may be different from the source network), many problems can occur at the interfaces between networks. To start with, when packets from a connection-oriented network must transit a connectionless one, they may be reordered, something the sender does not expect and the receiver is not prepared to deal with. Protocol conversions will often be needed, which can be difficult if the required functionality cannot be expressed. Address conversions will also be needed, which may require some kind of directory system. Passing multicast packets through a network that does not support multicasting requires generating separate packets for each destination.

The differing maximum packet sizes used by different networks is a major headache. How do you pass an 8000-byte packet through a network whose maximum size is 1500 bytes? Differing qualities of service is an issue when a packet that has real-time delivery constraints passes through a network that does offer any real-time guarantees.

Error, flow, and congestion control frequently differ among different networks. If the source and destination both expect all packets to be delivered in sequence without error, yet an intermediate network just discards packets whenever it smells congestion on the horizon, or packets can wander around aimlessly for a while and then suddenly emerge and be delivered, many applications will break. Different security mechanisms, parameter settings, and accounting rules, and even national privacy laws also can cause problems.

5.4.2. Concatenated Virtual Circuits

Two styles of internetworking are common: a connection-oriented concatenation of virtual circuit subnets, and a datagram internet style. We will now examine these in turn. In the concatenated virtual circuit model, shown in Fig. 5-36, a connection to a host in a distant network is set up in a way similar to the way connections are normally established. The subnet sees that the destination is remote and builds a virtual circuit to the router nearest the destination network. Then it constructs a virtual circuit from that router to an external "gateway" (multiprotocol router). This gateway records the existence of the virtual circuit in its tables and proceeds to build another virtual circuit to a router in the next subnet. This process continues until the destination host has been reached.




Fig. 5-36. Internetworking using concatenated virtual circuits.

Once data packets begin flowing along the path, each gateway relays incoming packets, converting between packet formats and virtual circuit numbers as needed. Clearly, all data packets must traverse the same sequence of gateways, and thus arrive in order.

The essential feature of this approach is that a sequence of virtual circuits is set up from the source through one or more gateways to the destination. Each gateway maintains tables telling which virtual circuits pass through it, where they are to be routed, and what the new virtual circuit number is.

Although Fig. 5-36 shows the connection made with a full gateway, it could equally well be done with half-gateways.

This scheme works best when all the networks have roughly the same properties. For example, if all of them guarantee reliable delivery of network layer packets, then barring a crash somewhere along the route, the flow from source to destination will also be reliable. Similarly, if none of them guarantee reliable delivery, then the concatenation of the virtual circuits is not reliable either. On

the other hand, if the source machine is on a network that does guarantee reliable delivery, but one of the intermediate networks can lose packets, the concatenation has fundamentally changed the nature of the service.

Concatenated virtual circuits are also common in the transport layer. In particular, it is possible to build a bit pipe using, say, OSI, which terminates in a gateway, and have a TCP connection go from the gateway to the next gateway. In this manner, an end-to-end virtual circuit can be built spanning different networks and protocols.

5.4.3. Connectionless Internetworking

The alternative internetwork model is the datagram model, shown in Fig. 5-37. In this model, the only service the network layer offers to the transport layer is the ability to inject datagrams into the subnet and hope for the best. There is no notion of a virtual circuit at all in the network layer, let alone a concatenation of them. This model does not require all packets belonging to one connection to traverse the same sequence of gateways. In Fig. 5-37 datagrams from host 1 to host 2 are shown taking different routes through the internetwork. A routing decision is made separately for each packet, possibly depending on the traffic at the moment the packet is sent. This strategy can use multiple routes and thus achieve a higher bandwidth than the concatenated virtual circuit model. On the other hand, there is no guarantee that the packets arrive at the destination in order, assuming that they arrive at all.




Fig. 5-37. A connectionless internet.

The model of Fig. 5-37 is not quite as simple as it looks. For one thing, if each network has its own network layer protocol, it is not possible for a packet from one network to transit another one. One could imagine the multiprotocol routers actually trying to translate from one format to another, but unless the two

formats are close relatives with the same information fields, such conversions will always be incomplete and often doomed to failure. For this reason, conversion is rarely attempted.

A second, and more serious problem, is addressing. Imagine a simple case: a host on the Internet is trying to send an IP packet to a host on an adjoining OSI host. The OSI datagram protocol, CLNP, was based on IP and is close enough to it that a conversion might well work. The trouble is that IP packets all carry the 32-bit Internet address of the destination host in a header field. OSI hosts do not have 32-bit Internet addresses. They use decimal addresses similar to telephone numbers.

To make it possible for the multiprotocol router to convert between formats, someone would have to assign a 32-bit Internet address to each OSI host. Taken to the limit, this approach would mean assigning an Internet address to every machine in the world that an Internet host might want to talk to. It would also mean assigning an OSI address to every machine in the world that an OSI host might want to talk to. The same problem occurs with every other address space (SNA, AppleTalk, etc.). The problems here are insurmountable. In addition, someone would have to maintain a database mapping everything to everything.

Another idea is to design a universal "internet" packet and have all routers recognize it. This approach is, in fact, what IP is—a packet designed to be carried through many networks. The only problem is that IPX, CLNP, and other "universal" packets exist too, making all of them less than universal. Getting everybody to agree to a single format is just not possible.

Let us now briefly recap the two ways internetworking can be attacked. The concatenated virtual circuit model has essentially the same advantages as using virtual circuits within a single subnet: buffers can be reserved in advance, sequencing can be guaranteed, short headers can be used, and the troubles caused by delayed duplicate packets can be avoided.

It also has the same disadvantages: table space required in the routers for each open connection, no alternate routing to avoid congested areas, and vulnerability to router failures along the path. It also has the disadvantage of being difficult, if not impossible, to implement if one of the networks involved is an unreliable datagram network.

The properties of the datagram approach to internetworking are the same as those of datagram subnets: more potential for congestion, but also more potential for adapting to it, robustness in the face of router failures, and longer headers needed. Various adaptive routing algorithms are possible in an internet, just as they are within a single datagram network.

A major advantage of the datagram approach to internetworking is that it can be used over subnets that do not use virtual circuits inside. Many LANs, mobile networks (e.g., aircraft and naval fleets), and even some WANs fall into this category. When an internet includes one of these, serious problems occur if the internetworking strategy is based on virtual circuits.

5.4.4. Tunneling

Handling the general case of making two different networks interwork is exceedingly difficult. However, there is a common special case that is manageable. This case is where the source and destination hosts are on the same type of network, but there is a different network in between. As an example, think of an international bank with a TCP/IP based Ethernet in Paris, a TCP/IP based Ethernet in London, and a PTT WAN in between, as shown in Fig. 5-38.




Fig. 5-38. Tunneling a packet from Paris to London.

The solution to this problem is a technique called **tunneling**. To send an IP packet to host 2, host 1 constructs the packet containing the IP address of host 2, inserts it into an Ethernet frame addressed to the Paris multiprotocol router, and puts it on the Ethernet. When the multiprotocol router gets the frame, it removes the IP packet, inserts it in the payload field of the WAN network layer packet, and addresses the latter to the WAN address of the London multiprotocol router. When it gets there, the London router removes the IP packet and sends it to host 2 inside an Ethernet frame.

The WAN can be seen as a big tunnel extending from one multiprotocol router to the other. The IP packet just travels from one end of the tunnel to the other, snug in its nice box. It does not have to worry about dealing with the WAN at all. Neither do the hosts on either Ethernet. Only the multiprotocol router has to understand IP and WAN packets. In effect, the entire distance from the middle of one multiprotocol router to the middle of the other acts like a serial line.

An analogy may make tunneling clearer. Consider a person driving her car from Paris to London. Within France, the car moves under its own power, but when it hits the English Channel, it is loaded into a high-speed train and transported to England through the Chunnel (cars are not permitted to drive through the Chunnel). Effectively, the car is being carried as freight, as depicted in Fig. 5-39.

At the far end, the car is let loose on the English roads and once again continues to move under its own power. Tunneling of packets through a foreign network works the same way.




Fig. 5-39. Tunneling a car from France to England.

5.4.5. Internetwork Routing

Routing through an internetwork is similar to routing within a single subnet, but with some added complications. Consider, for example, the internetwork of Fig. 5-40(a) in which five networks are connected by six multiprotocol routers. Making a graph model of this situation is complicated by the fact that every multiprotocol router can directly access (i.e., send packets to) every other router connected to any network to which it is connected. For example, B in Fig. 5-40(a) can directly access A and C via network 2 and also D via network 3. This leads to the graph of Fig. 5-40(b).

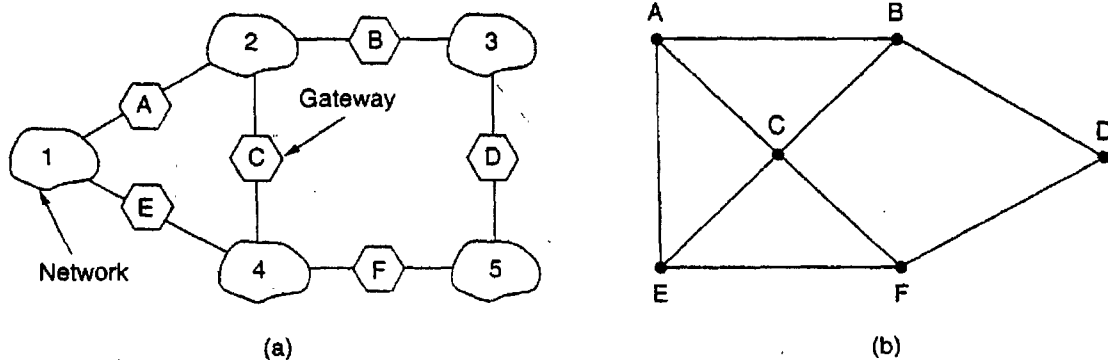


Fig. 5-40. (a) An internetwork. (b) A graph of the internetwork.

Once the graph has been constructed, known routing algorithms, such as the distance vector and link state algorithms, can be applied to the set of multiprotocol routers. This gives a two-level routing algorithm: within each network an interior gateway protocol is used, but between the networks, an exterior gateway protocol is used ("gateway" is an older term for "router"). In fact, since

each network is independent, they may all use different algorithms. Because each network in an internetwork is independent of all the others, it is often referred to as an **Autonomous System (AS)**.

A typical internet packet starts out on its LAN addressed to the local multiprotocol router (in the MAC layer header). After it gets there, the network layer code decides which multiprotocol router to forward the packet to, using its own routing tables. If that router can be reached using the packet's native network protocol, it is forwarded there directly. Otherwise it is tunneled there, encapsulated in the protocol required by the intervening network. This process is repeated until the packet reaches the destination network.

One of the differences between internetwork routing and intranetwork routing is that internetwork routing often requires crossing international boundaries. Various laws suddenly come into play, such as Sweden's strict privacy laws about exporting personal data about Swedish citizens from Sweden. Another example is the Canadian law saying that data traffic originating in Canada and ending in Canada may not leave the country. This law means that traffic from Windsor, Ontario to Vancouver may not be routed via nearby Detroit.

Another difference between interior and exterior routing is the cost. Within a single network, a single charging algorithm normally applies. However, different networks may be under different managements, and one route may be less expensive than another. Similarly, the quality of service offered by different networks may be different, and this may be a reason to choose one route over another.

In a large internetwork, choosing the best route may be a time-consuming operation. Estrin et al. (1992) have proposed dealing with this problem by precomputing routes for popular (source, destination) pairs and storing them in a database to be consulted at route selection time.

5.4.6. Fragmentation

Each network imposes some maximum size on its packets. These limits have various causes, among them:

1. Hardware (e.g., the width of a TDM transmission slot).
2. Operating system (e.g., all buffers are 512 bytes).
3. Protocols (e.g., the number of bits in the packet length field).
4. Compliance with some (inter)national standard.
5. Desire to reduce error induced retransmissions to some level.
6. Desire to prevent one packet from occupying the channel too long.

The result of all these factors is that the network designers are not free to choose any maximum packet size they wish. Maximum payloads range from 48 bytes

**BUSINESS
COMMUNICATIONS
REVIEW**

**BUSINESS
COMMUNICATIONS
REVIEW**

Save 25% Off Cover Price. Get 12 Monthly Issues for \$49

Current Articles
Older Articles
Index of Contents

SUBSCRIBE

Subscriber Services
(Change of Address, etc.)

BCR eWeekly:

Free weekly
newsletter on
networking & telecom
trends. No hype. Just
information.

your email

Subscribe

Read it in the BCR
eForum

Search In:

BCR Magazine

Go

Advanced Search
BCR Acronym Guide

Supplements:

BCR ACCESS
Voice 2000

JULY 1999 ISSUE:

- Frame Relay and IP
VPNs: Compete or
Coexist?

IN BRIEF:

- Inside the CLEC/ILEC
Service Order Process
- VPNs: At Least A
Remote Possibility
- Multiservice Networks:
Converging on IP
VPNs
- Will Carrier Equipment
Vendors Strike Gold?
- Token Ring For
Gamblers
- Building and Retaining
Staff—It's a Whole
New World
- 12k Telephone Calls:
Prepare for the Tidal
Wave

OPINION:

- Customers Need
Performance-Based
Network ROI Analysis
- The New Customer

CARRIER SERVICES

**Frame Relay And IP VPNs: Compete Or
Coexist?**

from the July 1999 issue of Business Communications Review, pp. 28-32

by Joanie Wexler, an independent networking editor/writer in Campbell, CA.

As an access service, frame relay continues to flourish. In the core of many public data networks, though, alternatives to pure frame switching are being put in place to support the trend toward IP virtual private networks (VPNs). Because of these changes, enterprise users may soon find themselves buying hybrid services that look and feel like frame relay at the customer premises site, but offer broader, less expensive connectivity across the carrier backbone—albeit with potentially inferior service guarantees.

Frame relay services grew by a healthy 46 percent from 1998 to 1999, according to Distributed Networking Associates, Inc., a network consulting firm in Greensboro, NC, that surveys the market each year on behalf of the Frame Relay Forum and a consortium of equipment makers and service providers. Users have long been satisfied with the cost benefits and networking stability frame relay provides, particularly for LAN-to-LAN traffic, and the market shows no signs of slowing down.

In addition, new capabilities continue to be specified. An Implementation Agreement (IA) currently under development for MultiLink Frame Relay, for example, will soon help bridge the bandwidth gaps between 56/64 kbps and T1 access speeds and between T1 and T3 access speeds. The IA is intended to bring a standard for inverse multiplexing implementations to the frame relay industry.

However, in carrier backbones, frame switching is starting to give way to other technologies. Many carriers already convert frames to ATM cells for more efficient transport over higher-speed ATM networks, though this does not change the look and feel of the frame relay service. Now, Layer 3 switching, based on the emerging Multiprotocol Label Switching (MPLS) standard, is also beginning to crop up in network cores to improve the scalability of both frame relay- and ATM-based IP networks. The result, depending on a particular carrier's marketing plans, might be an alternative to frame relay or, as is the case with a service being turned up by AT&T this month, a hybrid that offers frame relay access to an IP VPN backbone.

MPLS is a leading contender as a base platform for IP VPNs, which carry a slightly different set of components than frame relay service. For example, as with frame relay, users can specify the access and port rate at which traffic is pumped into and out of the WAN cloud. But the idea of a committed information rate (CIR) disappears for IP VPNs, because it is very difficult for a carrier to guarantee a specific rate across a network cloud when no permanent virtual circuits (PVCs) exist in that portion of the network.

Mock-Mesh Networking

So why go to all this trouble? Why not simply continue to run IP encapsulated in frame relay packets?

The main reason to consider an IP VPN over "vanilla" frame-relay service is to

Hierarchy

- Meanwhile, Back at the Ranch...

REVIEW:

- NBase-Xyplex's Linux Router

TECHnary

the animated technical dictionary

BCR White Papers

Search for White Papers and Analyst Research:

Search Help
Advanced Search

Advertising Information
BCR Edit Calendar 2003
BCR Edit Calendar 2002

About BCR Magazine
Email the Editor

achieve broader connectivity at a lower cost and with fewer management headaches. A growing number of enterprises are seeking to move away from star network topologies, and have direct "any-to-any" connectivity among their frame relay sites. However, they would like to achieve this connectivity without having to purchase a full n -squared mesh of PVCs and manage the associated data link connection identifiers (DLCIs). One way carriers are looking to deliver these capabilities is via connections, MPLS-based IP VPN architectures that are more scalable than criss-crosses of frame relay PVCs.

"A single DLCI can be used to represent the entire user group; one DLCI can communicate with multiple sites," explained Bill Flanagan, program director at NetReference, Inc., a network architecture consulting firm in Sterling, VA.

This cuts down on administrative work. "In an MPLS network, 'soft PVCs' are created by routing updates, rather than setting DLCIs by hand," said Fred Sammartino, director of IP marketing at Ascend Communications, which is being acquired by Lucent Technologies. Ascend equipment powers the lion's share of public frame relay networks, and its IP Navigator switches enable MPLS networking based on either ATM or frame relay technology.

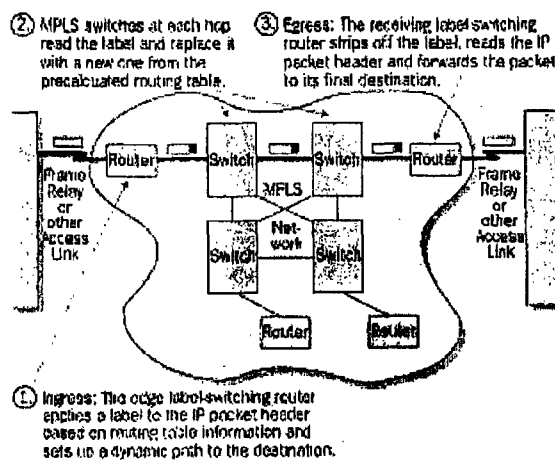
In addition, IP VPNs are a boon to mobile networkers, who have found it challenging to access the corporate frame relay network. With an IP VPN, they simply use their TCP/IP software to access the network, rather than having to run special frame relay software, Sammartino notes.

Finally, in an MPLS-enabled VPN, the Border Gateway Protocol (BGP) wide-area routing algorithm distributes information about VPNs only to members of the same VPN. This traffic separation provides base-level, native security.

The first version of the MPLS standard—which addresses packet forwarding—is still being finalized by the IETF. However, prestandard versions are available from Ascend, as well as from Cisco Systems, which pioneered the concept of MPLS with its Tag Switching technology several years ago. Switching gear from Nortel Networks will support prestandard MPLS by September, according to the company.

MPLS blends the intelligence of routing—the ability to choose the optimal path through the network to a destination—with the speed of switching. In an MPLS network, a routing decision is made based on the IP destination address. MPLS maps IP addressing and routing information directly into ATM or frame-relay switching tables. Traffic is then switched across the backbone with no more IP address lookups or processing required in the backbone at Layer 3 (Figure 1).

FIGURE 1. The Basics of an MPLS Network



In traditional frame relay using end-to-end PVCs, on the other hand, traffic tends to be switched between two customer sites across a fixed Layer 2 path across the network backbone and access links. This approach works well, but, from a performance point of view, does not take into account congestion or failed interim switches to optimize the delivery path.

Though the use of MPLS is far from widespread, many carriers' provisioning plans call for MPLS to blend IP with ATM (rather than frame) in the backbone, so as to piggyback on ATM's strong class-of-service (COS) capabilities. It is possible to use MPLS with frame relay switches as well, a strategy likely to be more common in local-exchange provider networks, which tend to use pure frame switching in their backbones more often than long-haul networks do.

AT&T's IPFR: Frame Relay or IP VPN?

Sometimes it seems that frame relay and IP VPN services are becoming indistinguishable, depending on the marketing approach taken by the provisioning carrier. Regardless of the backbone technology used to provision IP VPNs, customers can still use frame relay to access the cloud.

Consider, for example, AT&T's IP-Enabled Frame Relay (IPFR) service, announced in January and scheduled for general availability this month. The service represents frame-relay access to an IP VPN service based on a prestandard version of MPLS running in Cisco's MGX 8800 ATM switches.


"I believe AT&T is using frame relay access because their target customers already have frame relay and thus can simply allocate a PVC on an existing interface to the VPN access mission," said Tom Nolle, president of CIMI Corp., a network consulting firm in Voorhees, NJ. "However, any other interface, such as ATM or a leased line, would also work." Others point out, though, that at T1-and-below speeds, ATM's overhead often makes it an uneconomical access alternative.

"IPFR enables existing frame relay customers to set up IPFR PVCs across the same access link as traditional frame relay PVCs," said Keith Falter, AT&T frame relay marketing manager. The port and access link components of IPFR remain the same as with traditional frame relay service.

"The benefit of the IPFR service is that it reduces complexity for users seeking a high degree of connectivity inside the network cloud, breaking down the performance barriers of hub-and-spoke configurations," said Falter.

By enabling sites to cost-effectively communicate directly instead of going through a hub site first, users can realize some performance improvements, although exactly how much will depend on the application, number of sites and IP parameters. A 25-site IPFR network is likely to see a 25-30 percent performance improvement by eliminating the latency introduced by a stop-off at the hub site, said Tim Halpin, AT&T product manager, frame relay and ATM services.

AT&T claims that the IPFR service also can sharply reduce the number of PVCs required by a customer to achieve full mesh connectivity. This is because only one PVC is required to run from the customer premises site to the WAN cloud in order for



communicate with any other site within the user group (Figure 2). By contrast, in a traditional frame relay network, users would have to purchase PVCs among all sites to achieve the same level of meshed connectivity—or, where available, they'd have to purchase switched virtual circuit (SVC) services.

In a 25-site IPFR network, 25 PVCs are required—one from each site to the

WAN cloud, according to Halpin. With IPFR, two priority classes are available, so having both a high- and low-priority PVC at each location brings the PVC count to 50. The same configuration using meshed, direct, frame-relay connections would require 300 PVCs for a single-priority network and 600 PVCs for a two-priority network.

QOS Considerations

From a QOS perspective, however, IPFR has a drawback: There is no PVC or CIR across the backbone. Instead, the service uses a concept called *committed delivery rate* (CDR), meaning that each site receives a minimum, committed amount of traffic as an aggregate from all sites connected to it. There is no commitment of bandwidth between any particular two sites and there are currently no SLAs available for this service, Halpin acknowledged.

"Since PVCs with CIRs don't exist [in IPFR], CIRs aren't available for traffic engineering," said Halpin. "Therefore, a new traffic engineering variable called CDR defines the minimum amount of IP traffic that can be guaranteed for delivery across the network."

To determine the proper CDR at a given site (which is equal to or less than the speed of the access link), users must analyze their IP traffic flows and configure their links accordingly, Halpin explained. Customers choose the size of the CDR (as they once chose CIR) based on the volume of incoming traffic; they sign up with the carrier for that amount and tune their equipment to support it.

For now, the lack of an SLA could prove to be a deterrent to such services. "When you throw in IP, SLAs get much more complex," said Liza Henderson, director of consulting at TeleChoice, Inc., a telecommunications consultancy headquartered in Boston. "There is not enough technology in place today to guarantee complete SLAs with IP VPNs."

Eventually, MPLS will bring higher levels of QOS to frame relay networks by specifying priority information in the MPLS label that is attached to each packet. The priority information will be derived from IP Precedence information that is contained in the IP packet header. However, QOS capabilities are not being addressed in version 1 of the MPLS standard.

Pre-standard MPLS QOS capabilities will be available soon, however. Cisco's IOS switch/ router software will support MPLS QOS this summer, according to the company, and Ascend and Nortel Networks said they expect their switching equipment to support MPLS QOS capabilities by the end of the year. "MPLS, because it brings circuit-mode properties to an IP network, [will be] as QOS-capable as a virtual circuit," said analyst Tom Nolle.

FRF.14: A Boon to Frame Relay in the Carrier Backbone?

Other developments in frame relay may affect IP VPNs indirectly, as well as changing the landscape for native frame-relay services. For example, while most MPLS work is happening in conjunction with ATM backbones, the recent ratification of FRF.14, the Forum's Physical Layer IA, could rejuvenate frame switching as an MPLS adjunct.

FRF.14 specifies a common way to support various physical interface types in frame relay switches, and it includes a specification for SONET/SDH. This will enable frame relay to scale to higher speeds—OC-3 (155 Mbps) and OC-12 (622 Mbps)—than today's limit of T3 (45 Mbps). This could pose some fresh competition to frame-to-ATM interworking services, as well as render frame switches a more viable platform for MPLS software.

Three key infrastructure providers—Ascend, Cisco and Nortel Networks—have indicated plans to provide the capabilities in their switches to enable such high-

speed frame relay services, which would extend the scalability of carriers' frame relay networks.

MPLS versus SVCs

As mentioned, Layer 3 switching enabled by pre-standard MPLS implementations in carrier switches enables dynamic connectivity among all sites within a user group running connectionless IP traffic. No special pre-provisioning is necessary. This renders it an attractive alternative to frame relay switched virtual circuits (SVCs), which, while scarce in the marketplace, are available today.

SVCs have the advantage of switching any Layer 3 protocol. Technically, MPLS could do the same, however to date, it is only being developed for blending IP with Layer 2 switched networks. In addition, SVCs retain CIRs, so customers know exactly how much minimum bandwidth, end to end, will be available for their traffic.

SVC services are available from MCI WorldCom and Qwest Communications International. As voice and other interactive applications join the network, either SVCs or MPLS will be needed to deliver the dynamic and widespread connectivity required for collaborative sessions.

Whether it makes sense for a user to purchase these services depends on the carrier's pricing model and associated fees, considered along with the enterprise's connectivity needs. For example, Qwest's SVC service is based solely on usage; a customer who does not transmit any SVC traffic in a given month pays nothing. In contrast, MCI WorldCom's SVC services, as of press time, were still flat-rated. Enterprise customers must do the calculations, based on their traffic patterns, to determine if these services match their any-to-any needs.

Frame relay SVCs are not available from AT&T because "SVCs entail setup time and tend to be usage-based [in pricing], which is not attractive to some budget-conscious users," Halpin said. The carrier, however, has long offered SVCs for ATM services.

Frame Relay SLAs Get Stronger

While the industry awaits stronger QOS capabilities for IP VPNs, the carriers have begun bringing more comprehensive QOS offerings to their frame relay services. Frame relay has long been a successful Layer 2 VPN that maintained the privacy of communications among sites within a user group. It also has always been capable of delivering a degree of QOS via the user-specified CIR.

Not until recently, however, have the latency and packet-loss components of the QOS package been available to frame-relay customers. Part of the reason for the holdup is that, unlike ATM classes of service and MPLS, "there is no standard QOS for frame relay," noted Henderson of TeleChoice. "The various priority-queuing techniques in use are vendor-specific implementations."

Delay and throughput metrics become important as low-latency applications such as voice and video join the traffic mix, because network congestion could degrade the quality of a session if traffic is not delivered within a set of strict parameters. Advances in switching equipment that make it possible to prioritize frame relay PVCs and to map frame relay PVCs to ATM COS PVCs, enable most carriers to now offer SLAs (Table 1). In addition, Sprint and Qwest each offer three classes of frame-relay service with different SLAs, based on application. The highest class is recommended for voice over frame relay, the middle for LAN and SNA traffic, and the lowest for Internet access or other lower-priority applications.

TABLE 1 Standard Frame Relay Guarantees

	Availability	Maximum Delay	Data Delivery Ratio
AT&T	99.99%	60 msec (one way)	99.99% at or below CIR
MCI	99.99% (core only)	60 msec (one way)	99.99% at or below CIR
Worldcom	99.9% (with access link included)		
Sprint	100% if Sprint provides access link; 99.9% if alternative access is used	55-130 msec one way, end to end, depending on CIR and service class	99% (with no CIR); 99.9% (with CIR specified)

Sources: AT&T, MCI Worldcom, Sprint

Conclusion

Frame relay continues to hold its own in popularity, and is continually getting important new enhancements, such as higher-speed physical interface standards and the emerging multilink capabilities. As customers expand their connectivity and grow their use of the IP protocol, however, IP VPNs based on MPLS in the backbone will emerge as an alternative to frame-relay SVCs. IP VPNs will enable this scalability at affordable costs to customers and carriers with a minimum of operational overhead.

As time marches on, frame relay will not only serve as an access technology to frame- and ATM-switched networks, but, depending on carrier implementation, also to MPLS-switched networks based on ATM or frame-relay technology and optimized for IP traffic. Ultimately, users must factor in the degree of connectivity they require, pricing models, SLA availability and the Layer 3 protocols they run to determine which service—or blend of services—makes sense in their specific environments.

Top of Page

[BCR Home](#) | [Site Map](#) | [Contact Us](#) | [Search](#) | [BCR Magazine](#)
[NGN](#) | [NGN Ventures](#) | [Opticon](#) | [VoiceCon](#)
[Instructor-Led Training/Seminars](#)
[Subscribe to BCR Magazine](#) | [Register for Conference](#) | [Register for Seminar](#)

All contents of this site copyright © 1995-2002 Key3Media BCR Events, Inc. All Rights Reserved

Please direct any comments or questions to: webmaster@bcr.com.

A4

Network Working Group
Request for Comments: 2764
Category: Informational

B. Gleason
A. Lin
Nortel Networks
J. Heinanen
Telia Finland
G. Armitage
A. Malis
Lucent Technologies
February 2000

A Framework for IP Based Virtual Private Networks

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

IESG Note

This document is not the product of an IETF Working Group. The IETF currently has no effort underway to standardize a specific VPN framework.

Abstract

This document describes a framework for Virtual Private Networks (VPNs) running across IP backbones. It discusses the various different types of VPNs, their respective requirements, and proposes specific mechanisms that could be used to implement each type of VPN using existing or proposed specifications. The objective of this document is to serve as a framework for related protocol development in order to develop the full set of specifications required for widespread deployment of interoperable VPN solutions.

Table of Contents

1.0 Introduction	4
2.0 VPN Application and Implementation Requirements	5
2.1 General VPN Requirements	5
2.1.1 Opaque Packet Transport:	6
2.1.2 Data Security	7
2.1.3 Quality of Service Guarantees	7
2.1.4 Tunneling Mechanism	8
2.2 CPE and Network Based VPNs	8
2.3 VPNs and Extranets	9
3.0 VPN Tunneling	10
3.1 Tunneling Protocol Requirements for VPNs	11
3.1.1 Multiplexing	11
3.1.2 Signalling Protocol	12
3.1.3 Data Security	13
3.1.4 Multiprotocol Transport	14
3.1.5 Frame Sequencing	14
3.1.6 Tunnel Maintenance	15
3.1.7 Large MTUs	16
3.1.8 Minimization of Tunnel Overhead	16
3.1.9 Flow and congestion control	17
3.1.10 QoS / Traffic Management	17
3.2 Recommendations	18
4.0 VPN Types: Virtual Leased Lines	18
5.0 VPN Types: Virtual Private Routed Networks	20
5.1 VPRN Characteristics	20
5.1.1 Topology	23
5.1.2 Addressing	24
5.1.3 Forwarding	24
5.1.4 Multiple concurrent VPRN connectivity	24
5.2 VPRN Related Work	24
5.3 VPRN Generic Requirements	25
5.3.1 VPN Identifier	26
5.3.2 VPN Membership Information Configuration	27
5.3.2.1 Directory Lookup	27
5.3.2.2 Explicit Management Configuration	28
5.3.2.3 Piggybacking in Routing Protocols	28
5.3.3 Stub Link Reachability Information	30
5.3.3.1 Stub Link Connectivity Scenarios	30
5.3.3.1.1 Dual VPRN and Internet Connectivity	30
5.3.3.1.2 VPRN Connectivity Only	30
5.3.3.1.3 Multihomed Connectivity	31
5.3.3.1.4 Backdoor Links	31
5.3.3.1 Routing Protocol Instance	31
5.3.3.2 Configuration	33
5.3.3.3 ISP Administered Addresses	33
5.3.3.4 MPLS Label Distribution Protocol	33

5.3.4 Intra-VPN Reachability Information	34
5.3.4.1 Directory Lookup	34
5.3.4.2 Explicit Configuration	34
5.3.4.3 Local Intra-VPN Routing Instantiations	34
5.3.4.4 Link Reachability Protocol	35
5.3.4.5 Piggybacking in IP Backbone Routing Protocols	36
5.3.5 Tunneling Mechanisms	36
5.4 Multihomed Stub Routers	37
5.5 Multicast Support	38
5.5.1 Edge Replication	38
5.5.2 Native Multicast Support	39
5.6 Recommendations	40
6.0 VPN Types: Virtual Private Dial Networks	41
6.1 L2TP protocol characteristics	41
6.1.1 Multiplexing	41
6.1.2 Signalling	42
6.1.3 Data Security	42
6.1.4 Multiprotocol Transport	42
6.1.5 Sequencing	42
6.1.6 Tunnel Maintenance	43
6.1.7 Large MTUs	43
6.1.8 Tunnel Overhead	43
6.1.9 Flow and Congestion Control	43
6.1.10 QoS / Traffic Management	43
6.1.11 Miscellaneous	44
6.2 Compulsory Tunneling	44
6.3 Voluntary Tunnels	46
6.3.1 Issues with Use of L2TP for Voluntary Tunnels	46
6.3.2 Issues with Use of IPSec for Voluntary Tunnels	48
6.4 Networked Host Support	49
6.4.1 Extension of PPP to Hosts Through L2TP	49
6.4.2 Extension of PPP Directly to Hosts:	49
6.4.3 Use of IPSec	50
6.5 Recommendations	50
7.0 VPN Types: Virtual Private LAN Segment	50
7.1 VPLS Requirements	51
7.1.1 Tunneling Protocols	51
7.1.2 Multicast and Broadcast Support	52
7.1.3 VPLS Membership Configuration and Topology	52
7.1.4 CPE Stub Node Types	52
7.1.5 Stub Link Packet Encapsulation	53
7.1.5.1 Bridge CPE	53
7.1.5.2 Router CPE	53
7.1.6 CPE Addressing and Address Resolution	53
7.1.6.1 Bridge CPE	53
7.1.6.2 Router CPE	54
7.1.7 VPLS Edge Node Forwarding and Reachability Mechanisms	54
7.1.7.1 Bridge CPE	54

7.1.7.2 Router CPE	54
7.2 Recommendations	55
8.0 Summary of Recommendations	55
9.0 Security Considerations	56
10.0 Acknowledgements	56
11.0 References	56
12.0 Author Information	61
13.0 Full Copyright Statement	62

1.0 Introduction

This document describes a framework for Virtual Private Networks (VPNs) running across IP backbones. It discusses the various different types of VPNs, their respective requirements, and proposes specific mechanisms that could be used to implement each type of VPN using existing or proposed specifications. The objective of this document is to serve as a framework for related protocol development in order to develop the full set of specifications required for widespread deployment of interoperable VPN solutions.

There is currently significant interest in the deployment of virtual private networks across IP backbone facilities. The widespread deployment of VPNs has been hampered, however, by the lack of interoperable implementations, which, in turn, derives from the lack of general agreement on the definition and scope of VPNs and confusion over the wide variety of solutions that are all described by the term VPN. In the context of this document, a VPN is simply defined as the 'emulation of a private Wide Area Network (WAN) facility using IP facilities' (including the public Internet, or private IP backbones). As such, there are as many types of VPNs as there are types of WANs, hence the confusion over what exactly constitutes a VPN.

In this document a VPN is modeled as a connectivity object. Hosts may be attached to a VPN, and VPNs may be interconnected together, in the same manner as hosts today attach to physical networks, and physical networks are interconnected together (e.g., via bridges or routers). Many aspects of networking, such as addressing, forwarding mechanism, learning and advertising reachability, quality of service (QoS), security, and firewalling, have common solutions across both physical and virtual networks, and many issues that arise in the discussion of VPNs have direct analogues with those issues as implemented in physical networks. The introduction of VPNs does not create the need to reinvent networking, or to introduce entirely new paradigms that have no direct analogue with existing physical networks. Instead it is often useful to first examine how a particular issue is handled in a physical network environment, and then apply the same principle to an environment which contains

virtual as well as physical networks, and to develop appropriate extensions and enhancements when necessary. Clearly having mechanisms that are common across both physical and virtual networks facilitates the introduction of VPNs into existing networks, and also reduces the effort needed for both standards and product development, since existing solutions can be leveraged.

This framework document proposes a taxonomy of a specific set of VPN types, showing the specific applications of each, their specific requirements, and the specific types of mechanisms that may be most appropriate for their implementation. The intent of this document is to serve as a framework to guide a coherent discussion of the specific modifications that may be needed to existing IP mechanisms in order to develop a full range of interoperable VPN solutions.

The document first discusses the likely expectations customers have of any type of VPN, and the implications of these for the ways in which VPNs can be implemented. It also discusses the distinctions between Customer Premises Equipment (CPE) based solutions, and network based solutions. Thereafter it presents a taxonomy of the various VPN types and their respective requirements. It also outlines suggested approaches to their implementation, hence also pointing to areas for future standardization.

Note also that this document only discusses implementations of VPNs across IP backbones, be they private IP networks, or the public Internet. The models and mechanisms described here are intended to apply to both IPV4 and IPV6 backbones. This document specifically does not discuss means of constructing VPNs using native mappings onto switched backbones - e.g., VPNs constructed using the LAN Emulation over ATM (LANE) [1] or Multiprotocol over ATM (MPOA) [2] protocols operating over ATM backbones. Where IP backbones are constructed using such protocols, by interconnecting routers over the switched backbone, the VPNs discussed operate on top of this IP network, and hence do not directly utilize the native mechanisms of the underlying backbone. Native VPNs are restricted to the scope of the underlying backbone, whereas IP based VPNs can extend to the extent of IP reachability. Native VPN protocols are clearly outside the scope of the IETF, and may be tackled by such bodies as the ATM Forum.

2.0 VPN Application and Implementation Requirements

2.1 General VPN Requirements

There is growing interest in the use of IP VPNs as a more cost effective means of building and deploying private communication networks for multi-site communication than with existing approaches.

Existing private networks can be generally categorized into two types - dedicated WANs that permanently connect together multiple sites, and dial networks, that allow on-demand connections through the Public Switched Telephone Network (PSTN) to one or more sites in the private network.

WANs are typically implemented using leased lines or dedicated circuits - for instance, Frame Relay or ATM connections - between the multiple sites. CPE routers or switches at the various sites connect these dedicated facilities together and allow for connectivity across the network. Given the cost and complexity of such dedicated facilities and the complexity of CPE device configuration, such networks are generally not fully meshed, but instead have some form of hierarchical topology. For example remote offices could be connected directly to the nearest regional office, with the regional offices connected together in some form of full or partial mesh:

Private dial networks are used to allow remote users to connect into an enterprise network using PSTN or Integrated Services Digital Network (ISDN) links. Typically, this is done through the deployment of Network Access Servers (NASs) at one or more central sites. Users dial into such NASs, which interact with Authentication, Authorization, and Accounting (AAA) servers to verify the identity of the user, and the set of services that the user is authorized to receive.

In recent times, as more businesses have found the need for high speed Internet connections to their private corporate networks, there has been significant interest in the deployment of CPE based VPNs running across the Internet. This has been driven typically by the ubiquity and distance insensitive pricing of current Internet services, that can result in significantly lower costs than typical dedicated or leased line services.

The notion of using the Internet for private communications is not new, and many techniques, such as controlled route leaking, have been used for this purpose [3]. Only in recent times, however, have the appropriate IP mechanisms needed to meet customer requirements for VPNs all come together. These requirements include the following:

2.1.1 Opaque Packet Transport:

The traffic carried within a VPN may have no relation to the traffic on the IP backbone, either because the traffic is multiprotocol, or because the customer's IP network may use IP addressing unrelated to that of the IP backbone on which the traffic is transported. In particular, the customer's IP network may use non-unique, private IP addressing [4].

2.1.2 Data Security

In general customers using VPNs require some form of data security. There are different trust models applicable to the use of VPNs. One such model is where the customer does not trust the service provider to provide any form of security, and instead implements a VPN using CPE devices that implement firewall functionality and that are connected together using secure tunnels. In this case the service provider is used solely for IP packet transport.

An alternative model is where the customer trusts the service provider to provide a secure managed VPN service. This is similar to the trust involved when a customer utilizes a public switched Frame Relay or ATM service, in that the customer trusts that packets will not be misdirected, injected into the network in an unauthorized manner, snooped on, modified in transit, or subjected to traffic analysis by unauthorized parties.

With this model providing firewall functionality and secure packet transport services is the responsibility of the service provider. Different levels of security may be needed within the provider backbone, depending on the deployment scenario used. If the VPN traffic is contained within a single provider's IP backbone then strong security mechanisms, such as those provided by the IP Security protocol suite (IPSec) [5], may not be necessary for tunnels between backbone nodes. If the VPN traffic traverses networks or equipment owned by multiple administrations then strong security mechanisms may be appropriate. Also a strong level of security may be applied by a provider to customer traffic to address a customer perception that IP networks, and particularly the Internet, are insecure. Whether or not this perception is correct it is one that must be addressed by the VPN implementation.

2.1.3 Quality of Service Guarantees

In addition to ensuring communication privacy, existing private networking techniques, building upon physical or link layer mechanisms, also offer various types of quality of service guarantees. In particular, leased and dial up lines offer both bandwidth and latency guarantees, while dedicated connection technologies like ATM and Frame Relay have extensive mechanisms for similar guarantees. As IP based VPNs become more widely deployed, there will be market demand for similar guarantees, in order to ensure end to end application transparency. While the ability of IP based VPNs to offer such guarantees will depend greatly upon the commensurate capabilities of the underlying IP backbones, a VPN framework must also address the means by which VPN systems can utilize such capabilities, as they evolve.

2.1.4 Tunneling Mechanism

Together, the first two of the requirements listed above imply that VPNs must be implemented through some form of IP tunneling mechanism, where the packet formats and/or the addressing used within the VPN can be unrelated to that used to route the tunneled packets across the IP backbone. Such tunnels, depending upon their form, can provide some level of intrinsic data security, or this can also be enhanced using other mechanisms (e.g., IPSec).

Furthermore, as discussed later, such tunneling mechanisms can also be mapped into evolving IP traffic management mechanisms. There are already defined a large number of IP tunneling mechanisms. Some of these are well suited to VPN applications, as discussed in section 3.0.

2.2 CPE and Network Based VPNs

Most current VPN implementations are based on CPE equipment. VPN capabilities are being integrated into a wide variety of CPE devices, ranging from firewalls to WAN edge routers and specialized VPN termination devices. Such equipment may be bought and deployed by customers, or may be deployed (and often remotely managed) by service providers in an outsourcing service.

There is also significant interest in 'network based VPNs', where the operation of the VPN is outsourced to an Internet Service Provider (ISP), and is implemented on network as opposed to CPE equipment. There is significant interest in such solutions both by customers seeking to reduce support costs and by ISPs seeking new revenue sources. Supporting VPNs in the network allows the use of particular mechanisms which may lead to highly efficient and cost effective VPN solutions, with common equipment and operations support amortized across large numbers of customers.

Most of the mechanisms discussed below can apply to either CPE based or network based VPNs. However particular mechanisms are likely to prove applicable only to the latter, since they leverage tools (e.g., piggybacking on routing protocols) which are accessible only to ISPs and which are unlikely to be made available to any customer, or even hosted on ISP owned and operated CPE, due to the problems of coordinating joint management of the CPE gear by both the ISP and the customer. This document will indicate which techniques are likely to apply only to network based VPNs.

2.3 VPNs and Extranets

The term 'extranet' is commonly used to refer to a scenario whereby two or more companies have networked access to a limited amount of each other's corporate data. For example a manufacturing company might use an extranet for its suppliers to allow it to query databases for the pricing and availability of components, and then to order and track the status of outstanding orders. Another example is joint software development, for instance, company A allows one development group within company B to access its operating system source code, and company B allows one development group in company A to access its security software. Note that the access policies can get arbitrarily complex. For example company B may internally restrict access to its security software to groups in certain geographic locations to comply with export control laws, for example.

A key feature of an extranet is thus the control of who can access what data, and this is essentially a policy decision. Policy decisions are typically enforced today at the interconnection points between different domains, for example between a private network and the Internet, or between a software test lab and the rest of the company network. The enforcement may be done via a firewall, router with access list functionality, application gateway, or any similar device capable of applying policy to transit traffic. Policy controls may be implemented within a corporate network, in addition to between corporate networks. Also the interconnections between networks could be a set of bilateral links, or could be a separate network, perhaps maintained by an industry consortium. This separate network could itself be a VPN or a physical network.

Introducing VPNs into a network does not require any change to this model. Policy can be enforced between two VPNs, or between a VPN and the Internet, in exactly the same manner as is done today without VPNs. For example two VPNs could be interconnected, with each administration locally imposing its own policy controls, via a firewall, on all traffic that enters its VPN from the outside, whether from another VPN or from the Internet.

This model of a VPN provides for a separation of policy from the underlying mode of packet transport used. For example, a router may direct voice traffic to ATM Virtual Channel Connections (VCCs) for guaranteed QoS, non-local internal company traffic to secure tunnels, and other traffic to a link to the Internet. In the past the secure tunnels may have been frame relay circuits, now they may also be secure IP tunnels or MPLS Label Switched Paths (LSPs)

Other models of a VPN are also possible. For example there is a model whereby a set of application flows is mapped into a VPN. As the policy rules imposed by a network administrator can get quite complex, the number of distinct sets of application flows that are used in the policy rulebase, and hence the number of VPNs, can thus grow quite large, and there can be multiple overlapping VPNs. However there is little to be gained by introducing such new complexity into a network. Instead a VPN should be viewed as a direct analogue to a physical network, as this allows the leveraging of existing protocols and procedures, and the current expertise and skill sets of network administrators and customers.

3.0 VPN Tunneling

As noted above in section 2.1, VPNs must be implemented using some form of tunneling mechanism. This section looks at the generic requirements for such VPN tunneling mechanisms. A number of characteristics and aspects common to any link layer protocol are taken and compared with the features offered by existing tunneling protocols. This provides a basis for comparing different protocols and is also useful to highlight areas where existing tunneling protocols could benefit from extensions to better support their operation in a VPN environment.

An IP tunnel connecting two VPN endpoints is a basic building block from which a variety of different VPN services can be constructed. An IP tunnel operates as an overlay across the IP backbone, and the traffic sent through the tunnel is opaque to the underlying IP backbone. In effect the IP backbone is being used as a link layer technology, and the tunnel forms a point-to-point link.

A VPN device may terminate multiple IP tunnels and forward packets between these tunnels and other network interfaces in different ways. In the discussion of different types of VPNs, in later sections of this document, the primary distinguishing characteristic of these different types is the manner in which packets are forwarded between interfaces (e.g., bridged or routed). There is a direct analogy with how existing networking devices are characterized today. A two-port repeater just forwards packets between its ports, and does not examine the contents of the packet. A bridge forwards packets using Media Access Control (MAC) layer information contained in the packet, while a router forwards packets using layer 3 addressing information contained in the packet. Each of these three scenarios has a direct VPN analogue, as discussed later. Note that an IP tunnel is viewed as just another sort of link, which can be concatenated with another link, bound to a bridge forwarding table, or bound to an IP forwarding table, depending on the type of VPN.

The following sections look at the requirements for a generic IP tunneling protocol that can be used as a basic building block to construct different types of VPNs.

3.1 Tunneling Protocol Requirements for VPNs

There are numerous IP tunneling mechanisms, including IP/IP [6], Generic Routing Encapsulation (GRE) tunnels [7], Layer 2 Tunneling Protocol (L2TP) [8], IPsec [5], and Multiprotocol Label Switching (MPLS) [9]. Note that while some of these protocols are not often thought of as tunneling protocols, they do each allow for opaque transport of frames as packet payload across an IP network, with forwarding disjoint from the address fields of the encapsulated packets.

Note, however, that there is one significant distinction between each of the IP tunneling protocols mentioned above, and MPLS. MPLS can be viewed as a specific link layer for IP, insofar as MPLS specific mechanisms apply only within the scope of an MPLS network, whereas IP based mechanisms extend to the extent of IP reachability. As such, VPN mechanisms built directly upon MPLS tunneling mechanisms cannot, by definition, extend outside the scope of MPLS networks, any more so than, for instance, ATM based mechanisms such as LANE can extend outside of ATM networks. Note however, that an MPLS network can span many different link layer technologies, and so, like an IP network, its scope is not limited by the specific link layers used. A number of proposals for defining a set of mechanisms to allow for interoperable VPNs specifically over MPLS networks have also been produced ([10] [11] [12] [13], [14] and [15]).

There are a number of desirable requirements for a VPN tunneling mechanism, however, that are not all met by the existing tunneling mechanisms. These requirements include:

3.1.1 Multiplexing

There are cases where multiple VPN tunnels may be needed between the same two IP endpoints. This may be needed, for instance, in cases where the VPNs are network based, and each end point supports multiple customers. Traffic for different customers travels over separate tunnels between the same two physical devices. A multiplexing field is needed to distinguish which packets belong to which tunnel. Sharing a tunnel in this manner may also reduce the latency and processing burden of tunnel set up. Of the existing IP tunneling mechanisms, L2TP (via the tunnel-id and session-id fields), MPLS (via the label) and IPsec (via the Security Parameter Index (SPI) field) have a multiplexing mechanism. Strictly speaking GRE does not have a multiplexing field. However the key field, which was

intended to be used for authenticating the source of a packet, has sometimes been used as a multiplexing field. IP/IP does not have a multiplexing field.

The IETF [16] and the ATM Forum [17] have standardized on a single format for a globally unique identifier used to identify a VPN (a VPN-ID). A VPN-ID can be used in the control plane, to bind a tunnel to a VPN at tunnel establishment time, or in the data plane, to identify the VPN associated with a packet, on a per-packet basis. In the data plane a VPN encapsulation header can be used by MPLS, MPOA and other tunneling mechanisms to aggregate packets for different VPNs over a single tunnel. In this case an explicit indication of VPN-ID is included with every packet, and no use is made of any tunnel specific multiplexing field. In the control plane a VPN-ID field can be included in any tunnel establishment signalling protocol to allow for the association of a tunnel (e.g., as identified by the SPI field) with a VPN. In this case there is no need for a VPN-ID to be included with every data packet. This is discussed further in section 5.3.1.

3.1.2 Signalling Protocol

There is some configuration information that must be known by an end point in advance of tunnel establishment, such as the IP address of the remote end point, and any relevant tunnel attributes required, such as the level of security needed. Once this information is available, the actual tunnel establishment can be completed in one of two ways - via a management operation, or via a signalling protocol that allows tunnels to be established dynamically.

An example of a management operation would be to use an SNMP Management Information Base (MIB) to configure various tunneling parameters, e.g., MPLS labels, source addresses to use for IP/IP or GRE tunnels, L2TP tunnel-ids and session-ids, or security association parameters for IPsec.

Using a signalling protocol can significantly reduce the management burden however, and as such, is essential in many deployment scenarios. It reduces the amount of configuration needed, and also reduces the management co-ordination needed if a VPN spans multiple administrative domains. For example, the value of the multiplexing field, described above, is local to the node assigning the value, and can be kept local if distributed via a signalling protocol, rather than being first configured into a management station and then distributed to the relevant nodes. A signalling protocol also allows nodes that are mobile or are only intermittently connected to establish tunnels on demand.

When used in a VPN environment a signalling protocol should allow for the transport of a VPN-ID to allow the resulting tunnel to be associated with a particular VPN. It should also allow tunnel attributes to be exchanged or negotiated, for example the use of frame sequencing or the use of multiprotocol transport. Note that the role of the signalling protocol need only be to negotiate tunnel attributes, not to carry information about how the tunnel is used, for example whether the frames carried in the tunnel are to be forwarded at layer 2 or layer 3. (This is similar to Q.2931 ATM signalling - the same signalling protocol is used to set up Classical IP logical subnetworks as well as for LANE emulated LANs.

Of the various IP tunneling protocols, the following ones support a signalling protocol that could be adapted for this purpose: L2TP (the L2TP control protocol), IPSec (the Internet Key Exchange (IKE) protocol [18]), and GRE (as used with mobile-ip tunneling [19]). Also there are two MPLS signalling protocols that can be used to establish LSP tunnels. One uses extensions to the MPLS Label Distribution Protocol (LDP) protocol [20], called Constraint-Based Routing LDP (CR-LDP) [21], and the other uses extensions to the Resource Reservation Protocol (RSVP) for LSP tunnels [22].

3.1.3 Data Security

A VPN tunneling protocol must support mechanisms to allow for whatever level of security may be desired by customers, including authentication and/or encryption of various strengths. None of the tunneling mechanisms discussed, other than IPSec, have intrinsic security mechanisms, but rely upon the security characteristics of the underlying IP backbone. In particular, MPLS relies upon the explicit labeling of label switched paths to ensure that packets cannot be misdirected, while the other tunneling mechanisms can all be secured through the use of IPSec. For VPNs implemented over non-IP backbones (e.g., MPOA, Frame Relay or ATM virtual circuits), data security is implicitly provided by the layer two switch infrastructure.

Overall VPN security is not just a capability of the tunnels alone, but has to be viewed in the broader context of how packets are forwarded onto those tunnels. For example with VPRNs implemented with virtual routers, the use of separate routing and forwarding table instances ensures the isolation of traffic between VPNs. Packets on one VPN cannot be misrouted to a tunnel on a second VPN since those tunnels are not visible to the forwarding table of the first VPN.

If some form of signalling mechanism is used by one VPN end point to dynamically establish a tunnel with another endpoint, then there is a requirement to be able to authenticate the party attempting the tunnel establishment. IPsec has an array of schemes for this purpose, allowing, for example, authentication to be based on pre-shared keys, or to use digital signatures and certificates. Other tunneling schemes have weaker forms of authentication. In some cases no authentication may be needed, for example if the tunnels are provisioned, rather than dynamically established, or if the trust model in use does not require it.

Currently the IPsec Encapsulating Security Payload (ESP) protocol [23] can be used to establish SAs that support either encryption or authentication or both. However the protocol specification precludes the use of an SA where neither encryption or authentication is used. In a VPN environment this "null/null" option is useful, since other aspects of the protocol (e.g., that it supports tunneling and multiplexing) may be all that is required. In effect the "null/null" option can be viewed as just another level of data security.

3.1.4 Multiprotocol Transport

In many applications of VPNs, the VPN may carry opaque, multiprotocol traffic. As such, the tunneling protocol used must also support multiprotocol transport. L2TP is designed to transport Point-to-Point Protocol (PPP) [24] packets, and thus can be used to carry multiprotocol traffic since PPP itself is multiprotocol. GRE also provides for the identification of the protocol being tunneled. IP/IP and IPsec tunnels have no such protocol identification field, since the traffic being tunneled is assumed to be IP.

It is possible to extend the IPsec protocol suite to allow for the transport of multiprotocol packets. This can be achieved, for example, by extending the signalling component of IPsec - IKE, to indicate the protocol type of the traffic being tunneled, or to carry a packet multiplexing header (e.g., an LLC/SNAP header or GRE header) with each tunneled packet. This approach is similar to that used for the same purpose in ATM networks, where signalling is used to indicate the encapsulation used on the VCC, and where packets sent on the VCC can use either an LLC/SNAP header or be placed directly into the AAL5 payload, the latter being known as VC-multiplexing (see [25]).

3.1.5 Frame Sequencing

One quality of service attribute required by customers of a VPN may be frame sequencing, matching the equivalent characteristic of physical leased lines or dedicated connections. Sequencing may be

required for the efficient operation of particular end to end protocols or applications. In order to implement frame sequencing, the tunneling mechanism must support a sequencing field. Both L2TP and GRE have such a field. IPsec has a sequence number field, but it is used by a receiver to perform an anti-replay check, not to guarantee in-order delivery of packets.

It is possible to extend IPsec to allow the use of the existing sequence field to guarantee in-order delivery of packets. This can be achieved, for example, by using IKE to negotiate whether or not sequencing is to be used, and to define an end point behaviour which preserves packet sequencing.

3.1.6. Tunnel Maintenance

The VPN end points must monitor the operation of the VPN tunnels to ensure that connectivity has not been lost, and to take appropriate action (such as route recalculation) if there has been a failure.

There are two approaches possible. One is for the tunneling protocol itself to periodically check in-band for loss of connectivity, and to provide an explicit indication of failure. For example L2TP has an optional keep-alive mechanism to detect non-operational tunnels.

The other approach does not require the tunneling protocol itself to perform this function, but relies on the operation of some out-of-band mechanism to determine loss of connectivity. For example if a routing protocol such as Routing Information Protocol (RIP) [26] or Open Shortest Path First (OSPF) [27] is run over a tunnel mesh, a failure to hear from a neighbor within a certain period of time will result in the routing protocol declaring the tunnel to be down. Another out-of-band approach is to perform regular ICMP pings with a peer. This is generally sufficient assurance that the tunnel is operational, due to the fact the tunnel also runs across the same IP backbone.

When tunnels are established dynamically a distinction needs to be drawn between the static and dynamic tunnel information needed. Before a tunnel can be established some static information is needed by a node, such as the identify of the remote end point and the attributes of the tunnel to propose and accept. This is typically put in place as a result of a configuration operation. As a result of the signalling exchange to establish a tunnel, some dynamic state is established in each end point, such as the value of the multiplexing field or keys to be used. For example with IPsec, the establishment of a Security Association (SA) puts in place the keys to be used for the lifetime of that SA.

Different policies may be used as to when to trigger the establishment of a dynamic tunnel. One approach is to use a data-driven approach and to trigger tunnel establishment whenever there is data to be transferred, and to timeout the tunnel due to inactivity. This approach is particularly useful if resources for the tunnel are being allocated in the network for QoS purposes. Another approach is to trigger tunnel establishment whenever the static tunnel configuration information is installed, and to attempt to keep the tunnel up all the time.

3.1.7 Large MTUs

An IP tunnel has an associated Maximum Transmission Unit (MTU), just like a regular link. It is conceivable that this MTU may be larger than the MTU of one or more individual hops along the path between tunnel endpoints. If so, some form of frame fragmentation will be required within the tunnel.

If the frame to be transferred is mapped into one IP datagram, normal IP fragmentation will occur when the IP datagram reaches a hop with an MTU smaller than the IP tunnel's MTU. This can have undesirable performance implications at the router performing such mid-tunnel fragmentation.

An alternative approach is for the tunneling protocol itself to incorporate a segmentation and reassembly capability that operates at the tunnel level, perhaps using the tunnel sequence number and an end-of-message marker of some sort. (Note that multilink PPP uses a mechanism similar to this to fragment packets). This avoids IP level fragmentation within the tunnel itself. None of the existing tunneling protocols support such a mechanism.

3.1.8 Minimization of Tunnel Overhead

There is clearly benefit in minimizing the overhead of any tunneling mechanisms. This is particularly important for the transport of jitter and latency sensitive traffic such as packetized voice and video. On the other hand, the use of security mechanisms, such as IPsec, do impose their own overhead, hence the objective should be to minimize overhead over and above that needed for security, and to not burden those tunnels in which security is not mandatory with unnecessary overhead.

One area where the amount of overhead may be significant is when voluntary tunneling is used for dial-up remote clients connecting to a VPN, due to the typically low bandwidth of dial-up links. This is discussed further in section 6.3.

3.1.9 Flow and congestion control

During the development of the L2TP protocol procedures were developed for flow and congestion control. These were necessitated primarily because of the need to provide adequate performance over lossy networks when PPP compression is used, which, unlike IP Payload Compression Protocol (IPComp) [28], is stateful across packets. Another motivation was to accommodate devices with very little buffering, used for example to terminate low speed dial-up lines. However the flow and congestion control mechanisms defined in the final version of the L2TP specification are used only for the control channels, and not for data traffic.

In general the interactions between multiple layers of flow and congestion control schemes can be very complex. Given the predominance of TCP traffic in today's networks and the fact that TCP has its own end-to-end flow and congestion control mechanisms, it is not clear that there is much benefit to implementing similar mechanisms within tunneling protocols. Good flow and congestion control schemes, that can adapt to a wide variety of network conditions and deployment scenarios are complex to develop and test, both in themselves and in understanding the interaction with other schemes that may be running in parallel. There may be some benefit, however, in having the capability whereby a sender can shape traffic to the capacity of a receiver in some manner, and in providing the protocol mechanisms to allow a receiver to signal its capabilities to a sender. This is an area that may benefit from further study.

Note also the work of the Performance Implications of Link Characteristics (PILC) working group of the IETF, which is examining how the properties of different network links can have an impact on the performance of Internet protocols operating over those links.

3.1.10 QoS / Traffic Management

As noted above, customers may require that VPNs yield similar behaviour to physical leased lines or dedicated connections with respect to such QoS parameters as loss rates, jitter, latency and bandwidth guarantees. How such guarantees could be delivered will, in general, be a function of the traffic management characteristics of the VPN nodes themselves, and the access and backbone networks across which they are connected.

A full discussion of QoS and VPNs is outside the scope of this document, however by modeling a VPN tunnel as just another type of link layer, many of the existing mechanisms developed for ensuring QoS over physical links can also be applied. For example at a VPN node, the mechanisms of policing, marking, queuing, shaping and

scheduling can all be applied to VPN traffic with VPN-specific parameters, queues and interfaces, just as for non-VPN traffic. The techniques developed for Diffserv, Intserv and for traffic engineering in MPLS are also applicable. See also [29] for a discussion of QoS and VPNs.

It should be noted, however, that this model of tunnel operation is not necessarily consistent with the way in which specific tunneling protocols are currently modeled. While a model is an aid to comprehension, and not part of a protocol specification, having differing models can complicate discussions, particularly if a model is misinterpreted as being part of a protocol specification or as constraining choice of implementation method. For example, IPsec tunnel processing can be modeled both as an interface and as an attribute of a particular packet flow.

3.2 Recommendations

IPsec is needed whenever there is a requirement for strong encryption or strong authentication. It also supports multiplexing and a signalling protocol - IKE. However extending the IPsec protocol suite to also cover the following areas would be beneficial, in order to better support the tunneling requirements of a VPN environment.

- the transport of a VPN-ID when establishing an SA (3.1.2)
- a null encryption and null authentication option (3.1.3)
- multiprotocol operation (3.1.4)
- frame sequencing (3.1.5)

L2TP provides no data security by itself, and any PPP security mechanisms used do not apply to the L2TP protocol itself, so that in order for strong security to be provided L2TP must run over IPsec. Defining specific modes of operation for IPsec when it is used to support L2TP traffic will aid interoperability. This is currently a work item for the proposed L2TP working group.

4.0 VPN Types: Virtual Leased Lines

The simplest form of a VPN is a 'Virtual Leased Line' (VLL) service. In this case a point-to-point link is provided to a customer, connecting two CPE devices, as illustrated below. The link layer type used to connect the CPE devices to the ISP nodes can be any link layer type, for example an ATM VCC or a Frame Relay circuit. The CPE devices can be either routers bridges or hosts.

The two ISP nodes are both connected to an IP network, and an IP tunnel is set up between them. Each ISP node is configured to bind the stub link and the IP tunnel together at layer 2 (e.g., an ATM VCC and the IP tunnel). Frames are relayed between the two links. For example the ATM Adaptation Layer 5 (AAL5) payload is taken and encapsulated in an IPSec tunnel, and vice versa. The contents of the AAL5 payload are opaque to the ISP node, and are not examined there.




Figure 4.1: VLL Example

To a customer it looks the same as if a single ATM VCC or Frame Relay circuit were used to interconnect the CPE devices, and the customer could be unaware that part of the circuit was in fact implemented over an IP backbone. This may be useful, for example, if a provider wishes to provide a LAN interconnect service using ATM as the network interface, but does not have an ATM network that directly interconnects all possible customer sites.

It is not necessary that the two links used to connect the CPE devices to the ISP nodes be of the same media type, but in this case the ISP nodes cannot treat the traffic in an opaque manner, as described above. Instead the ISP nodes must perform the functions of an interworking device between the two media types (e.g., ATM and Frame Relay), and perform functions such as LLC/SNAP to NLPID conversion, mapping between ARP protocol variants and performing any media specific processing that may be expected by the CPE devices (e.g., ATM OAM cell handling or Frame Relay XID exchanges).

The IP tunneling protocol used must support multiprotocol operation and may need to support sequencing, if that characteristic is important to the customer traffic. If the tunnels are established using a signalling protocol, they may be set up in a data driven manner, when a frame is received from a customer link and no tunnel exists, or the tunnels may be established at provisioning time and kept up permanently.

Note that the use of the term 'VLL' in this document is different to that used in the definition of the Diffserv Expedited Forwarding Per Hop Behaviour (EF-PHB) [30]. In that document a VLL is used to mean a low latency, low jitter, assured bandwidth path, which can be provided using the described PHB. Thus the focus there is primarily on link characteristics that are temporal in nature. In this document the term VLL does not imply the use of any specific QoS mechanism, Diffserv or otherwise. Instead the focus is primarily on link characteristics that are more topological in nature, (e.g., such as constructing a link which includes an IP tunnel as one segment of the link). For a truly complete emulation of a link layer both the temporal and topological aspects need to be taken into account.

5.0 VPN Types: Virtual Private Routed Networks

5.1 VPRN Characteristics

A Virtual Private Routed Network (VPRN) is defined to be the emulation of a multi-site wide area routed network using IP facilities. This section looks at how a network-based VPRN service can be provided. CPE-based VPRNs are also possible, but are not specifically discussed here. With network-based VPRNs many of the issues that need to be addressed are concerned with configuration and operational issues, which must take into account the split in administrative responsibility between the service provider and the service user.

The distinguishing characteristic of a VPRN, in comparison to other types of VPNs, is that packet forwarding is carried out at the network layer. A VPRN consists of a mesh of IP tunnels between ISP routers, together with the routing capabilities needed to forward traffic received at each VPRN node to the appropriate destination site. Attached to the ISP routers are CPE routers connected via one or more links, termed 'stub' links. There is a VPRN specific forwarding table at each ISP router to which members of the VPRN are connected. Traffic is forwarded between ISP routers, and between ISP routers and customer sites, using these forwarding tables, which contain network layer reachability information (in contrast to a Virtual Private LAN Segment type of VPN (VPLS) where the forwarding tables contain MAC layer reachability information - see section 7.0).

An example VPRN is illustrated in the following diagram, which shows 3 ISP edge routers connected via a full mesh of IP tunnels, used to interconnect 4 CPE routers. One of the CPE routers is multihomed to the ISP network. In the multihomed case, all stub links may be active, or, as shown, there may be one primary and one or more backup links to be used in case of failure of the primary. The term 'backdoor' link is used to refer to a link between two customer sites

that does not traverse the ISP network.

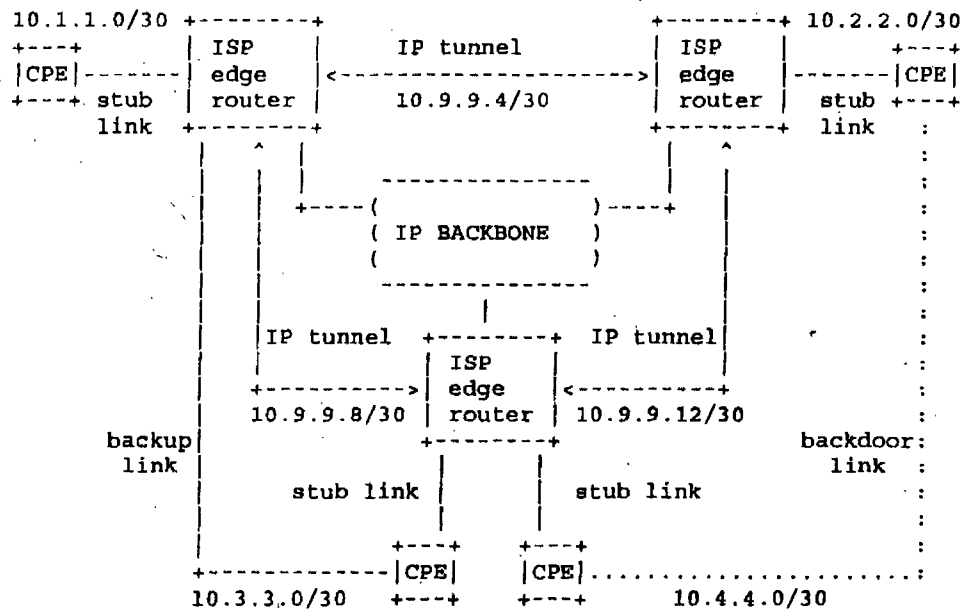


Figure 5.1: VPRN Example

The principal benefit of a VPRN is that the complexity and the configuration of the CPE routers is minimized. To a CPE router, the ISP edge router appears as a neighbor router in the customer's network, to which it sends all traffic, using a default route. The tunnel mesh that is set up to transfer traffic extends between the ISP edge routers, not the CPE routers. In effect the burden of tunnel establishment and maintenance and routing configuration is outsourced to the ISP. In addition other services needed for the operation of a VPN such as the provision of a firewall and QoS processing can be handled by a small number of ISP edge routers, rather than a large number of potentially heterogeneous CPE devices. The introduction and management of new services can also be more easily handled, as this can be achieved without the need to upgrade any CPE equipment. This latter benefit is particularly important when there may be large numbers of residential subscribers using VPN services to access private corporate networks. In this respect the model is somewhat akin to that used for telephony services, whereby new services (e.g., call waiting) can be introduced with no change in subscriber equipment.

The VPRN type of VPN is in contrast to one where the tunnel mesh extends to the CPE routers, and where the ISP network provides layer 2 connectivity alone. The latter case can be implemented either as a set of VLLs between CPE routers (see section 4.0), in which case the ISP network provides a set of layer 2 point-to-point links, or as a VPLS (see section 7.0), in which case the ISP network is used to emulate a multiaccess LAN segment. With these scenarios a customer may have more flexibility (e.g., any IGP or any protocol can be run across all customer sites) but this usually comes at the expense of a more complex configuration for the customer. Thus, depending on customer requirements, a VPRN or a VPLS may be the more appropriate solution.

Because a VPRN carries out forwarding at the network layer, a single VPRN only directly supports a single network layer protocol. For multiprotocol support, a separate VPRN for each network layer protocol could be used, or one protocol could be tunneled over another (e.g., non-IP protocols tunneled over an IP VPRN) or alternatively the ISP network could be used to provide layer 2 connectivity only, such as with a VPLS as mentioned above.

The issues to be addressed for VPRNs include initial configuration, determination by an ISP edge router of the set of links that are in each VPRN, the set of other routers that have members in the VPRN, and the set of IP address prefixes reachable via each stub link, determination by a CPE router of the set of IP address prefixes to be forwarded to an ISP edge router, the mechanism used to disseminate stub reachability information to the correct set of ISP routers, and the establishment and use of the tunnels used to carry the data traffic. Note also that, although discussed first for VPRNs, many of these issues also apply to the VPLS scenario described later, with the network layer addresses being replaced by link layer addresses.

Note that VPRN operation is decoupled from the mechanisms used by the customer sites to access the Internet. A typical scenario would be for the ISP edge router to be used to provide both VPRN and Internet connectivity to a customer site. In this case the CPE router just has a default route pointing to the ISP edge router, with the latter being responsible for steering private traffic to the VPRN and other traffic to the Internet, and providing firewall functionality between the two domains. Alternatively a customer site could have Internet connectivity via an ISP router not involved in the VPRN, or even via a different ISP. In this case the CPE device is responsible for splitting the traffic into the two domains and providing firewall functionality.

5.1.1 Topology

The topology of a VPRN may consist of a full mesh of tunnels between each VPRN node, or may be an arbitrary topology, such as a set of remote offices connected to the nearest regional site, with these regional sites connected together via a full or partial mesh. With VPRNs using IP tunnels there is much less cost assumed with full meshing than in cases where physical resources (e.g., a leased line) must be allocated for each connected pair of sites, or where the tunneling method requires resources to be allocated in the devices used to interconnect the edge routers (e.g., Frame Relay DLCIs). A full mesh topology yields optimal routing, since it precludes the need for traffic between two sites to traverse a third. Another attraction of a full mesh is that there is no need to configure topology information for the VPRN. Instead, given the member routers of a VPRN, the topology is implicit. If the number of ISP edge routers in a VPRN is very large, however, a full mesh topology may not be appropriate, due to the scaling issues involved, for example, the growth in the number of tunnels needed between sites, (which for n sites is $n(n-1)/2$), or the number of routing peers per router. Network policy may also lead to non full mesh topologies, for example an administrator may wish to set up the topology so that traffic between two remote sites passes through a central site, rather than go directly between the remote sites. It is also necessary to deal with the scenario where there is only partial connectivity across the IP backbone under certain error conditions (e.g. A can reach B, and B can reach C, but A cannot reach C directly), which can occur if policy routing is being used.

For a network-based VPRN, it is assumed that each customer site CPE router connects to an ISP edge router through one or more point-to-point stub links (e.g. leased lines, ATM or Frame Relay connections). The ISP routers are responsible for learning and disseminating reachability information amongst themselves. The CPE routers must learn the set of destinations reachable via each stub link, though this may be as simple as a default route.

The stub links may either be dedicated links, set up via provisioning, or may be dynamic links set up on demand, for example using PPP, voluntary tunneling (see section 6.3), or ATM signalling. With dynamic links it is necessary to authenticate the subscriber, and determine the authorized resources that the subscriber can access (e.g. which VPRNs the subscriber may join). Other than the way the subscriber is initially bound to the VPRN, (and this process may involve extra considerations such as dynamic IP address assignment), the subsequent VPRN mechanisms and services can be used for both types of subscribers in the same way.

5.1.2 Addressing

The addressing used within a VPRN may have no relation to the addressing used on the IP backbone over which the VPRN is instantiated. In particular non-unique private IP addressing may be used [4]. Multiple VPRNs may be instantiated over the same set of physical devices, and they may use the same or overlapping address spaces.

5.1.3 Forwarding

For a VPRN the tunnel mesh forms an overlay network operating over an IP backbone. Within each of the ISP edge routers there must be VPN specific forwarding state to forward packets received from stub links ('ingress traffic') to the appropriate next hop router, and to forward packets received from the core ('egress traffic') to the appropriate stub link. For cases where an ISP edge router supports multiple stub links belonging to the same VPRN, the tunnels can, as a local matter, either terminate on the edge router, or on a stub link. In the former case a VPN specific forwarding table is needed for egress traffic, in the latter case it is not. A VPN specific forwarding table is generally needed in the ingress direction, in order to direct traffic received on a stub link onto the correct IP tunnel towards the core.

Also since a VPRN operates at the internetwork layer, the IP packets sent over a tunnel will have their Time to Live (TTL) field decremented in the normal manner, preventing packets circulating indefinitely in the event of a routing loop within the VPRN.

5.1.4 Multiple concurrent VPRN connectivity

Note also that a single customer site may belong concurrently to multiple VPRNs and may want to transmit traffic both onto one or more VPRNs and to the default Internet, over the same stub link. There are a number of possible approaches to this problem, but these are outside the scope of this document.

5.2 VPRN Related Work

VPRN requirements and mechanisms have been discussed previously in a number of different documents. One of the first was [10], which showed how the same VPN functionality can be implemented over both MPLS and non-MPLS networks. Some others are briefly discussed below.

There are two main variants as regards the mechanisms used to provide VPRN membership and reachability functionality, - overlay and piggybacking. These are discussed in greater detail in sections

5.3.2, 5.3.3 and 5.3.4 below. An example of the overlay model is described in [14], which discusses the provision of VPRN functionality by means of a separate per-VPN routing protocol instance and route and forwarding table instantiation, otherwise known as virtual routing. Each VPN routing instance is isolated from any other VPN routing instance, and from the routing used across the backbone. As a result any routing protocol (e.g. OSPF, RIP2, IS-IS) can be run with any VPRN, independently of the routing protocols used in other VPRNs, or in the backbone itself. The VPN model described in [12] is also an overlay VPRN model using virtual routing. That document is specifically geared towards the provision of VPRN functionality over MPLS backbones, and it describes how VPRN membership dissemination can be automated over an MPLS backbone, by performing VPN neighbor discovery over the base MPLS tunnel mesh. [31] extends the virtual routing model to include VPN areas, and VPN border routers which route between VPN areas. VPN areas may be defined for administrative or technical reasons, such as different underlying network infrastructures (e.g. ATM, MPLS, IP).

In contrast [15] describes the provision of VPN functionality using a piggybacking approach for membership and reachability dissemination, with this information being piggybacked in Border Gateway Protocol 4 (BGP) [32] packets. VPNs are constructed using BGP policies, which are used to control which sites can communicate with each other. [13] also uses BGP for piggybacking membership information, and piggybacks reachability information on the protocol used to establish MPLS LSPs (CR-LDP or extended RSVP). Unlike the other proposals, however, this proposal requires the participation on the CPE router to implement the VPN functionality.

5.3 VPRN Generic Requirements

There are a number of common requirements which any network-based VPRN solution must address, and there are a number of different mechanisms that can be used to meet these requirements. These generic issues are

- 1) The use of a globally unique VPN identifier in order to be able to refer to a particular VPN.
- 2) VPRN membership determination. An edge router must learn of the local stub links that are in each VPRN, and must learn of the set of other routers that have members in that VPRN.
- 3) Stub link reachability information. An edge router must learn the set of addresses and address prefixes reachable via each stub link.

- 4) Intra-VPRN reachability information. Once an edge router has determined the set of address prefixes associated with each of its stub links, then this information must be disseminated to each other edge router in the VPRN.
- 5) Tunneling mechanism. An edge router must construct the necessary tunnels to other routers that have members in the VPRN, and must perform the encapsulation and decapsulation necessary to send and receive packets over the tunnels.

5.3.1 VPN Identifier

The IETF [16] and the ATM Forum [17] have standardized on a single format for a globally unique identifier used to identify a VPN - a VPN-ID. Only the format of the VPN-ID has been defined, not its semantics or usage. The aim is to allow its use for a wide variety of purposes, and to allow the same identifier to be used with different technologies and mechanisms. For example a VPN-ID can be included in a MIB to identify a VPN for management purposes. A VPN-ID can be used in a control plane protocol, for example to bind a tunnel to a VPN at tunnel establishment time. All packets that traverse the tunnel are then implicitly associated with the identified VPN. A VPN-ID can be used in a data plane encapsulation, to allow for an explicit per-packet identification of the VPN associated with the packet. If a VPN is implemented using different technologies (e.g., IP and ATM) in a network, the same identifier can be used to identify the VPN across the different technologies. Also if a VPN spans multiple administrative domains the same identifier can be used everywhere.

Most of the VPN schemes developed (e.g. [11], [12], [13], [14]) require the use of a VPN-ID that is carried in control and/or data packets, which is used to associate the packet with a particular VPN. Although the use of a VPN-ID in this manner is very common, it is not universal. [15] describes a scheme where there is no protocol field used to identify a VPN in this manner. In this scheme the VPNs as understood by a user, are administrative constructs, built using BGP policies. There are a number of attributes associated with VPN routes, such as a route distinguisher, and origin and target "VPN", that are used by the underlying protocol mechanisms for disambiguation and scoping, and these are also used by the BGP policy mechanism in the construction of VPNs, but there is nothing corresponding with the VPN-ID as used in the other documents.

Note also that [33] defines a multiprotocol encapsulation for use over ATM AAL5 that uses the standard VPN-ID format.

5.3.2 VPN Membership Information Configuration and Dissemination

In order to establish a VPRN, or to insert new customer sites into an established VPRN, an ISP edge router must determine which stub links are associated with which VPRN. For static links (e.g. an ATM VCC) this information must be configured into the edge router, since the edge router cannot infer such bindings by itself. An SNMP MIB allowing for bindings between local stub links and VPN identities is one solution.

For subscribers that attach to the network dynamically (e.g. using PPP or voluntary tunneling) it is possible to make the association between stub link and VPRN as part of the end user authentication processing that must occur with such dynamic links. For example the VPRN to which a user is to be bound may be derived from the domain name the used as part of PPP authentication. If the user is successfully authenticated (e.g. using a Radius server), then the newly created dynamic link can be bound to the correct VPRN. Note that static configuration information is still needed, for example to maintain the list of authorized subscribers for each VPRN, but the location of this static information could be an external authentication server rather than on an ISP edge router. Whether the link was statically or dynamically created, a VPN-ID can be associated with that link to signify to which VPRN it is bound.

After learning which stub links are bound to which VPRN, each edge router must learn either the identity of, or, at least, the route to, each other edge router supporting other stub links in that particular VPRN. Implicit in the latter is the notion that there exists some mechanism by which the configured edge routers can then use this edge router and/or stub link identity information to subsequently set up the appropriate tunnels between them. The problem of VPRN member dissemination between participating edge routers, can be solved in a variety of ways, discussed below.

5.3.2.1 Directory Lookup

The members of a particular VPRN, that is, the identity of the edge routers supporting stub links in the VPRN, and the set of static stub links bound to the VPRN per edge router, could be configured into a directory, which edge routers could query, using some defined mechanism (e.g. Lightweight Directory Access Protocol (LDAP) [34]), upon startup.

Using a directory allows either a full mesh topology or an arbitrary topology to be configured. For a full mesh, the full list of member routers in a VPRN is distributed everywhere. For an arbitrary topology, different routers may receive different member lists.

Using a directory allows for authorization checking prior to disseminating VPRN membership information, which may be desirable where VPRNs span multiple administrative domains. In such a case, directory to directory protocol mechanisms could also be used to propagate authorized VPRN membership information between the directory systems of the multiple administrative domains.

There also needs to be some form of database synchronization mechanism (e.g. triggered or regular polling of the directory by edge routers, or active pushing of update information to the edge routers by the directory) in order for all edge routers to learn the identity of newly configured sites inserted into an active VPRN, and also to learn of sites removed from a VPRN.

5.3.2.2 Explicit Management Configuration

A VPRN MIB could be defined which would allow a central management system to configure each edge router with the identities of each other participating edge router and the identity of each of the static stub links bound to the VPRN. Like the use of a directory, this mechanism allows both full mesh and arbitrary topologies to be configured. Another mechanism using a centralized management system is to use a policy server and use the Common Open Policy Service (COPS) protocol [35] to distribute VPRN membership and policy information, such as the tunnel attributes to use when establishing a tunnel, as described in [36].

Note that this mechanism allows the management station to impose strict authorization control; on the other hand, it may be more difficult to configure edge routers outside the scope of the management system. The management configuration model can also be considered a subset of the directory method, in that the management directories could use MIBs to push VPRN membership information to the participating edge routers, either subsequent to, or as part of, the local stub link configuration process.

5.3.2.3 Piggybacking in Routing Protocols

VPRN membership information could be piggybacked into the routing protocols run by each edge router across the IP backbone, since this is an efficient means of automatically propagating information throughout the network to other participating edge routers. Specifically, each route advertisement by each edge router could include, at a minimum, the set of VPN identifiers associated with each edge router, and adequate information to allow other edge routers to determine the identity of, and/or, the route to, the particular edge router. Other edge routers would examine received route advertisements to determine if any contained information was

relevant to a supported (i.e., configured) VPRN; this determination could be done by looking for a VPN identifier matching a locally configured VPN. The nature of the piggybacked information, and related issues, such as scoping, and the means by which the nodes advertising particular VPN memberships will be identified, will generally be a function both of the routing protocol and of the nature of the underlying transport.

Using this method all the routers in the network will have the same view of the VPRN membership information, and so a full mesh topology is easily supported. Supporting an arbitrary topology is more difficult, however, since some form of pruning would seem to be needed.

The advantage of the piggybacking scheme is that it allows for efficient information dissemination, but it does require that all nodes in the path, and not just the participating edge routers, be able to accept such modified route advertisements. A disadvantage is that significant administrative complexity may be required to configure scoping mechanisms so as to both permit and constrain the dissemination of the piggybacked advertisements, and in itself this may be quite a configuration burden, particularly if the VPRN spans multiple routing domains (e.g. different autonomous systems / ISPs).

Furthermore, unless some security mechanism is used for routing updates so as to permit only all relevant edge routers to read the piggybacked advertisements, this scheme generally implies a trust model where all routers in the path must perforce be authorized to know this information. Depending upon the nature of the routing protocol, piggybacking may also require intermediate routers, particularly autonomous system (AS) border routers, to cache such advertisements and potentially also re-distribute them between multiple routing protocols.

Each of the schemes described above have merit in particular situations. Note that, in practice, there will almost always be some centralized directory or management system which will maintain VPRN membership information, such as the set of edge routers that are allowed to support a certain VPRN, the bindings of static stub links to VPRNs, or authentication and authorization information for users that access the network via dynamics links. This information needs to be configured and stored in some form of database, so that the additional steps needed to facilitate the configuration of such information into edge routers, and/or, facilitate edge router access to such information, may not be excessively onerous.

5.3.3 Stub Link Reachability Information

There are two aspects to stub site reachability - the means by which VPRN edge routers determine the set of VPRN addresses and address prefixes reachable at each stub site, and the means by which the CPE routers learn the destinations reachable via each stub link. A number of common scenarios are outlined below. In each case the information needed by the ISP edge router is the same - the set of VPRN addresses reachable at the customer site, but the information needed by the CPE router differs.

5.3.3.1 Stub Link Connectivity Scenarios

5.3.3.1.1 Dual VPRN and Internet Connectivity

The CPE router is connected via one link to an ISP edge router, which provides both VPRN and Internet connectivity.

This is the simplest case for the CPE router, as it just needs a default route pointing to the ISP edge router.

5.3.3.1.2 VPRN Connectivity Only

The CPE router is connected via one link to an ISP edge router, which provides VPRN, but not Internet, connectivity.

The CPE router must know the set of non-local VPRN destinations reachable via that link. This may be a single prefix, or may be a number of disjoint prefixes. The CPE router may be either statically configured with this information, or may learn it dynamically by running an instance of an Interior Gateway Protocol (IGP). For simplicity it is assumed that the IGP used for this purpose is RIP, though it could be any IGP. The ISP edge router will inject into this instance of RIP the VPRN routes which it learns by means of one of the intra-VPRN reachability mechanisms described in section 5.3.4. Note that the instance of RIP run to the CPE, and any instance of a routing protocol used to learn intra-VPRN reachability (even if also RIP) are separate, with the ISP edge router redistributing the routes from one instance to another.

5.3.3.1.3 Multihomed Connectivity

The CPE router is multihomed to the ISP network, which provides VPRN connectivity.

In this case all the ISP edge routers could advertise the same VPRN routes to the CPE router, which then sees all VPRN prefixes equally reachable via all links. More specific route redistribution is also possible, whereby each ISP edge router advertises a different set of prefixes to the CPE router.

5.3.3.1.4 Backdoor Links

The CPE router is connected to the ISP network, which provides VPRN connectivity, but also has a backdoor link to another customer site

In this case the ISP edge router will advertise VPRN routes as in case 2 to the CPE device. However now the same destination is reachable via both the ISP edge router and via the backdoor link. If the CPE routers connected to the backdoor link are running the customer's IGP, then the backdoor link may always be the favored link as it will appear as an 'internal' path, whereas the destination as injected via the ISP edge router will appear as an 'external' path (to the customer's IGP). To avoid this problem, assuming that the customer wants the traffic to traverse the ISP network, then a separate instance of RIP should be run between the CPE routers at both ends of the backdoor link, in the same manner as an instance of RIP is run on a stub or backup link between a CPE router and an ISP edge router. This will then also make the backdoor link appear as an external path, and by adjusting the link costs appropriately, the ISP path can always be favored, unless it goes down, when the backdoor link is then used.

The description of the above scenarios covers what reachability information is needed by the ISP edge routers and the CPE routers, and discusses some of the mechanisms used to convey this information. The sections below look at these mechanisms in more detail.

5.3.3.1 Routing Protocol Instance

A routing protocol can be run between the CPE edge router and the ISP edge router to exchange reachability information. This allows an ISP edge router to learn the VPRN prefixes reachable at a customer site, and also allows a CPE router to learn the destinations reachable via the provider network.

The extent of the routing domain for this protocol instance is generally just the ISP edge router and the CPE router although if the customer site is also running the same protocol as its IGP, then the domain may extend into customer site. If the customer site is running a different routing protocol then the CPE router redistributes the routes between the instance running to the ISP edge router, and the instance running into the customer site.

Given the typically restricted scope of this routing instance, a simple protocol will generally suffice. RIP is likely to be the most common protocol used, though any routing protocol, such as OSPF, or BGP run in internal mode (IBGP), could also be used.

Note that the instance of the stub link routing protocol is different from any instance of a routing protocol used for intra-VPRN reachability. For example, if the ISP edge router uses routing protocol piggybacking to disseminate VPRN membership and reachability information across the core, then it may redistribute suitably labeled routes from the CPE routing instance to the core routing instance. The routing protocols used for each instance are decoupled, and any suitable protocol can be used in each case. There is no requirement that the same protocol, or even the same stub link reachability information gathering mechanism, be run between each CPE router and associated ISP edge router in a particular VPRN, since this is a purely local matter.

This decoupling allows ISPs to deploy a common (across all VPRNs) intra-VPRN reachability mechanism, and a common stub link reachability mechanism, with these mechanisms isolated both from each other, and from the particular IGP used in a customer network. In the first case, due to the IGP-IGP boundary implemented on the ISP edge router, the ISP can insulate the intra-VPRN reachability mechanism from misbehaving stub link protocol instances. In the second case the ISP is not required to be aware of the particular IGP running in a customer site. Other scenarios are possible, where the ISP edge routers are running a routing protocol in the same instance as the customer's IGP, but are unlikely to be practical, since it defeats the purpose of a VPRN simplifying CPE router configuration. In cases where a customer wishes to run an IGP across multiple sites, a VPLS solution is more suitable.

Note that if a particular customer site concurrently belongs to multiple VPRNs (or wishes to concurrently communicate with both a VPRN and the Internet), then the ISP edge router must have some means of unambiguously mapping stub link address prefixes to particular VPRNs. A simple way is to have multiple stub links, one per VPRN. It is also possible to run multiple VPRNs over one stub link. This could be done either by ensuring (and appropriately configuring the

ISP edge router to know) that particular disjoint address prefixes are mapped into separate VPRNs, or by tagging the routing advertisements from the CPE router with the appropriate VPN identifier. For example if MPLS was being used to convey stub link reachability information, different MPLS labels would be used to differentiate the disjoint prefixes assigned to particular VPRNs. In any case, some administrative procedure would be required for this coordination.

5.3.3.2 Configuration

The reachability information across each stub link could be manually configured, which may be appropriate if the set of addresses or prefixes is small and static.

5.3.3.3 ISP Administered Addresses

The set of addresses used by each stub site could be administered and allocated via the VPRN edge router, which may be appropriate for small customer sites, typically containing either a single host, or a single subnet. Address allocation can be carried out using protocols such as PPP or DHCP [37], with, for example, the edge router acting as a Radius client and retrieving the customer's IP address to use from a Radius server, or acting as a DHCP relay and examining the DHCP reply message as it is relayed to the customer site. In this manner the edge router can build up a table of stub link reachability information. Although these address assignment mechanisms are typically used to assign an address to a single host, some vendors have added extensions whereby an address prefix can be assigned, with, in some cases, the CPE device acting as a "mini-DHCP" server and assigning addresses for the hosts in the customer site.

Note that with these schemes it is the responsibility of the address allocation server to ensure that each site in the VPN received a disjoint address space. Note also that an ISP would typically only use this mechanism for small stub sites, which are unlikely to have backdoor links.

5.3.3.4 MPLS Label Distribution Protocol

In cases where the CPE router runs MPLS, LDP can be used to convey the set of prefixes at a stub site to a VPRN edge router. Using the downstream unsolicited mode of label distribution the CPE router can distribute a label for each route in the stub site. Note however that the processing carried out by the edge router in this case is more than just the normal LDP processing, since it is learning new routes via LDP, rather than the usual case of learning labels for existing routes that it has learned via standard routing mechanisms.

5.3.4 Intra-VPN Reachability Information

Once an edge router has determined the set of prefixes associated with each of its stub links, then this information must be disseminated to each other edge router in the VPRN. Note also that there is an implicit requirement that the set of reachable addresses within the VPRN be locally unique that is, each VPRN stub link (not performing load sharing) maintain an address space disjoint from any other, so as to permit unambiguous routing. In practical terms, it is also generally desirable, though not required, that this address space be well partitioned i.e., specific, disjoint address prefixes per edge router, so as to preclude the need to maintain and disseminate large numbers of host routes.

The problem of intra-VPN reachability information dissemination can be solved in a number of ways, some of which include the following:

5.3.4.1 Directory Lookup

Along with VPRN membership information, a central directory could maintain a listing of the address prefixes associated with each customer site. Such information could be obtained by the server through protocol interactions with each edge router. Note that the same directory synchronization issues discussed above in section 5.3.2 also apply in this case.

5.3.4.2 Explicit Configuration

The address spaces associated with each edge router could be explicitly configured into each other router. This is clearly a non-scalable solution, particularly when arbitrary topologies are used, and also raises the question of how the management system learns such information in the first place.

5.3.4.3 Local Intra-VPRN Routing Instantiations

In this approach, each edge router runs an instance of a routing protocol (a 'virtual router') per VPRN, running across the VPRN tunnels to each peer edge router, to disseminate intra-VPRN reachability information. Both full-mesh and arbitrary VPRN topologies can be easily supported, since the routing protocol itself can run over any topology. The intra-VPRN routing advertisements could be distinguished from normal tunnel data packets either by being addressed directly to the peer edge router, or by a tunnel specific mechanism.

Note that this intra-VPRN routing protocol need have no relationship either with the IGP of any customer site or with the routing protocols operated by the ISPs in the IP backbone. Depending on the size and scale of the VPRNs to be supported either a simple protocol like RIP or a more sophisticated protocol like OSPF could be used. Because the intra-VPRN routing protocol operates as an overlay over the IP backbone it is wholly transparent to any intermediate routers, and to any edge routers not within the VPRN. This also implies that such routing information can remain opaque to such routers, which may be a necessary security requirements in some cases. Also note that if the routing protocol runs directly over the same tunnels as the data traffic, then it will inherit the same level of security as that afforded the data traffic, for example strong encryption and authentication.

If the tunnels over which an intra-VPRN routing protocol runs are dedicated to a specific VPN (e.g. a different multiplexing field is used for each VPN) then no changes are needed to the routing protocol itself. On the other hand if shared tunnels are used, then it is necessary to extend the routing protocol to allow a VPN-ID field to be included in routing update packets, to allow sets of prefixes to be associated with a particular VPN.

5.3.4.4 Link Reachability Protocol

By link reachability protocol is meant a protocol that allows two nodes, connected via a point-to-point link, to exchange reachability information. Given a full mesh topology, each edge router could run a link reachability protocol, for instance some variation of MPLS CR-LDP, across the tunnel to each peer edge router in the VPRN, carrying the VPN-ID and the reachability information of each VPRN running across the tunnel between the two edge routers. If VPRN membership information has already been distributed to an edge router, then the neighbor discovery aspects of a traditional routing protocol are not needed, as the set of neighbors is already known. TCP connections can be used to interconnect the neighbors, to provide reliability. This approach may reduce the processing burden of running routing protocol instances per VPRN, and may be of particular benefit where a shared tunnel mechanism is used to connect a set of edge routers supporting multiple VPRNs.

Another approach to developing a link reachability protocol would be to base it on IBGP. The problem that needs to be solved by a link reachability protocol is very similar to that solved by IBGP - conveying address prefixes reliably between edge routers.

Using a link reachability protocol it is straightforward to support a full mesh topology - each edge router conveys its own local reachability information to all other routers, but does not redistribute information received from any other router. However once an arbitrary topology needs to be supported, the link reachability protocol needs to develop into a full routing protocol, due to the need to implement mechanisms to avoid loops, and there would seem little benefit in reinventing another routing protocol to deal with this. Some reasons why partially connected meshes may be needed even in a tunneled environment are discussed in section 5.1.1.

5.3.4.5 Piggybacking in IP Backbone Routing Protocols

As with VPRN membership, the set of address prefixes associated with each stub interface could also be piggybacked into the routing advertisements from each edge router and propagated through the network. Other edge routers extract this information from received route advertisements in the same way as they obtain the VPRN membership information (which, in this case, is implicit in the identification of the source of each route advertisement). Note that this scheme may require, depending upon the nature of the routing protocols involved, that intermediate routers, e.g. border routers, cache intra-VPRN routing information in order to propagate it further. This also has implications for the trust model, and for the level of security possible for intra-VPRN routing information.

Note that in any of the cases discussed above, an edge router has the option of disseminating its stub link prefixes in a manner so as to permit tunneling from remote edge routers directly to the egress stub links. Alternatively, it could disseminate the information so as to associate all such prefixes with the edge router, rather than with specific stub links. In this case, the edge router would need to implement a VPN specific forwarding mechanism for egress traffic, to determine the correct egress stub link. The advantage of this is that it may significantly reduce the number of distinct tunnels or tunnel label information which need to be constructed and maintained. Note that this choice is purely a local manner and is not visible to remote edge routers.

5.3.5 Tunneling Mechanisms

Once VPRN membership information has been disseminated, the tunnels comprising the VPRN core can be constructed.

One approach to setting up the tunnel mesh is to use point-to-point IP tunnels, and the requirements and issues for such tunnels have been discussed in section 3.0. For example while tunnel establishment can be done through manual configuration, this is

clearly not likely to be a scalable solution, given the $O(n^2)$ problem of meshed links. As such, tunnel set up should use some form of signalling protocol to allow two nodes to construct a tunnel to each other knowing only each other's identity.

Another approach is to use the multipoint to point 'tunnels' provided by MPLS. As noted in [38], MPLS can be considered to be a form of IP tunneling, since the labels of MPLS packets allow for routing decisions to be decoupled from the addressing information of the packets themselves. MPLS label distribution mechanisms can be used to associate specific sets of MPLS labels with particular VPRN address prefixes supported on particular egress points (i.e., stub links of edge routers) and hence allow other edge routers to explicitly label and route traffic to particular VPRN stub links.

One attraction of MPLS as a tunneling mechanism is that it may require less processing within each edge router than alternative tunneling mechanisms. This is a function of the fact that data security within a MPLS network is implicit in the explicit label binding, much as with a connection oriented network, such as Frame Relay. This may hence lessen customer concerns about data security and hence require less processor intensive security mechanisms (e.g., IPSec). However there are other potential security concerns with MPLS. There is no direct support for security features such as authentication, confidentiality, and non-repudiation and the trust model for MPLS means that intermediate routers, (which may belong to different administrative domains), through which membership and prefix reachability information is conveyed, must be trusted, not just the edge routers themselves.

5.4 Multihomed Stub Routers

The discussion thus far has implicitly assumed that stub routers are connected to one and only one VPRN edge router. In general, this restriction should be capable of being relaxed without any change to VPRN operation, given general market interest in multihoming for reliability and other reasons. In particular, in cases where the stub router supports multiple redundant links, with only one operational at any given time, with the links connected either to the same VPRN edge router, or to two or more different VPRN edge routers, then the stub link reachability mechanisms will both discover the loss of an active link, and the activation of a backup link. In the former situation, the previously connected VPRN edge router will cease advertising reachability to the stub node, while the VPRN edge router with the now active link will begin advertising reachability, hence restoring connectivity.

An alternative scenario is where the stub node supports multiple active links, using some form of load sharing algorithm. In such a case, multiple VPRN edge routers may have active paths to the stub node, and may so advertise across the VPRN. This scenario should not cause any problem with reachability across the VPRN providing that the intra-VPRN reachability mechanism can accommodate multiple paths to the same prefix, and has the appropriate mechanisms to preclude looping - for instance, distance vector metrics associated with each advertised prefix.

5.5 Multicast Support

Multicast and broadcast traffic can be supported across VPRNs either by edge replication or by native multicast support in the backbone. These two cases are discussed below.

5.5.1 Edge Replication

This is where each VPRN edge router replicates multicast traffic for transmission across each link in the VPRN. Note that this is the same operation that would be performed by CPE routers terminating actual physical links or dedicated connections. As with CPE routers, multicast routing protocols could also be run on each VPRN edge router to determine the distribution tree for multicast traffic and hence reduce unnecessary flood traffic. This could be done by running instances of standard multicast routing protocols, e.g. Protocol Independent Multicast (PIM) [39] or Distance Vector Multicast Routing Protocol (DVMRP) [40], on and between each VPRN edge router, through the VPRN tunnels, in the same way that unicast routing protocols might be run at each VPRN edge router to determine intra-VPN unicast reachability, as discussed in section 5.3.4. Alternatively, if a link reachability protocol was run across the VPRN tunnels for intra-VPRN reachability, then this could also be augmented to allow VPRN edge routers to indicate both the particular multicast groups requested for reception at each edge node, and also the multicast sources at each edge site.

In either case, there would need to be some mechanism to allow for the VPRN edge routers to determine which particular multicast groups were requested at each site and which sources were present at each site. How this could be done would, in general, be a function of the capabilities of the CPE stub routers at each site. If these run multicast routing protocols, then they can interact directly with the equivalent protocols at each VPRN edge router. If the CPE device does not run a multicast routing protocol, then in the absence of Internet Group Management Protocol (IGMP) proxying [41] the customer site would be limited to a single subnet connected to the VPRN edge router via a bridging device, as the scope of an IGMP message is

limited to a single subnet. However using IGMP-proxying the CPE router can engage in multicast forwarding without running a multicast routing protocol, in constrained topologies. On its interfaces into the customer site the CPE router performs the router functions of IGMP, and on its interface to the VPRN edge router it performs the host functions of IGMP.

5.5.2 Native Multicast Support

This is where VPRN edge routers map intra-VPRN multicast traffic onto a native IP multicast distribution mechanism across the backbone. Note that intra-VPRN multicast has the same requirements for isolation from general backbone traffic as intra-VPRN unicast traffic. Currently the only IP tunneling mechanism that has native support for multicast is MPLS. On the other hand, while MPLS supports native transport of IP multicast packets, additional mechanisms would be needed to leverage these mechanisms for the support of intra-VPRN multicast.

For instance, each VPRN router could prefix multicast group addresses within each VPRN with the VPN-ID of that VPRN and then redistribute these, essentially treating this VPN-ID/intra-VPRN multicast address tuple as a normal multicast address, within the backbone multicast routing protocols, as with the case of unicast reachability, as discussed previously. The MPLS multicast label distribution mechanisms could then be used to set up the appropriate multicast LSPs to interconnect those sites within each VPRN supporting particular multicast group addresses. Note, however, that this would require each of the intermediate LSRs to not only be aware of each intra-VPRN multicast group, but also to have the capability of interpreting these modified advertisements. Alternatively, mechanisms could be defined to map intra-VPRN multicast groups into backbone multicast groups.

Other IP tunneling mechanisms do not have native multicast support. It may prove feasible to extend such tunneling mechanisms by allocating IP multicast group addresses to the VPRN as a whole and hence distributing intra-VPRN multicast traffic encapsulated within backbone multicast packets. Edge VPRN routers could filter out unwanted multicast groups. Alternatively, mechanisms could also be defined to allow for allocation of backbone multicast group addresses for particular intra-VPRN multicast groups, and to then utilize these, through backbone multicast protocols, as discussed above, to limit forwarding of intra-VPRN multicast traffic only to those nodes within the group.

A particular issue with the use of native multicast support is the provision of security for such multicast traffic. Unlike the case of edge replication, which inherits the security characteristics of the underlying tunnel, native multicast mechanisms will need to use some form of secure multicast mechanism. The development of architectures and solutions for secure multicast is an active research area, for example see [42] and [43]. The Secure Multicast Group (SMuG) of the IRTF has been set up to develop prototype solutions, which would then be passed to the IETF IPsec working group for standardization.

However considerably more development is needed before scalable secure native multicast mechanisms can be generally deployed.

5.6 Recommendations

The various proposals that have been developed to support some form of VPRN functionality can be broadly classified into two groups - those that utilize the router piggybacking approach for distributing VPN membership and/or reachability information ([13],[15]) and those that use the virtual routing approach ([12],[14]). In some cases the mechanisms described rely on the characteristics of a particular infrastructure (e.g. MPLS) rather than just IP.

Within the context of the virtual routing approach it may be useful to develop a membership distribution protocol based on a directory or MIB. When combined with the protocol extensions for IP tunneling protocols outlined in section 3.2, this would then provide the basis for a complete set of protocols and mechanisms that support interoperable VPRNs that span multiple administrations over an IP backbone. Note that the other major pieces of functionality needed - the learning and distribution of customer reachability information, can be performed by instances of standard routing protocols, without the need for any protocol extensions.

Also for the constrained case of a full mesh topology, the usefulness of developing a link reachability protocol could be examined, however the limitations and scalability issues associated with this topology may not make it worthwhile to develop something specific for this case, as standard routing will just work.

Extending routing protocols to allow a VPN-ID to be carried in routing update packets could also be examined, but is not necessary if VPN specific tunnels are used.

6.0 VPN Types: Virtual Private Dial Networks

A Virtual Private Dial Network (VPDN) allows for a remote user to connect on demand through an ad hoc tunnel into another site. The user is connected to a public IP network via a dial-up PSTN or ISDN link, and user packets are tunneled across the public network to the desired site, giving the impression to the user of being 'directly' connected into that site. A key characteristic of such ad hoc connections is the need for user authentication as a prime requirement, since anyone could potentially attempt to gain access to such a site using a switched dial network.

Today many corporate networks allow access to remote users through dial connections made through the PSTN, with users setting up PPP connections across an access network to a network access server, at which point the PPP sessions are authenticated using AAA systems running such standard protocols as Radius [44]. Given the pervasive deployment of such systems, any VPDN system must in practice allow for the near transparent re-use of such existing systems.

The IETF have developed the Layer 2 Tunneling Protocol (L2TP) [8] which allows for the extension of of user PPP sessions from an L2TP Access Concentrator (LAC) to a remote L2TP Network Server (LNS). The L2TP protocol itself was based on two earlier protocols, the Layer 2 Forwarding protocol (L2F) [45], and the Point-to-Point Tunneling Protocol (PPTP) [46], and this is reflected in the two quite different scenarios for which L2TP can be used - compulsory tunneling and voluntary tunneling, discussed further below in sections 6.2 and 6.3.

This document focuses on the use of L2TP over an IP network (using UDP), but L2TP may also be run directly over other protocols such as ATM or Frame Relay. Issues specifically related to running L2TP over non-IP networks, such as how to secure such tunnels, are not addressed here.

6.1 L2TP protocol characteristics

This section looks at the characteristics of the L2TP tunneling protocol using the categories outlined in section 3.0.

6.1.1 Multiplexing

L2TP has inherent support for the multiplexing of multiple calls from different users over a single link. Between the same two IP endpoints, there can be multiple L2TP tunnels, as identified by a tunnel-id, and multiple sessions within a tunnel, as identified by a session-id.

6.1.2 Signalling

This is supported via the inbuilt control connection protocol, allowing both tunnels and sessions to be established dynamically.

6.1.3 Data Security

By allowing for the transparent extension of PPP from the user, through the LAC to the LNS, L2TP allows for the use of whatever security mechanisms, with respect to both connection set up, and data transfer, may be used with normal PPP connections. However this does not provide security for the L2TP control protocol itself. In this case L2TP could be further secured by running it in combination with IPSec through IP backbones [47], [48], or related mechanisms on non-IP backbones [49].

The interaction of L2TP with AAA systems for user authentication and authorization is a function of the specific means by which L2TP is used, and the nature of the devices supporting the LAC and the LNS. These issues are discussed in depth in [50].

The means by which the host determines the correct LAC to connect to, and the means by which the LAC determines which users to further tunnel, and the LNS parameters associated with each user, are outside the scope of the operation of a VPDN, but may be addressed, for instance, by evolving Internet roaming specifications [51].

6.1.4 Multiprotocol Transport

L2TP transports PPP packets (and only PPP packets) and thus can be used to carry multiprotocol traffic since PPP itself is multiprotocol.

6.1.5 Sequencing

L2TP supports sequenced delivery of packets. This is a capability that can be negotiated at session establishment, and that can be turned on and off by an LNS during a session. The sequence number field in L2TP can also be used to provide an indication of dropped packets, which is needed by various PPP compression algorithms to operate correctly. If no compression is in use, and the LNS determines that the protocols in use (as evidenced by the PPP NCP negotiations) can deal with out of sequence packets (e.g. IP), then it may disable the use of sequencing.

6.1.6 Tunnel Maintenance

A keepalive protocol is used by L2TP in order to allow it to distinguish between a tunnel outage and prolonged periods of tunnel inactivity.

6.1.7 Large MTUs

L2TP itself has no inbuilt support for a segmentation and reassembly capability, but when run over UDP/IP IP fragmentation will take place if necessary. Note that a LAC or LNS may adjust the Maximum Receive Unit (MRU) negotiated via PPP in order to preclude fragmentation, if it has knowledge of the MTU used on the path between LAC and LNS. To this end, there is a proposal to allow the use of MTU discovery for cases where the L2TP tunnel transports IP frames [52].

6.1.8 Tunnel Overhead

L2TP as used over IP networks runs over UDP and must be used to carry PPP traffic. This results in a significant amount of overhead, both in the data plane with UDP, L2TP and PPP headers, and also in the control plane, with the L2TP and PPP control protocols. This is discussed further in section 6.3

6.1.9 Flow and Congestion Control

L2TP supports flow and congestion control mechanisms for the control protocol, but not for data traffic. See section 3.1.9 for more details.

6.1.10 QoS / Traffic Management

An L2TP header contains a 1-bit priority field, which can be set for packets that may need preferential treatment (e.g. keepalives) during local queuing and transmission. Also by transparently extending PPP, L2TP has inherent support for such PPP mechanisms as multi-link PPP [53] and its associated control protocols [54], which allow for bandwidth on demand to meet user requirements.

In addition L2TP calls can be mapped into whatever underlying traffic management mechanisms may exist in the network, and there are proposals to allow for requests through L2TP signalling for specific differentiated services behaviors [55].

6.1.11 Miscellaneous

Since L2TP is designed to transparently extend PPP, it does not attempt to supplant the normal address assignment mechanisms associated with PPP. Hence, in general terms the host initiating the PPP session will be assigned an address by the LNS using PPP procedures. This addressing may have no relation to the addressing used for communication between the LAC and LNS. The LNS will also need to support whatever forwarding mechanisms are needed to route traffic to and from the remote host.

6.2 Compulsory Tunneling

Compulsory tunneling refers to the scenario in which a network node - a dial or network access server, for instance - acting as a LAC, extends a PPP session across a backbone using L2TP to a remote LNS, as illustrated below. This operation is transparent to the user initiating the PPP session to the LAC. This allows for the decoupling of the location and/or ownership of the modem pools used to terminate dial calls, from the site to which users are provided access. Support for this scenario was the original intent of the L2F specification, upon which the L2TP specification was based.

There are a number of different deployment scenarios possible. One example, shown in the diagram below, is where a subscriber host dials into a NAS acting as a LAC, and is tunneled across an IP network (e.g. the Internet) to a gateway acting as an LNS. The gateway provides access to a corporate network, and could either be a device in the corporate network itself, or could be an ISP edge router, in the case where a customer has outsourced the maintenance of LNS functionality to an ISP. Another scenario is where an ISP uses L2TP to provide a subscriber with access to the Internet. The subscriber host dials into a NAS acting as a LAC, and is tunneled across an access network to an ISP edge router acting as an LNS. This ISP edge router then feeds the subscriber traffic into the Internet. Yet other scenarios are where an ISP uses L2TP to provide a subscriber with access to a VPRN, or with concurrent access to both a VPRN and the Internet.

A VPDN, whether using compulsory or voluntary tunneling, can be viewed as just another type of access method for subscriber traffic, and as such can be used to provide connectivity to different types of networks, e.g. a corporate network, the Internet, or a VPRN. The last scenario is also an example of how a VPN service as provided to a customer may be implemented using a combination of different types of VPN.

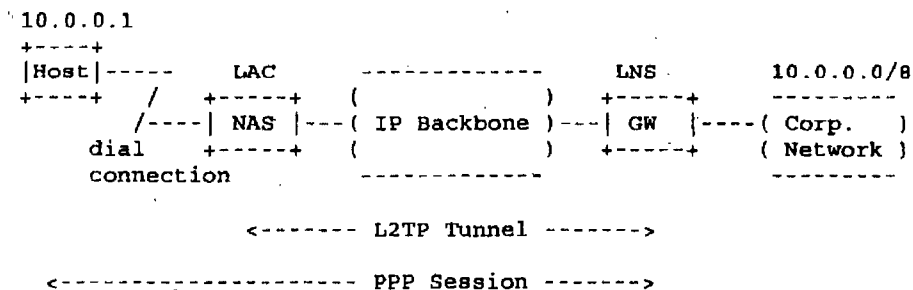


Figure 6.1: Compulsory Tunneling Example

Compulsory tunneling was originally intended for deployment on network access servers supporting wholesale dial services, allowing for remote dial access through common facilities to an enterprise site, while precluding the need for the enterprise to deploy its own dial servers. Another example of this is where an ISP outsources its own dial connectivity to an access network provider (such as a Local Exchange Carrier (LEC) in the USA) removing the need for an ISP to maintain its own dial servers and allowing the LEC to serve multiple ISPs. More recently, compulsory tunneling mechanisms have also been proposed for evolving Digital Subscriber Line (DSL) services [56], [57], which also seek to leverage the existing AAA infrastructure.

Call routing for compulsory tunnels requires that some aspect of the initial PPP call set up can be used to allow the LAC to determine the identity of the LNS. As noted in [50], these aspects can include the user identity, as determined through some aspect of the access network, including calling party number, or some attribute of the called party, such as the Fully Qualified Domain Name (FQDN) of the identity claimed during PPP authentication.

It is also possible to chain two L2TP tunnels together, whereby a LAC initiates a tunnel to an intermediate relay device, which acts as an LNS to this first LAC, and acts as a LAC to the final LNS. This may be needed in some cases due to administrative, organizational or regulatory issues pertaining to the split between access network provider, IP backbone provider and enterprise customer.

6.3 Voluntary Tunnels

Voluntary tunneling refers to the case where an individual host connects to a remote site using a tunnel originating on the host, with no involvement from intermediate network nodes, as illustrated below. The PPTP specification, parts of which have been incorporated into L2TP, was based upon a voluntary tunneling model.

As with compulsory tunneling there are different deployment scenarios possible. The diagram below shows a subscriber host accessing a corporate network with either L2TP or IPsec being used as the voluntary tunneling mechanism. Another scenario is where voluntary tunneling is used to provide a subscriber with access to a VPRN.

6.3.1 Issues with Use of L2TP for Voluntary Tunnels

The L2TP specification has support for voluntary tunneling, insofar as the LAC can be located on a host, not only on a network node. Note that such a host has two IP addresses - one for the LAC-LNS IP tunnel, and another, typically allocated via PPP, for the network to which the host is connecting. The benefits of using L2TP for voluntary tunneling are that the existing authentication and address assignment mechanisms used by PPP can be reused without modification. For example an LNS could also include a Radius client, and communicate with a Radius server to authenticate a PPP PAP or CHAP exchange, and to retrieve configuration information for the host such as its IP address and a list of DNS servers to use. This information can then be passed to the host via the PPP IPCP protocol.




Figure 6.2: Voluntary Tunneling Example

The above procedure is not without its costs, however. There is considerable overhead with such a protocol stack, particularly when IPsec is also needed for security purposes, and given that the host may be connected via a low-bandwidth dial up link. The overhead consists of both extra headers in the data plane and extra control protocols needed in the control plane. Using L2TP for voluntary tunneling, secured with IPsec, means a web application, for example, would run over the following stack

HTTP/TCP/IP/PPP/L2TP/UDP/ESP/IP/PPP/AHDLC

It is proposed in [58] that IPsec alone be used for voluntary tunnels reducing overhead, using the following stack.

HTTP/TCP/IP/ESP/IP/PPP/AHDLC

In this case IPsec is used in tunnel mode, with the tunnel terminating either on an IPsec edge device at the enterprise site, or on the provider edge router connected to the enterprise site. There are two possibilities for the IP addressing of the host. Two IP addresses could be used, in a similar manner to the L2TP case. Alternatively the host can use a single public IP address as the source IP address in both inner and outer IP headers, with the gateway performing Network Address Translation (NAT) before forwarding the traffic to the enterprise network. To other hosts in the enterprise network the host appears to have an 'internal' IP address. Using NAT has some limitations and restrictions, also pointed out in [58].

Another area of potential problems with PPP is due to the fact that the characteristics of a link layer implemented via an L2TP tunnel over an IP backbone are quite different to a link layer run over a serial line, as discussed in the L2TP specification itself. For example, poorly chosen PPP parameters may lead to frequent resets and timeouts, particularly if compression is in use. This is because an L2TP tunnel may misorder packets, and may silently drop packets, neither of which normally occurs on serial lines. The general packet loss rate could also be significantly higher due to network congestion. Using the sequence number field in an L2TP header addresses the misordering issue, and for cases where the LAC and LNS are coincident with the PPP endpoints, as in voluntary tunneling, the sequence number field can also be used to detect a dropped packet, and to pass a suitable indication to any compression entity in use, which typically requires such knowledge in order to keep the compression histories in synchronization at both ends. (In fact this is more of an issue with compulsory tunneling since the LAC may have to deliberately issue a corrupted frame to the PPP host, to give an indication of packet loss, and some hardware may not allow this).

6.3.2 Issues with Use of IPSec for Voluntary Tunnels

If IPSec is used for voluntary tunneling, the functions of user authentication and host configuration, achieved by means of PPP when using L2TP, still need to be carried out. A distinction needs to be drawn here between machine authentication and user authentication. Two factor authentication is carried out on the basis of both something the user has, such as a machine or smartcard with a digital certificate, and something the user knows, such as a password. (Another example is getting money from an bank ATM machine - you need a card and a PIN number). Many of the existing legacy schemes currently in use to perform user authentication are asymmetric in nature, and are not supported by IKE. For remote access the most common existing user authentication mechanism is to use PPP between the user and access server, and Radius between the access server and authentication server. The authentication exchanges that occur in this case, e.g. a PAP or CHAP exchange, are asymmetric. Also CHAP supports the ability for the network to reauthenticate the user at any time after the initial session has been established, to ensure that the current user is the same person that initiated the session.

While IKE provides strong support for machine authentication, it has only limited support for any form of user authentication and has no support for asymmetric user authentication. While a user password can be used to derive a key used as a preshared key, this cannot be used with IKE Main Mode in a remote access environment, as the user will not have a fixed IP address, and while Aggressive Mode can be used instead, this affords no identity protection. To this end there have been a number of proposals to allow for support of legacy asymmetric user level authentication schemes with IPSec. [59] defines a new IKE message exchange - the transaction exchange - which allows for both Request/Reply and Set/Acknowledge message sequences, and it also defines attributes that can be used for client IP stack configuration. [60] and [61] describe mechanisms that use the transaction message exchange, or a series of such exchanges, carried out between the IKE Phase 1 and Phase 2 exchanges, to perform user authentication. A different approach, that does not extend the IKE protocol itself, is described in [62]. With this approach a user establishes a Phase 1 SA with a security gateway and then sets up a Phase 2 SA to the gateway, over which an existing authentication protocol is run. The gateway acts as a proxy and relays the protocol messages to an authentication server.

In addition there have also been proposals to allow the remote host to be configured with an IP address and other configuration information over IPSec. For example [63] describes a method whereby a remote host first establishes a Phase 1 SA with a security gateway and then sets up a Phase 2 SA to the gateway, over which the DHCP

protocol is run. The gateway acts as a proxy and relays the protocol messages to the DHCP server. Again, like [62], this proposal does not involve extensions to the IKE protocol itself.

Another aspect of PPP functionality that may need to be supported is multiprotocol operation, as there may be a need to carry network layer protocols other than IP, and even to carry link layer protocols (e.g. ethernet) as would be needed to support bridging over IPsec. This is discussed in section 3.1.4.

The methods of supporting legacy user authentication and host configuration capabilities in a remote access environment are currently being discussed in the IPsec working group.

6.4 Networked Host Support

The current PPP based dial model assumes a host directly connected to a connection oriented dial access network. Recent work on new access technologies such as DSL have attempted to replicate this model [57], so as to allow for the re-use of existing AAA systems. The proliferation of personal computers, printers and other network appliances in homes and small businesses, and the ever lowering costs of networks, however, are increasingly challenging the directly connected host model. Increasingly, most hosts will access the Internet through small, typically Ethernet, local area networks.

There is hence interest in means of accommodating the existing AAA infrastructure within service providers, whilst also supporting multiple networked hosts at each customer site. The principal complication with this scenario is the need to support the login dialogue, through which the appropriate AAA information is exchanged. A number of proposals have been made to address this scenario:

6.4.1 Extension of PPP to Hosts Through L2TP

A number of proposals (e.g. [56]) have been made to extend L2TP over Ethernet so that PPP sessions can run from networked hosts out to the network, in much the same manner as a directly attached host.

6.4.2 Extension of PPP Directly to Hosts:

There is also a specification for mapping PPP directly onto Ethernet (PPPOE) [64] which uses a broadcast mechanism to allow hosts to find appropriate access servers with which to connect. Such servers could then further tunnel, if needed, the PPP sessions using L2TP or a similar mechanism.

6.4.3 Use of IPSec

The IPSec based voluntary tunneling mechanisms discussed above can be used either with networked or directly connected hosts.

Note that all of these methods require additional host software to be used, which implements either LAC, PPPOE client or IPSec client functionality.

6.5 Recommendations

The L2TP specification has been finalized and will be widely used for compulsory tunneling. As discussed in section 3.2, defining specific modes of operation for IPSec when used to secure L2TP would be beneficial.

Also, for voluntary tunneling using IPSec, completing the work needed to provide support for the following areas would be useful

- asymmetric / legacy user authentication (6.3)
- host address assignment and configuration (6.3)

along with any other issues specifically related to the support of remote hosts. Currently as there are many different non-interoperable proprietary solutions in this area.

7.0 VPN Types: Virtual Private LAN Segment

A Virtual Private LAN Segment (VPLS) is the emulation of a LAN segment using Internet facilities. A VPLS can be used to provide what is sometimes known also as a Transparent LAN Service (TLS), which can be used to interconnect multiple stub CPE nodes, either bridges or routers, in a protocol transparent manner. A VPLS emulates a LAN segment over IP, in the same way as protocols such as LANE emulate a LAN segment over ATM. The primary benefits of a VPLS are complete protocol transparency, which may be important both for multiprotocol transport and for regulatory reasons in particular service provider contexts.

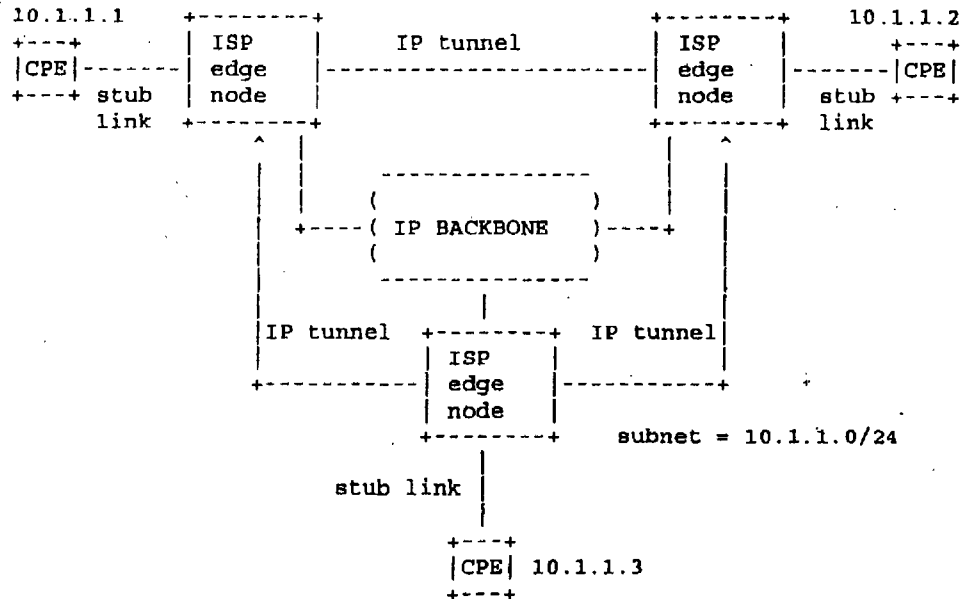


Figure 7.1: VPLS Example

7.1 VPLS Requirements

Topologically and operationally a VPLS can be most easily modeled as being essentially equivalent to a VPRN, except that each VPLS edge node implements link layer bridging rather than network layer forwarding. As such, most of the VPRN tunneling and configuration mechanisms discussed previously can also be used for a VPLS, with the appropriate changes to accommodate link layer, rather than network layer, packets and addressing information. The following sections discuss the primary changes needed in VPRN operation to support VPLSs.

7.1.1 Tunneling Protocols

The tunneling protocols employed within a VPLS can be exactly the same as those used within a VPRN, if the tunneling protocol permits the transport of multiprotocol traffic, and this is assumed below.

7.1.2 Multicast and Broadcast Support

A VPLS needs to have a broadcast capability. This is needed both for broadcast frames, and for link layer packet flooding, where a unicast frame is flooded because the path to the destination link layer address is unknown. The address resolution protocols that run over a bridged network typically use broadcast frames (e.g. ARP). The same set of possible multicast tunneling mechanisms discussed earlier for VPRNs apply also to a VPLS, though the generally more frequent use of broadcast in VPLSs may increase the pressure for native multicast support that reduces, for instance, the burden of replication on VPLS edge nodes.

7.1.3 VPLS Membership Configuration and Topology

The configuration of VPLS membership is analogous to that of VPRNs since this generally requires only knowledge of the local VPN link assignments at any given VPLS edge node, and the identity of, or route to, the other edge nodes in the VPLS; in particular, such configuration is independent of the nature of the forwarding at each VPN edge node. As such, any of the mechanisms for VPN member configuration and dissemination discussed for VPRN configuration can also be applied to VPLS configuration. Also as with VPRNs, the topology of the VPLS could be easily manipulated by controlling the configuration of peer nodes at each VPLS edge node, assuming that the membership dissemination mechanism was such as to permit this. It is likely that typical VPLSs will be fully meshed, however, in order to preclude the need for traffic between two VPLS nodes to transit through another VPLS node, which would then require the use of the Spanning Tree protocol [65] for loop prevention.

7.1.4 CPE Stub Node Types

A VPLS can support either bridges or routers as a CPE device.

CPE routers would peer transparently across a VPLS with each other without requiring any router peering with any nodes within the VPLS. The same scalability issues that apply to a full mesh topology for VPRNs, apply also in this case, only that now the number of peering routers is potentially greater, since the ISP edge device is no longer acting as an aggregation point.

With CPE bridge devices the broadcast domain encompasses all the CPE sites as well as the VPLS itself. There are significant scalability constraints in this case, due to the need for packet flooding, and

the fact that any topology change in the bridged domain is not localized, but is visible throughout the domain. As such this scenario is generally only suited for support of non-routable protocols.

The nature of the CPE impacts the nature of the encapsulation, addressing, forwarding and reachability protocols within the VPLS, and are discussed separately below.

7.1.5 Stub Link Packet Encapsulation

7.1.5.1 Bridge CPE

In this case, packets sent to and from the VPLS across stub links are link layer frames, with a suitable access link encapsulation. The most common case is likely to be ethernet frames, using an encapsulation appropriate to the particular access technology, such as ATM, connecting the CPE bridges to the VPLS edge nodes. Such frames are then forwarded at layer 2 onto a tunnel used in the VPLS. As noted previously, this does mandate the use of an IP tunneling protocol which can transport such link layer frames. Note that this does not necessarily mandate, however, the use of a protocol identification field in each tunnel packet, since the nature of the encapsulated traffic (e.g. ethernet frames) could be indicated at tunnel setup.

7.1.5.2 Router CPE

In this case, typically, CPE routers send link layer packets to and from the VPLS across stub links, destined to the link layer addresses of their peer CPE routers. Other types of encapsulations may also prove feasible in such a case, however, since the relatively constrained addressing space needed for a VPLS to which only router CPE are connected, could allow for alternative encapsulations, as discussed further below.

7.1.6 CPE Addressing and Address Resolution

7.1.6.1 Bridge CPE

Since a VPLS operates at the link layer, all hosts within all stub sites, in the case of bridge CPE, will typically be in the same network layer subnet. (Multinetting, whereby multiple subnets operate over the same LAN segment, is possible, but much less common). Frames are forwarded across and within the VPLS based upon the link layer addresses - e.g. IEEE MAC addresses - associated with the individual hosts. The VPLS needs to support broadcast traffic, such as that typically used for the address resolution mechanism used

to map the host network addresses to their respective link addresses. The VPLS forwarding and reachability algorithms also need to be able to accommodate flooded traffic.

7.1.6.2 Router CPE

A single network layer subnet is generally used to interconnect router CPE devices, across a VPLS. Behind each CPE router are hosts in different network layer subnets. CPE routers transfer packets across the VPLS by mapping next hop network layer addresses to the link layer addresses of a router peer. A link layer encapsulation is used, most commonly ethernet, as for the bridge case.

As noted above, however, in cases where all of the CPE nodes connected to the VPLS are routers, then it may be possible, due to the constrained addressing space of the VPLS, to use encapsulations that use a different address space than normal MAC addressing. See, for instance, [11], for a proposed mechanism for VPLSs over MPLS networks, leveraging earlier work on VPRN support over MPLS [38], which proposes MPLS as the tunneling mechanism, and locally assigned MPLS labels as the link layer addressing scheme to identify the CPE LSR routers connected to the VPLS.

7.1.7 VPLS Edge Node Forwarding and Reachability Mechanisms

7.1.7.1 Bridge CPE

The only practical VPLS edge node forwarding mechanism in this case is likely to be standard link layer packet flooding and MAC address learning, as per [65]. As such, no explicit intra-VPLS reachability protocol will be needed, though there will be a need for broadcast mechanisms to flood traffic, as discussed above. In general, it may not prove necessary to also implement the Spanning Tree protocol between VPLS edge nodes, if the VPLS topology is such that no VPLS edge node is used for transit traffic between any other VPLS edge nodes - in other words, where there is both full mesh connectivity and transit is explicitly precluded. On the other hand, the CPE bridges may well implement the spanning tree protocol in order to safeguard against 'backdoor' paths that bypass connectivity through the VPLS.

7.1.7.2 Router CPE

Standard bridging techniques can also be used in this case. In addition, the smaller link layer address space of such a VPLS may also permit other techniques, with explicit link layer routes between CPE routers. [11], for instance, proposes that MPLS LSPs be set up, at the insertion of any new CPE router into the VPLS, between all CPE

LSRs. This then precludes the need for packet flooding. In the more general case, if stub link reachability mechanisms were used to configure VPLS edge nodes with the link layer addresses of the CPE routers connected to them, then modifications of any of the intra-VPN reachability mechanisms discussed for VPRNs could be used to propagate this information to each other VPLS edge node. This would then allow for packet forwarding across the VPLS without flooding.

Mechanisms could also be developed to further propagate the link layer addresses of peer CPE routers and their corresponding network layer addresses across the stub links to the CPE routers, where such information could be inserted into the CPE router's address resolution tables. This would then also preclude the need for broadcast address resolution protocols across the VPLS.

Clearly there would be no need for the support of spanning tree protocols if explicit link layer routes were determined across the VPLS. If normal flooding mechanisms were used then spanning tree would only be required if full mesh connectivity was not available and hence VPLS nodes had to carry transit traffic.

7.2 Recommendations

There is significant commonality between VPRNs and VPLSs, and, where possible, this similarity should be exploited in order to reduce development and configuration complexity. In particular, VPLSs should utilize the same tunneling and membership configuration mechanisms, with changes only to reflect the specific characteristics of VPLSs.

8.0 Summary of Recommendations

In this document different types of VPNs have been discussed individually, but there are many common requirements and mechanisms that apply to all types of VPNs, and many networks will contain a mix of different types of VPNs. It is useful to have as much commonality as possible across these different VPN types. In particular, by standardizing a relatively small number of mechanisms, it is possible to allow a wide variety of VPNs to be implemented.

The benefits of adding support for the following mechanisms should be carefully examined.

For IKE/IPSec:

- the transport of a VPN-ID when establishing an SA (3.1.2)
- a null encryption and null authentication option (3.1.3)

- multiprotocol operation (3.1.4)
- frame sequencing (3.1.5)
- asymmetric / legacy user authentication (6.3)
- host address assignment and configuration (6.3)

For L2TP:

- defining modes of operation of IPSec when used to support L2TP (3.2)

For VPNs generally:

- defining a VPN membership information configuration and dissemination mechanism, that uses some form of directory or MIB (5.3.2)
- ensure that solutions developed, as far as possible, are applicable to different types of VPNs, rather than being specific to a single type of VPN.

9.0 Security Considerations

Security considerations are an integral part of any VPN mechanisms, and these are discussed in the sections describing those mechanisms.

10.0 Acknowledgements

Thanks to Anthony Alles, of Nortel Networks, for his invaluable assistance with the generation of this document, and who developed much of the material on which early versions of this document were based. Thanks also to Joel Halpern for his helpful review comments.

11.0 References

- [1] ATM Forum. "LAN Emulation over ATM 1.0", af-lanE-0021.000, January 1995.
- [2] ATM Forum. "Multi-Protocol Over ATM Specification v1.0", af-mpoa-0087.000, June 1997.
- [3] Ferguson, P. and Huston, G. "What is a VPN?", Revision 1, April 1 1998; <http://www.employees.org/~ferguson/vpn.pdf>.

- [4] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [5] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [6] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.
- [7] Hanks, S., Li, T., Farinacci, D. and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 1701, October 1994.
- [8] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G. and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.
- [9] Rosen, E., et al., "Multiprotocol Label Switching Architecture", Work in Progress.
- [10] Heinanen, J., et al., "MPLS Mappings of Generic VPN Mechanisms", Work in Progress.
- [11] Jamieson, D., et al., "MPLS VPN Architecture", Work in Progress.
- [12] Casey, L., et al., "IP VPN Realization using MPLS Tunnels", Work in Progress.
- [13] Li, T. "CPE based VPNs using MPLS", Work in Progress.
- [14] Muthukrishnan, K. and A. Malis, "Core MPLS IP VPN Architecture", Work in Progress.
- [15] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March 1999.
- [16] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", RFC 2685, September 1999.
- [17] Petri, B. (editor) "MPOA v1.1 Addendum on VPN support", ATM Forum, af-mpoa-0129.000.
- [18] Harkins, D. and C. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.
- [19] Calhoun, P., et al., "Tunnel Establishment Protocol", Work in Progress.

- [20] Andersson, L., et al., "LDP Specification", Work in Progress.
- [21] Jamoussi, B., et al., "Constraint-Based LSP Setup using LDP" Work in Progress.
- [22] Awduche, D., et al., "Extensions to RSVP for LSP Tunnels", Work in Progress.
- [23] Kent, S. and R. Atkinson, "IP Encapsulating Security Protocol (ESP)", RFC 2406, November 1998.
- [24] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [25] Perez, M., Liaw, F., Mankin, A., Hoffman, E., Grossman, D. and A. Malis, "ATM Signalling Support for IP over ATM", RFC 1755, February 1995.
- [26] Malkin, G. "RIP Version 2 Carrying Additional Information", RFC 1723, November 1994.
- [27] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [28] Shacham, A., Monsour, R., Pereira, R. and M. Thomas, "IP Payload Compression Protocol (IPComp)", RFC 2393, December 1998.
- [29] Duffield N., et al., "A Performance Oriented Service Interface for Virtual Private Networks", Work in Progress.
- [30] Jacobson, V., Nichols, K. and B. Poduri, "An Expedited Forwarding PHB", RFC 2598, June 1999.
- [31] Casey, L., "An extended IP VPN Architecture", Work in Progress.
- [32] Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [33] Grossman, D. and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 2684, September 1999.
- [34] Wahl, M., Howes, T. and S. Kille, "Lightweight Directory Access Protocol (v3)", RFC 2251, December 1997.
- [35] Boyle, J., et al., "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.
- [36] MacRae, M. and S. Ayandeh, "Using COPS for VPN Connectivity" Work in Progress.

- [37] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [38] Heinanen, J. and E. Rosen, "VPN Support with MPLS", Work in Progress.
- [39] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma, P. and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification", RFC 2362, June 1998.
- [40] Waitzman, D., Partridge, C., and S. Deering, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988.
- [41] Fenner, W., "IGMP-based Multicast Forwarding (IGMP Proxying)", Work in Progress.
- [42] Wallner, D., Harder, E. and R. Agee, "Key Management for Multicast: Issues and Architectures", RFC 2627, June 1999.
- [43] Hardjono, T., et al., "Secure IP Multicast: Problem areas, Framework, and Building Blocks", Work in Progress.
- [44] Rigney, C., Rubens, A., Simpson, W. and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2138, April 1997.
- [45] Valencia, A., Littlewood, M. and T. Kolar, "Cisco Layer Two Forwarding (Protocol) "L2F"", RFC 2341, May 1998.
- [46] Hamzeh, K., Pall, G., Verthein, W., Taarud, J., Little, W. and G. Zorn, "Point-to-Point Tunneling Protocol (PPTP)", RFC 2637, July 1999.
- [47] Patel, B., et al., "Securing L2TP using IPSEC", Work in Progress.
- [48] Srisuresh, P., "Secure Remote Access with L2TP", Work in Progress.
- [49] Calhoun, P., et al., "Layer Two Tunneling Protocol "L2TP" Security Extensions for Non-IP networks", Work in Progress.
- [50] Aboba, B. and Zorn, G. "Implementation of PPTP/L2TP Compulsory Tunneling via RADIUS", Work in progress.
- [51] Aboba, B. and G. Zorn, "Criteria for Evaluating Roaming Protocols", RFC 2477, January 1999.

- [52] Shea, R., "L2TP-over-IP Path MTU Discovery (L2TPMTU)", Work in Progress.
- [53] Sklower, K., Lloyd, B., McGregor, G., Carr, D. and T. Coradetti, "The PPP Multilink Protocol (MP)", RFC 1990, August 1996.
- [54] Richards, C. and K. Smith, "The PPP Bandwidth Allocation Protocol (BAP) The PPP Bandwidth Allocation Control Protocol (BACP)", RFC 2125, March 1997.
- [55] Calhoun, P. and K. Peirce, "Layer Two Tunneling Protocol "L2TP" IP Differential Services Extension", Work in Progress.
- [56] ADSL Forum. "An Interoperable End-to-end Broadband Service Architecture over ADSL Systems (Version 3.0)", ADSL Forum 97-215.
- [57] ADSL Forum. "Core Network Architectures for ADSL Access Systems (Version 1.01)", ADSL Forum 98-017.
- [58] Gupta, V., "Secure, Remote Access over the Internet using IPsec", Work in Progress.
- [59] Pereira, R., et al., "The ISAKMP Configuration Method", Work in Progress.
- [60] Pereira, R. and S. Beaulieu, "Extended Authentication Within ISAKMP/Oakley", Work in Progress.
- [61] Litvin, M., et al., "A Hybrid Authentication Mode for IKE", Work in Progress.
- [62] Kelly, S., et al., "User-level Authentication Mechanisms for IPsec", Work in Progress.
- [63] Patel, B., et al., "DHCP Configuration of IPSEC Tunnel Mode", Work in Progress.
- [64] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D. and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, February 1999.
- [65] ANSI/IEEE - 10038: 1993 (ISO/IEC) Information technology - Telecommunications and information exchange between systems - Local area networks - Media access control (MAC) bridges, ANSI/IEEE Std 802.1D, 1993 Edition.

12.0 Author Information

Bryan Gleeson
Nortel Networks
4500 Great America Parkway
Santa Clara CA 95054
USA

Phone: +1 (408) 548 3711
EMail: bgleeson@shastanets.com

Juha Heinanen
Telia Finland, Inc.
Myrmaentie 2
01600 VANTAA
Finland

Phone: +358 303 944 808
EMail: jh@telia.fi

Arthur Lin
Nortel Networks
4500 Great America Parkway
Santa Clara CA 95054
USA

Phone: +1 (408) 548 3788
EMail: alin@shastanets.com

Grenville Armitage
Bell Labs Research Silicon Valley
Lucent Technologies
3180 Porter Drive,
Palo Alto, CA 94304
USA

EMail: gja@lucent.com

Andrew G. Malis
Lucent Technologies
1 Robbins Road
Westford, MA 01886
USA

Phone: +1 978 952 7414
EMail: amalis@lucent.com

13.0 Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

3



Commissioner for Patents
Washington, DC 20231
www.uspto.gov

APPLICATION NUMBER	FILING/RECEIPT DATE	FIRST NAMED APPLICANT	ATTORNEY DOCKET NUMBER
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A

CONFIRMATION NO. 3645

23484
JOHN W L OGILVIE
COMPUTER LAW
1211 EAST YALE AVE
SALT LAKE CITY, UT 84105

FORMALITIES LETTER



OC00000009751373

Date Mailed: 04/02/2003

NOTICE TO FILE MISSING PARTS OF NONPROVISIONAL APPLICATION

FILED UNDER 37 CFR 1.53(b)

Filing Date Granted

Items Required To Avoid Abandonment:

An application number and filing date have been accorded to this application. The item(s) indicated below, however, are missing. Applicant is given **TWO MONTHS** from the date of this Notice within which to file all required items and pay any fees required below to avoid abandonment. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

- The statutory basic filing fee is missing.
Applicant must submit \$ 375 to complete the basic filing fee for a small entity.
- To avoid abandonment, a late filing fee or oath or declaration surcharge as set forth in 37 CFR 1.16(e) of \$65 for a small entity in compliance with 37 CFR 1.27, must be submitted with the missing items identified in this letter.

Items Required To Avoid Processing Delays:

The item(s) indicated below are also required and should be submitted with any reply to this notice to avoid further processing delays.

- Additional claim fees of \$387 as a small entity, including any required multiple dependent claim fee, are required. Applicant must submit the additional claim fees or cancel the additional claims for which fees are due.


SUMMARY OF FEES DUE:

Total additional fee(s) required for this application is \$827 for a Small Entity


- \$375 Statutory basic filing fee.
- \$65 Late oath or declaration Surcharge.
- Total additional claim fee(s) for this application is \$387
 - \$135 for 15 total claims over 20 .

- \$252 for 6 independent claims over 3 .

A copy of this notice MUST be returned with the reply.


Customer Service Center
Initial Patent Examination Division (703) 308-1202

PART 3 - OFFICE COPY



Commissioner for Patents
Washington, DC 20231
www.uspto.gov

APPLICATION NUMBER	FILING/RECEIPT DATE	FIRST NAMED APPLICANT	ATTORNEY DOCKET NUMBER
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A

23484
 JOHN W/L OGILVIE
 COMPUTER LAW
 1211 EAST YALE AVE
 SALT LAKE CITY, UT 84105

CONFIRMATION NO. 3645

FORMALITIES LETTER



OC00000009751373

Date Mailed: 04/02/2003

NOTICE TO FILE MISSING PARTS OF NONPROVISIONAL APPLICATION

04/16/2003 HNEKONEN 00000017 10361837

FILED UNDER 37 CFR 1.53(b)

01 FC:2001	375.00 DP
02 FC:2051	65.00 DP
03 FC:2202	135.00 DP
04 FC:2201	252.00 DP

Filing Date Granted

Items Required To Avoid Abandonment:

An application number and filing date have been accorded to this application. The item(s) indicated below, however, are missing. Applicant is given **TWO MONTHS** from the date of this Notice within which to file all required items and pay any fees required below to avoid abandonment. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

- The statutory basic filing fee is missing.
Applicant must submit \$ 375 to complete the basic filing fee for a small entity.
- To avoid abandonment, a late filing fee or oath or declaration surcharge as set forth in 37 CFR 1.16(e) of \$65 for a small entity in compliance with 37 CFR 1.27, must be submitted with the missing items identified in this letter.

Items Required To Avoid Processing Delays:

The item(s) indicated below are also required and should be submitted with any reply to this notice to avoid further processing delays.

- Additional claim fees of \$387 as a small entity, including any required multiple dependent claim fee, are required. Applicant must submit the additional claim fees or cancel the additional claims for which fees are due.

SUMMARY OF FEES DUE:

Total additional fee(s) required for this application is \$827 for a Small Entity

- \$375 Statutory basic filing fee.
- \$65 Late oath or declaration Surcharge.
- Total additional claim fee(s) for this application is \$387
 - \$135 for 15 total claims over 20 .

- \$252 for 6 independent claims over 3.

A copy of this notice MUST be returned with the reply.

Wilson
Customer Service Center
Initial Patent Examination Division (703) 308-1202

PART 2 - COPY TO BE RETURNED WITH RESPONSE

CERTIFICATE OF MAILING

I hereby certify that this Notice, check #3275 for \$827⁰⁰, and postcard are being deposited on April 11, 2003 as U.S. First Class Mail, postage paid, in an envelope addressed to: Commissioner of Patents, Box Missing Parts, P.O. Box 2327, Arlington, VA 22202.

John W. L. Gilvie
John W. L. Gilvie



263

PATENT APPLICATION
Docket No.: 3003.2.11A

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Sanchaita Datta and Ragula Bhaskar
Serial No.: 10/361837
Filed: February 7, 2003
For: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS
OVER DISPARATE NETWORKS

RECEIVED

SECOND INFORMATION DISCLOSURE STATEMENT

JUN 09 2003

Commissioner for Patents:

Technology Center 2600

This Information Disclosure Statement is filed in response to the duty of candor described in 37 C.F.R. §§ 1.56, 1.98, MPEP § 2001.06(c), and elsewhere. The references listed on the enclosed Form PTO-1449 (incorporated herein by reference) are respectfully submitted for consideration by the Office. They were first identified to the undersigned in an International Search Report in PCT/US03/03988 (Docket 3003.2.11B) received on June 2, 2003.

Dated June 3, 2003.

CERTIFICATE OF MAILING

I hereby certify that the correspondence listed below is being deposited with the United States Postal Service as Priority Mail, postage paid, on June 3, 2003 addressed to the Commissioner for Patents, Mail Stop Non-Fee Amendment, P.O. Box 1450, Alexandria, VA 22313-1450:

Postcard
Second IDS w/ PTO-1449 and 4 references

Respectfully submitted,

JOHN W.L. OGILVIE
Registration No. 37,987

COMPUTER LAW++
1211 East Yale Avenue
Salt Lake City, Utah 84105
801-582-2724 (voice)
801-583-1984 (fax)

p-ids-1

Form PTO-1449

Sheet 1 of 1

Applicant: Sanchaita Datta and Ragula Bhaskar

Serial No.: 10/361837

Att'y Docket No. 3003.2.11A

Filing Date: February 7, 2003

For: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS
OVER DISPARATE NETWORKS

LIST OF REFERENCES CITED BY APPLICANT

U.S. Patent Documents

Examiner Initial*	Document Number	Issue Date	Name	Class	Sub Class	Filing Date	
<u>mm</u>	A1	6,456,594	09/24/02	Kaplan et al.	370	238	07/24/00
<u>mm</u>	A2	6,449,259	09/10/02	Allain et al.	370	253	06/30/97
<u>mm</u>	A3	5,898,673	04/27/99	Riggan et al.	370	237	02/12/97

Other Document

(including Author, Title, Pertinent Pages, etc.)

<u>mm</u>	A4	B. Gleeson et al., "A Framework for IP Based Virtual Private Networks," RFC 2764 (February 2000)				
-----------	----	--	--	--	--	--

Examiner: Melvin Marcelo

Date Considered: 05-24-2004

*EXAMINER: Please initial if reference considered, whether or not citation is in conformance with MPEP 609; draw line through citation if not in conformance and not considered. Please include a copy of this form with the next communication to applicant.

Other Prior Art

According to the information contained in form PTO-1449 or PTO-892, there are one or more other prior art/non-patent literature documents missing from the original file history record obtained from the United States Patent and Trademark Office. Upon your request we will attempt to obtain these documents from alternative resources. Please note that additional charges will apply for this service.

PAGE 1/16 * RCVD AT 12/10/2003 10:32:41 AM [Eastern Standard Time] * SVR:USPTO-E-EXF-20 * DNS:746877 * CSID:8015831984 * DURATION (mm-ss):07:40

COMPUTER LAW++®

1211 East Yale Avenue
Salt Lake City, Utah 84105
USA

Voice: +1 801 582-2724
Fax: +1 801 583-1984
jwlo@lawplusplus.com

*#5 / Petition to Make
Special -
Accelerated
Examination
12-17-03*

CONFIDENTIAL

If this fax was not intended for you, please notify us immediately, and do not disclose the contents of this material any further. Thank you.

To: Chau M. Nguyen Date: October 28, 2003, *November 28, 2003,*
USPTO Voice#: 703-308-5340 *Dec 10, 2003*
Fax#: 703-872-9306; *703-746-6877* Pages: 16, including this cover sheet.
From: John W.L. Ogilvie
Subject: 10/361,837 (docket 3003.2.11A)

A confirmation copy of this fax will will not follow by mail.

*+ our discussion "12/10"
+ your voicemail 12/10/03*

Thank you for your voicemail this morning requesting another copy of the Petition to Make Special. In addition to that copy, I am enclosing a copy of the PTO-stamped postcard from the first filing of the Petition.

Please do not hesitate to contact me if you have any questions.

Respectfully submitted,

[Signature]
Computer Law++
John W.L. Ogilvie
Reg. No. 37,987

CERTIFICATE OF FAX TRANSMISSION

I hereby certify that this correspondence is being facsimile transmitted to the number indicated above on October 28, 2003: *for you's info & Dec 10, 2003 info*

[Signature]

1 of 16

PATENT APPLICATION
Docket No.: 3003.2.11A

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

RECEIVED
CENTRAL FAX CENTER

DEC 10 2003

In re application of: Sanchaita Datta and Ragula Bhaskar
Serial No.: 10/361837
Filed: February 7, 2003
For: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS
OVER DISPARATE NETWORKS

OFFICIAL

PETITION FOR SPECIAL EXAMINING PROCEDURE
(Accelerated Examination Of New Application)

The Honorable Commissioner of
Patents & Trademarks
Washington, D.C. 20231

Commissioner:

Pursuant to M.P.E.P. § 708.02 VIII, Applicants and Assignee respectfully petition the Office for accelerated examination of the above-identified patent application.

As required, a statement regarding pre-examination search for this application and a detailed discussion of references are submitted below. Copies of the references identified in the search and deemed most closely related to the subject matter encompassed by the claims were filed in a First Information Disclosure Statement on 12 March 2003. Other references, including those made of record in the parent application, are not discussed here, but they are of record and they may be cited by the Examiner as the Examiner sees fit. A petition for accelerated examination has also been filed in the parent application.

If the Office determines that the claims should be made subject to a restriction requirement, an oral election of claims to be initially examined will be made without traverse.

Pre-examination Search

A pre-examination search was made both for relevant patents and for relevant non-patent references, comprising a search of references from the parent application 10/361,837

2.F16

of which this current application is a continuation-in-part. References discussed below are also discussed in a petition to accelerate examination of that parent application.

With respect to U.S. patents, the classes and subclasses of patents identified in the search are as follows:

<u>Class</u>	<u>Sub-class(es)</u>
340	825.03
370	16, 85.13, 218, 392, 409
395	200.06
709	237, 240

Detailed Discussion of the References

Several points should be noted in connection with the references. First, some of the claimed subject matter was used to guide the search. It does not follow from the mere fact that certain references are listed here that one of ordinary skill in the art would have combined these or similar references without the benefit of seeing the claims. In the event it makes a rejection under § 103 using these or any other references, the Office must identify a suggestion or motivation in the art for combining the references.

Second, the discussion below tries to be both complete and concise. By necessity, however, the discussion rests on a good-faith prediction as to which topics the Office will find of interest in examining this application. All participants in the examination process are free to decide later that other aspects of these references and/or other references also merit attention. Of course, the Office will also notify Applicants if examination indicates that the claims and/or references should be interpreted or characterized in some way different from that now presented.

Third, the pre-examination search is not a substitute for the Examiner's search. Likewise, the information provided here is meant to be an aid to the Examiner; it is not meant to be a substitute for the Examiner's own independent review and analysis of the references. In particular, the fact that some of the references discussed below are

emphasized more than the other references does not imply that the Examiner's review of the other references will be cursory or non-existent. Although the information given here is believed to be accurate, errors may nonetheless be present. Also, points whose significance is not currently understood may be discussed here inadequately or not at all.

Fourth, to promote conciseness this initial discussion of the patentability of the claims focuses on certain features of the independent claims. However, other features and combinations of features in both the independent claims and the dependent claims also provide proper grounds for allowing the claims. A lack of patentability will not automatically follow from some later determination (either before or after issuance) that the claim features discussed expressly below are insufficient. Each claim must be viewed as a whole.

Fifth, the technical background of the invention is also discussed in the Technical Background of the Invention portion of the application, and that discussion is incorporated herein by this reference.

Sixth, citation of a reference does not imply adoption of all definitions given in the reference, or agreement with all assertions made in (or implied from) the reference. In particular and without limitation, terms may be used differently in a reference than in the present application; in the event of a conflict, the meaning given to a term (expressly or implicitly) in the application and/or in other statements by Assignee should govern.

Seventh, the dates in reference citations are merely presumptions based on copyright notices, retrieval dates, and/or similar indicia. A document's actual publication date, for instance, may be different than the date printed on the document. Indicia in a single document may specify multiple dates, or a range of dates, with only some of the dates qualifying the document as prior art. A document may also be submitted, even though submission is not required because the document's stated date makes it presumptively not prior art, if the document contains information that might be helpful, such as technical background or a discussion of work that may have been done earlier than the document's stated date.

4 of 16

Finally, a failure to expressly state here that a given reference does not teach a certain claim element does not mean that the reference teaches the claim element. If the Office takes the position that a claim element is taught by reference, then the Office must identify to Assignee the location(s) in the reference which support that position.

Datta '197: U.S. Patent Application No. 10/034,197 filed December 28, 2001

The present application is a continuation-in-part of the '197 application. The Datta '197 application may be of interest to the Examiner as background information and/or for other reasons. For instance, although the undersigned does not believe this reference would support an obviousness-type double patenting rejection of the current application, or vice versa, the Office will make its own independent initial decision regarding that possibility. **If this reference is not cited after this specific invitation to consider the Datta '197 reference, it will be understood that the Examiner has determined the reference is not a basis for rejection.**

Casey: U.S. Patent No. 6,493,349 to Casey

This reference discusses a virtual private network infrastructure and a method of configuring such an infrastructure. Separate VPN areas are discussed, for instance, at column 3 line 27 through column 4 line 26. Claims 7-9, for instance, also indicate that different areas may use different network protocols. However, it appears to the undersigned that the areas are in series, not in parallel. Parallel networks are expressly required by each of the present application's independent claims (claims 1, 10, 11, 12, 26, 30, 31, 32, 33).

Note that "private network" as used in the present application refers to frame relay and point-to-point networks (see the application at page 3 lines 13-20), whereas "private network" as used in Casey apparently refers to customer sites (column 1 lines 22-24). Applicants are entitled to be their own lexicographers, and any confusion over this (or other) terminology used in the application should be resolved in favor of the meaning intended by Applicants even if that meaning conflicts with other possible meanings.

5 of 16

Datta '341: U.S. Patent No. 6,493,341 to Datta et al.

The inventors of this patent are the same as in the present application. This reference claims priority to provisional application no. 60/174,114 filed on December 31, 1999. This patent may be of interest to the Examiner as background information and/or for other reasons. For instance, although the undersigned does not believe this reference would support an obviousness-type double patenting rejection of the current application, or vice versa, the Office will make its own independent initial decision regarding that possibility. **If this reference is not cited after this specific invitation to consider the Datta '341 reference, it will be understood that the Examiner has determined the reference is not a basis for rejection.**

Halpern: U.S. Patent No. 6,438,100 to Halpern et al.

This reference apparently deals mainly with routing inside a Carrier Scale Internetworking system. Frame relay is mentioned in column 2 lines 10-29, 47, and at column 6 line 30. This reference has some discussion of VPNs, e.g., in column 6 lines 14-40, so the remarks made above about the meaning of "private network" in discussing the Casey reference may also be noted here. A keyword search of this reference failed to disclose any use of "parallel" and the reference accordingly does not appear to the undersigned to teach the claimed access to parallel networks.

Rekhter: U.S. Patent No. 6,339,595 to Rekhter et al.

This reference deals with virtual private networks (VPNs), so the remarks made above about the meaning of "private network" in discussing the Casey reference may also be noted here. Note also that the present application defines "disparate" (see page 2 lines 15-18), so the ordinary dictionary definition does not govern in the pending claims. A keyword search of this reference failed to disclose any use of "frame relay" but there are several instances of "point-to-point". Only a single instance of "parallel" was found, at

6 of 16

column 4 line 3. This reference discusses addresses at length, and "address" is used in several independent claims of the current application, namely, claims 12, 26, 30, 31, 32, 33. This reference does not appear to the undersigned to teach the claimed access to parallel networks using a packet path selector to select between parallel disparate networks, or to teach using addresses to combine connections for access to parallel networks as claimed. However, analysis by the Examiner is called for at this point to determine whether the Office agrees with that conclusion. The Examiner is specifically requested to perform a detailed comparison of Rekhter with the pending claims, and to then take such action as the Examiner deems appropriate.

Datta '276: U.S. Patent No. 6,295,276 to Datta et al.

The inventors of this patent are the same as in the present application. As indicated in the Abstract, this patent describes methods, configured storage media, and systems for increasing bandwidth between a local area network ("LAN") and other networks by using multiple routers on the given LAN; Figures 2 and 3 each show a configuration with multiple routers in parallel. Data packets are multiplexed between the routers using a novel variation on the standard address resolution protocol, and other components. On receiving data destined for an external network, a controller or gateway computer will direct the data to the appropriate router. In addition to providing higher speed connections, the invention described in the '276 patent provides better fault tolerance in the form of redundant connections from the originating LAN to a wide area network such as the Internet.

The invention described in the present application is directed to configurations involving parallel networks; every independent claim calls expressly for parallel networks. Although the '276 invention might be usable in a parallel network configuration, that particular type of use is not required by, nor discussed in, the '276 patent.

7 of 16

Bhaskar: U.S. Patent No. 6,253,247 to Bhaskar et al.

The inventors of this patent are the same as in the present application. As indicated in the Abstract, this patent describes methods and systems for transmitting a user's data between two computer networks over physically separate telephone line connections which are allocated exclusively to the user. The user's data is placed in data packets, which are multiplexed onto the separate connections and sent concurrently to a demultiplexer. The data packets contain a computer network address such as an Internet protocol address. A dynamic address and sequence table allows the demultiplexer operation to restore the original order of the data after receiving the packets. The set of connections constitutes a virtual "fat pipe" connection through which the user's data is transmitted more rapidly. Additional users may be given their own dedicated "fat pipe" connections.

As noted above, each independent claim of the present application assumes parallel networks are involved; the invention is not those networks themselves, but it does provide tools and techniques for controlling access to parallel disparate networks. Although the '247 invention might be usable in a parallel network configuration, that particular type of use is not required by, nor discussed, in the '247 patent.

Kitai: U.S. Patent No. 5,948,069 to Kitai et al.

As indicated in the Abstract, this reference discusses a networking method and system for performing data communication to a client computer from a server computer having a plurality of network interfaces through a network. A LAN switch is provided between the network and the server computer. The LAN switch includes a plurality of communication paths correspondingly connected to the network interfaces of the server computer. Any one of the communication paths are usable to connect the client computer with the server computer. A selector is provided for selecting one of the communication paths in accordance with a quality of service (QOS) requested by the client computer. The selector selects the communication path using information contained in a routing table in the server computer based on a network address of the network connected to the client

8 of 16

computer. The routing table includes the address of the network connected to the client computer and addresses of network interfaces of the server computer correspondingly connected to the communication path.

However, a keyword search of this reference disclosed no instances of "frame relay", "point-to-point", "T1", or "T3", which are private networks according to the present application. Although Kitai appears to the undersigned to be one of the references that is closest to the present invention, analysis by the Examiner is called for at this point to determine whether these or other differences or similarities merit further attention. Figure 3 may be of particular interest, since it shows two public networks 3070 and 3080 on what are apparently parallel communication paths between a client 3101 and a server 3000. It may also be important that the choice between network interfaces in Kitai is apparently made at the server (see, e.g., column 10 lines 13-65) rather than elsewhere; in the present application independent claims 1, 10, 11, 26 each refer to a "site interface". Claim 12 involves "address ranges", a term not found in Kitai. Claims 10 and 26 involve "session", a term not found in Kitai. Although Kitai makes repeated references to "destination address" they do not appear to the undersigned to involve modifying the destination address as called for in claims 30, 31, 33 of the present application. In view of all this, the Examiner is specifically requested to perform a detailed comparison of Kitai with the pending claims, and to then take such action as the Examiner deems appropriate.

Periasamy: U.S. Patent No. 5,737,526 to Periasamy et al.

This reference discloses a hierarchical wide area network architecture in which multiple routers having a logical connection to one another are designated as a peer group. Column 3 states that more than one border peer can be included in each group, to share the transmission workload and act as a backup. In another statement (column 3 lines 16-23), two or more routers are connected to a network in parallel to provide back-up facilities. When both of the routers are operating, conditional filters cause one of the routers to drop selected network frames, which are handled by the other router to avoid duplicate frames.

8

9 16

A keyword search revealed no instance of "frame relay", no instance of "point-to-point", one instance of "T1" (column 5 line 55), and no instances of "T3". The parallelism taught is apparently parallelism of routers, not of networks.

Perlman: U.S. Patent No. 5,420,862 to Perlman

This reference discloses a "bridge router (brox)" which functions as a bridge under some conditions and as a router under other conditions. As illustrated in Figure 8, for instance, the broxes connect LANs. However, keyword searches of this reference failed to find any instances of "parallel", "frame relay", "T1", or "T3".

Derby: U.S. Patent No. 5,398,012 to Derby et al.

This reference discloses a process for determining the best communication route from a source end station to a destination end station, using network nodes at the interface between a wide area network and each sub-network. The network nodes contain access agents which control communication flow between the wide area network and an end station in the sub-network.

This reference discusses "parallel links" and "parallel transmission groups", e.g., at column 1 line 43, column 2 lines 48-55. Keyword searching revealed no instance of "frame relay", but "point-to-point" occurs at column 5 line 31, column 7 line 49, column 8 lines 18-20, and column 9 line 43. It also discusses a route selection apparatus, see, e.g., claim 1, for use with subnetworks, which the Examiner may consider different from the claimed invention's selection between parallel networks. Analysis by the Examiner is called for at this point to determine whether these or other differences or similarities merit further attention. The Examiner is specifically requested to perform a detailed comparison of Derby with the pending claims, and to then take such action as the Examiner deems appropriate.

10 F 10

Liao: T. Liao et al., "Using multiple links to interconnect LANs and public circuit switched data networks," *Proc. Int. Conference on Communications Systems: Towards Global Integration, Vol. 1*, Singapore, 59 November 1990, pp. 289-293

This reference discusses design of a gateway that interconnects a TCP/IP-based LAN and a public circuit switched data network using multilinks. Although it discusses use of multiple data links, the parallelism taught is apparently parallelism of multiple physical links to a single network, not parallelism of networks as claimed in the present application.

Coyotepoint: Press release from www.coyotepoint.com, September 8, 1997

This reference discusses a form of load-balancing, and the present application also refers to "load-balancing", see, e.g., claims 3, 18-20, 26, 34. However, this reference does not appear to the undersigned to teach the claimed access to parallel networks.

NAT: Network Address Translation Technical Discussion, from safety.net; no later than 05/07/1999

This reference discusses parallelism in servers and a form of load-balancing, see, e.g., the paragraph on "Mux Server Mapping Mode" on page 4. But it does not appear to the undersigned to teach the claimed access to parallel networks.

Higginson: Higginson et al., "Development of Router Clusters to Provide Fast Failover in IP Networks," from www.asia-pacific.digital.com; no later than 9/29/98

This reference discusses failover, which is related to a form of reliability, and the present application also refers to "reliability", see, e.g., claims 4 and 21. However, this reference does not appear to the undersigned to teach the claimed access to parallel networks.

11 F 16

Navpoint: Pages from www.navpoint.com; no later than 12/24/2001

This reference discusses frame relay and point-to-point connections. However, it does not appear to the undersigned to teach the claimed access to parallel networks.

Guide: "The Basic Guide to Frame Relay Networking", pp. 1-85, copyright date 1998

This reference discusses frame relay networks in depth. Point-to-point and other network technologies are also discussed. However, this reference does not appear to the undersigned to teach selection between parallel networks as called for by the present invention. A keyword search reveals that the word "parallel" is used only in connection with the example shown in Figure 13 on page 51 of the reference. That figure shows "Parallel SNA, BSC, Alarm and LAN Branch networks", as opposed to parallel frame relay and Internet-based networks. Figure 13 also fails to show a packet path selector. Moreover, the parallel nature of the SNA, BSC, and LAN networks is characterized as undesirable; one obtains "better performance, greater reliability and lower costs" by consolidating the data from these networks onto one frame relay-based WAN. By teaching away from parallelism, this reference teaches away from the present invention.

NNI & UNI: "NNI & UNI", pp. 1-2, Nov 16, 2001

This reference gives a definition for a network-to-network interface (NNI) and a definition for a user-to-network interface (UNI). It does not appear to the undersigned to teach parallel networks.

Disaster Recovery: "Disaster Recovery for Frame Relay Networks", pp. 1-14, no later than 12/7/2001

This reference discusses various options for increasing reliability in networking configurations that include a frame relay network. In particular, page 3 identifies "disaster recovery options" that each add something to a frame relay network; these options are discussed later in the reference, including without limitation on page 11. Keyword searches

12 F 16

found no instances of "parallel" in this reference. The reference also does not appear to the undersigned to teach the present invention's packet path selector. However, the Examiner is specifically requested to perform a detailed comparison of this reference with the pending claims, and to then take such action as the Examiner deems appropriate.

Nolle: T. Nolle, "Watching Your Back", pp. 1-3, 11/01/99

This reference discusses frame relay network outages. The first full paragraph on page 2 presents "multiple frame relay carriers" as an option. However, this reference does not appear to the undersigned to teach the present invention's packet path selector.

Multi: "Multi-Attached and Multi-Homed Dedicated Access", pp. 1-5, no later than 12/8/2001

This reference discusses multi-attached and multi-homed access for increased reliability. Frame relay is discussed. Load-balancing is mentioned on page 5; load-balancing is expressly called for in claims 3, 18-20, 26, and 34 of the present application. A keyword search found one use of the word "parallel", on page 3: "Using two parallel circuits between a customer's network and different CLIX routers will satisfy most customers high-availability requirements. For optimum resilience, you should ensure that the two CLIX access circuits do not share any common elements (e.g. a single unprotected tail circuit, a single CLEAR Frame AXIS shelf, or a single mux card), and use separate routers for each access circuit, powered from separate protected power sources if possible." The accompanying diagram on page 3 of the reference is reminiscent of Figure 1 of the present application; a similar but more general diagram shown on page 4 of the reference also resembles Figure 1. However, the Examiner is specifically requested to perform a detailed comparison of this reference with the pending claims, and to then take such action as the Examiner deems appropriate.

13.F.16

Feibel: Feibel, "Internetwork Link," Novell's® Complete Encyclopedia of Networking, copyright date 1995

This reference discusses connections between networks. However, it does not appear to the undersigned to teach the claimed access to parallel networks.

Tanenbaum: Tanenbaum, Computer Networks (3rd Ed.), pp. 396-406; copyright date 1996

This reference discusses connections between networks, and ways in which networks differ from one another. Figures 5-36 and 5-37 may also be of interest. However, this reference does not appear to the undersigned to teach the claimed access to parallel networks using a packet path selector.

Wexler: Wexler, "Frame Relay and IP VPNs: Compete Or Coexist?", from www.bcr.com; July 1999

This reference discusses frame relay and VPNs. In particular, an apparent blurring of the line between the two technologies is discussed, see, e.g., page 3. It does not appear to the undersigned to teach the claimed access to parallel networks.

Conclusion

In view of the above, Assignee respectfully petitions the Office for accelerated examination of the claims. In the event of any questions, the undersigned invites a telephone call from the Office.

Dated May 17, 2003.

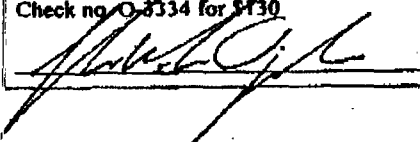
14 of 16

Enclosures
p-petn-MakeSpecial11A

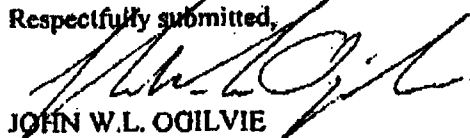
CERTIFICATE OF MAILING

I hereby certify that the correspondence listed below is being deposited with the United States Postal Service as Priority Mail, postage paid, on May 17, 2003 addressed to the Commissioner for Patents, Mail Stop Petition, P.O. Box 1450, Alexandria, VA 22313-1450:

Petition for Special Examining Procedure
Postcard
Check no. 0-3334 for \$130



Respectfully submitted,



JOHN W.L. OGILVIE
Attorney for Applicants & Assignee
Registration No. 37,987

COMPUTER LAW++
1211 East Yale Avenue
Salt Lake City, Utah 84105
801-582-2724 (voice)
801-583-1984 (fax)

15 F 16

USPTO: Please stamp and return
Mailed May 17, 2003 by Priority Mail
Docket No. 3003.2.11A:
Certificate of Mailing
Petition for Special Examining Procedure
Check no. O-3334 for \$130



16 of 16



UNITED STATES PATENT AND TRADEMARK OFFICE

COMMISSIONER FOR PATENTS
UNITED STATES PATENT AND TRADEMARK OFFICE
P.O. Box 1450
ALEXANDRIA, VA 22313-1450
www.uspto.gov

MAIL

Paper No. 6

JOHN W. L. OGILVIE
1320 EAST LAIRD AVENUE
SALT LAKE CITY UT 84105

JAN 26 2004

DIRECTOR OFFICE
TECHNOLOGY CENTER 2600

In re Application of :
Sanchaita DATTA, et al. :
Application No. 10/361,837 :
Filed: February 7, 2003 :
For: TOOLS AND TECHNIQUES FOR DIRECTING :
PACKETS OVER DISPARATE NETWORKS :

DECISION ON PETITION
TO MAKE SPECIAL


This is a decision on the petition filed December 10, 2003 under Manual of Patent Examination Procedure §708.02, VIII requesting accelerated examination.

The petition under Manual of Patent Examination Procedure §708.02, VIII, must:

- (1) be filed prior to receiving any examination by the examiner,
- (2) be accompanied by the required fee- \$130,
- (3) the claims should be directed to a single invention (if it is determined that the claims pertain to more than one invention, then applicant will have to make an election without traverse or forfeit accelerated examination status),
- (4) state that a pre-examination search was made, and fully discuss the search method employed, such as classes and subclasses searched, publications, Chemical abstracts, patents, etc. A search made by a foreign patent office satisfies this requirement,
- (5) be accompanied by a copy of each of the references most closely related to the subject matter encompassed by the claims if said references are not already of record,
- (6) fully discuss the references, pointing out with the particularity required by 37 C.F.R. §1.111 (b) and (c), how the claimed subject matter is patentable over the references.

The petitioner meets all the above-listed requirements. Accordingly, the petition is **GRANTED**.

The application will retain its special status throughout its entire prosecution, including any appeal to the Board of Patent Appeals and Interferences, subject only to diligent prosecution by the applicant. The application file is being forwarded to the examiner for appropriate action in due course.


Kenneth A. Wieder
Special Program Examiner
Technology Center 2600
Communications



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A	3645
23484	7590	02/25/2004	EXAMINER	
JOHN W. L. OGILVIE 1320 EAST LAIRD AVENUE SALT LAKE CITY, UT 84105			MARCELO, MELVIN C	
			ART UNIT	PAPER NUMBER
			2663	7

DATE MAILED: 02/25/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

DETAILED ACTION

Claim Rejections - 35 USC § 112

1. Claims 23-26 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Claim 23 lacks a proper antecedent basis to claim 12 since "the modifying step" first appears in claim 13.

Claim 24 lacks a proper antecedent basis to claim 12 since "the modifying step" first appears in claim 13.

Claim 25 lacks a proper antecedent basis to claim 12 since "the modifying step" first appears in claim 13.

Claim 26, line 7, "possibly with a modified destination address" is indefinite since no clear guidelines are recited for defining the conditions to determine the different possibilities.

Claim Rejections - 35 USC § 102

2. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless --

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

3. Claims 1, 3, 4, 9, 10, 26, 29 and 32 are rejected under 35 U.S.C. 102(b) as being anticipated by Kaplan et al. (6,016,307).

Kaplan teaches the subject matter of the following claims, wherein references to Kaplan appear in parenthesis.

1. A controller (Kaplan, switching system 10 in Figure 1) which controls access to multiple independent disparate networks in a parallel network configuration (T1 12, LAN 14, WAN 16, POTS 18, WIRELESS 20), the disparate networks comprising at least one private network (T1 is a private network) and at least one network based on the Internet (WAN is a collection of computer networks based on the Internet; see also column 1, lines 22-26), the controller comprising: a site interface connecting the controller to a site (User Interface 34 in Figure 1); at least two network interfaces which send packets toward the disparate networks (T1 12 and WAN 16); and a packet path selector (Routing Optimization 26) which selects between network interfaces according to at least: a destination of the packet, an optional presence of alternate paths to that destination (Column 3, line 55 to column 4, line 11), and at least one specified criterion for selecting between alternate paths when such alternate paths are present (Column 4, lines 12-64); wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector (Column 7, lines 39-44).

3. The controller of claim 1, wherein the packet path selector selects between network interfaces according to a load-balancing criterion, thereby promoting balanced loads on devices that carry packets on the selected path after the packets leave the selected network interfaces (Latency is a load-balancing criterion, since it is a measure of the amount of traffic on the particular networks; see column 6, lines 1-9).

4. The controller of claim 1, wherein the packet path selector selects between network interfaces according to a reliability criterion, thereby promoting use of devices that will still carry packets on the selected path after the packets leave the selected network interfaces, when other devices on a path not selected are not functioning (Kaplan's criteria include reliability in Table A and presentstate (operational state) in Table B, both on column 4).

9. The controller of claim 1, wherein the controller sends packets from a selected network interface to a point-to-point private network connection (**T1 is a point-to-point private network**).

10. A controller which controls access to multiple networks in a parallel network configuration (**Kaplan, switching system 10 in Figure 1**), suitable networks comprising Internet-based networks (**WAN 16**) and private networks from at least one more provider (**T1 12**), in combination, the controller comprising: a site interface connecting the controller to a site (**User interface 34**); at least two network interfaces which send packets toward the networks (**T1 12 and WAN 16**); and a packet path selector (**Routing Optimization 26**) which selects between network interfaces on granularity which is at least as fine as session-by-session (**Kaplan's granularity is session-by-session since path selection occurs in order to transfer a single data file during the user session; see column 6, lines 54-65**) according to at least: a destination of the packet, an optional presence of alternate paths to that destination, and at least one specified criterion for selecting between alternate paths when such alternate paths are present (**Column 3, line 55 to column 64**); wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector (**Column 7, lines 39-44**).

26. A method for combining connections for access to parallel networks (**Kaplan, column 3, line 55 to column 4, line 11**), the method comprising the steps of: sending a packet to a site interface of a controller (**In Figure 1, Data File 30 is entered via the User Interface 34**), the controller (**Switching system 10**) comprising the site interface (**User Interface 34**) which receives packets, at least two network interfaces to parallel networks (**T1 12 and WAN 16**), and a packet path selector (**Routing Optimization 26**) which selects between the network interfaces on a per-session basis (**Selection occurs during an user session; see column 6, lines 54-65**) to promote load-balancing (**Selection criteria includes latency which is a measure of load in the networks; see column 6, lines 1-9**); and forwarding the packet, possibly with a modified destination address, through the network interface selected by the packet path selector (**column 7, lines 39-44, wherein destination address formats for the various**

network interfaces 12-20 may differ from the destination address of the data file (column 1, lines 11-26) such that the address is modified prior to transmission over the network).

29. The method of claim 26, further comprising the step of sensing failure of one of parallel disparate networks and automatically sending traffic through at least one other parallel disparate network (Kaplan's criteria includes presentstate (Table B on column 4) which indicates the not operational/failure of the particular network).

32. A method for combining connections for access to disparate parallel networks (Column 3, line 55 to column 4, line 11), the method comprising the steps of: receiving at a controller (Switching system 10 in Figure 1) a packet which has a first site IP address as source address and a second site IP address as destination address (LAN 14 and WAN 16 transmit packets using the IP (Internet) protocol, wherein the source address is the IP address of the source site which is the switching system 10 and the destination address is the IP address of a corresponding destination switching system 10); selecting, within the controller (Routing Optimization 26), between a path through an Internet-based network (WAN 16) and a path through a private network that is not Internet-based (T1); and forwarding the packet along the selected path toward the second site (Column 7, lines 39-44).

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 2, 8 and 28 are rejected under 35 U.S.C. 103(a) as being unpatentable over Kaplan et al. as applied to the above claims and further in view of applicants' admitted prior art.

Kaplan does not mention the frame relay private network or a VPN network interface. However, Kaplan teaches that a great variety of networks exist (column 1, lines 19-26) and that a user should use the networks that are available to them (column 1, lines 40-50). Kaplan explicitly teaches to incorporate networks provided by common carriers such as MCI and AT&T (column 4, lines 4-9). Applicants have admitted that Frame relay and /or point-to-point network services are provided by carriers such as AT&T and MCI (specification, page 2, lines 13-14). Therefore, it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate a frame relay network in Kaplan, since a skilled artisan would have been motivated by Kaplan's explicit teaching to incorporate the admitted prior art networks provided by AT&T and MCI.

Similarly, applicants admit that the VPN network is prior art (see Figure 5). Therefore, it would have been obvious to incorporate a VPN network in Kaplan, since a skilled artisan would have been motivated to use any prior art network that is available to them as suggested by Kaplan.

With respect to the claims, references to the prior art appear in parenthesis.

2. The controller of claim 1, wherein the controller (Kaplan, switching system 10 in Figure 1) controls access to a frame relay private network (Frame relay networks are admitted prior art) through a first network interface of the controller, and the controller controls access to the Internet through a second network interface (WAN 16) of the controller.

8. The controller of claim 1, wherein the controller sends packets from a selected network interface to a VPN (VPN is admitted prior art).

28. *The method of claim 26, wherein the step of sending a packet to the controller site interface is repeated as multiple packets are sent (In Kaplan, the data entering the site interface (User Interface 34) can be formatted as packets (column 1, lines 16-19), wherein multiple packets correspond to multiple data), the network interfaces include at least two VPN line interfaces (VPN is admitted prior art) and a private network interface (Kaplan, T1 16 in Figure 1), and the packet path selector selects between at least those three interfaces (It would have been obvious to provide any known combinations of the prior art network interfaces since Kaplan suggests that "any number of the aforementioned interfaces may be used alone or in any combination as required by the user" (column 4, lines 2-4) and that an user may have more than one of a particular interface (column 1, lines 40-45)).*

Allowable Subject Matter

6. Claims 5-7 and 27 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

7. Claims 11-22, 30, 31 and 34-35 are allowed.

8. Claims 23-25 would be allowable if rewritten or amended to overcome the rejection(s) under 35 U.S.C. 112, second paragraph, set forth in this Office action.

9. The following is a statement of reasons for the indication of allowable subject matter: the prior art of record fails to anticipate or make obvious the additional features of the claimed invention. With respect to claim 11, the additional feature of the per-packet selection is not taught. With respect to claims 12 and 33, the feature of the accessing the multiple parallel disparate networks using at least two known location address ranges is not taught. With respect to claims 30 and 31, the modification of the packet to both the source and destination controller IP addresses is not taught, wherein

controller is interpreted as a device which functions as a source or destination on either ends of the combined connections for access to parallel networks from which a path is selected.

5. *The controller of claim 1, wherein the packet path selector selects between network interfaces according to a security criterion, thereby promoting use of multiple disparate networks to carry different pieces of a given message so that unauthorized interception of packets on fewer than all of the disparate networks used to carry the message will not provide the total content of the message.*

6. *The controller of claim 1, wherein the controller sends packets out of sequence over the parallel disparate networks.*

7. *The controller of claim 6, wherein the controller places an encrypted sequence number in at least some of the packets which are sent out of sequence.*

11. *A controller which controls access to multiple networks in a parallel network configuration, suitable networks comprising Internet-based networks and private networks from at least one more provider, in combination, the controller comprising: a site interface connecting the controller to a site; at least two network interfaces which send packets toward the networks; and a packet path selector which selects between network interfaces on a per-packet basis according to at least: a destination of the packet, an optional presence of alternate paths to that destination, and at least one specified criterion for selecting between alternate paths when such alternate paths are present; wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector.*

12. *A method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of: obtaining at least two known location address ranges which have associated networks; obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location; receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location; determining whether the destination address lies*

within a known location address range; selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location; forwarding the packet on the selected network path.

13. The method of claim 12, further comprising the step of modifying the packet destination address to lie within a known location address range associated with the selected network before the forwarding step.

14. The method of claim 12, wherein the forwarding step forwards the packet toward the Internet when the packet's destination address does not lie within any known location address range.

15. The method of claim 12, wherein the destination address identifies a destination location to which only a single associated network provides connectivity from the current location, and the forwarding step forwards the packet to that single associated network.

16. The method of claim 12, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

17. The method of claim 12, wherein repeated instances of the selecting step make network path selections on a per session basis.

18. The method of claim 12, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.

19. The method of claim 18, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance line loads by distributing packets between lines.

20. The method of claim 18, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance network loads by distributing packets between disparate networks.

21. The method of claim 12, wherein the selecting step selects the network path at least in part on the basis of a reliability criterion.

22. *The method of claim 12, wherein the selecting step selects the network path at least in part on the basis of a security criterion.*

23. *The method of claim 12, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a private network such that the modified packet destination address lies instead in a known location address range associated with a VPN.*

24. *The method of claim 12, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a VPN such that the modified packet destination address lies instead in a known location address range associated with a private network.*

25. *The method of claim 12, wherein the modifying step modifies a packet destination address corresponding to one of: the Internet, a private network, thereby making the modified packet destination address correspond to the other of: the Internet, a private network.*

27. *The method of claim 26, wherein the step of sending a packet to the controller site interface is repeated as multiple packets are sent, and the controller sends different packets of a given message to different parallel networks.*

30. *A method for combining connections for access to parallel networks, the method comprising the steps of: receiving at a first controller a packet which has a first site IP address as source address and a second site IP address as destination address; modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address; and forwarding the modified packet along a selected path toward the second site.*

31. *A method for combining connections for access to parallel networks, the method comprising the steps of: receiving at a first controller a packet which has a first VPN IP address as source address and a second VPN IP address as destination address; modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address; and forwarding the modified packet along a selected path toward the second VPN.*

33. A computer storage medium having a configuration that represents data and instructions which will cause performance of a method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of: obtaining at least two known location address ranges which have associated networks; obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location; receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location; determining whether the destination address lies within a known location address range; selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location; modifying the packet destination address to lie within a known location address range associated with the selected network if it does not already do so; and forwarding the packet on the selected network path.

34. The configured storage medium of claim 33, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.

35. The configured storage medium of claim 33, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

Conclusion

10. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Schoffelman et al. (US 6119170 A) teach a multi-homed system which provides connections over parallel networks. Wootten et al. (US 6128298 A) teach address modification using address translation between private and public networks.

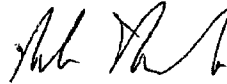
Application/Control Number: 10/361,837
Art Unit: 2663

Page 12

11. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Melvin Marcelo whose telephone number is 703-305-4373. The examiner can normally be reached on Monday-Friday, 8:30 to 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Chau Nguyen can be reached on 703-308-5340. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).



Melvin Marcelo
Primary Examiner
Art Unit 2663

Mm
February 23, 2004

Notice of References Cited	Application/Control No. 10/361,837	Applicant(s)/Patent Under Reexamination DATTA ET AL.	
	Examiner Melvin Marcelo	Art Unit 2663	Page 1 of 1

U.S. PATENT DOCUMENTS

*	Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
A	US-6,016,307 A	01-2000	Kaplan et al.	370/238
B	US-6,119,170 A	09-2000	Schoffelman et al.	709/244
C	US-6,128,298 A	10-2000	Wootton et al.	370/392
D	US-			
E	US-			
F	US-			
G	US-			
H	US-			
I	US-			
J	US-			
K	US-			
L	US-			
M	US-			

FOREIGN PATENT DOCUMENTS

*	Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
N					
O					
P					
Q					
R					
S					
T					

NON-PATENT DOCUMENTS

*	Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
U	
V	
W	
X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.



04-07-04

2663

Attorney Docket No. RADW 21.090 (101092-00074)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventor: SANCHAITA DATTA

Confirmation No.: 3645

Serial No.: 10/361,837

Filed: February 7, 2003

Title: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS...

Examiner: MELVIN C. MARCELO

Group Art Unit: 2663

RECEIVED #8
APR 09 2004
Technology Center 2600

April 5, 2004

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

THIRD PARTY SUBMISSION

SIR:

Please withdraw the fees for this third party submission from deposit account 50-1290, as set forth in 37 CFR 1.17(p) and 37 CFR 1.17(i).

Submitted for consideration is the following documents and publication date:

- 1) U.S. Patent No. 6,665,702B1 Issued December 16, 2003;
- 2) "Radware announces LinkProof: The first IP Load Balancing Solution for networks with multiple ISP connection" Published October 7, 1999;
- 3) "Radware Balances the Network" Published January 7, 2000;
- 4) "Global Product Spotlight: Radware Linkproof" Published December 1, 1999;
- 5) "Radware Seeks Solutions to Easy-Access Problems" Published December 1, 1999;

11180751.01

Filed by Express Mail
(Receipt No. EL979034379 US)
on 4-5-2004
pursuant to 37 C.F.R. 1.10.
by Francis Boyle

This submission has been served upon the applicant in accordance with 37 CFR 1.248.

Proof of service is attached.

This submission is after the two months from the time the application was published because:

1. The publication of the application only became known to the third party submitter on or about January 30, 2004; and
2. The U.S. patent issued on December 16, 2003, which was after the two month period had expired and therefore could not have been submitted within the time period.

Respectfully submitted,



Brian S. Myers
Reg. No. 46,947

CUSTOMER NUMBER 026304
Telephone: (212) 940-8703
Fax: (212) 940-8986 or 8987
Docket No.: RADW 21.090 (101092-00074)
BSM:fd

11180751.01

CERTIFICATE OF SERVICE

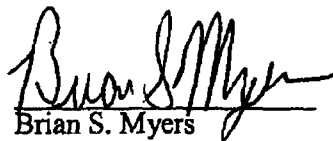
I hereby certify that on April 2, 2004, I caused the foregoing **THIRD PARTY SUBMISSION** to be served as follows:

by U.S. Mail, first class, by depositing the same in a depository of the United States

Postal Service, on:

John W. L. Ogilvie
Computer Law
1211 East Yale Ave.
Salt Lake City, UT 84105

Attorney for Applicant


Brian S. Myers

MAY-18-04 TUE 02:46 PM

FAX NO. RECEIVED
CENTRAL FAX CENTER P. 01

MAY 18 2004

#9

OFFICIAL

FACSIMILE MEMORANDUM

THORPE NORTH & WESTERN L.L.P.
SINCE 1979
THE TEAM APPROACH TO PREMIER PERFORMANCE ®

P.O. Box 1219
SANDY, UTAH 84091-1219
USA
TELEPHONE 1.801.566.6633
FACSIMILE 1.801.566.6673
FACSIMILE 1.801.566.0750

DATE: MAY 18, 2004
TO: MELVIN C. MARCELO
ASSISTANT COMMISSIONER FOR PATENTS
FROM: JOHN W. L. OGILVIE
OUR DOCKET NUMBER: 22973.NP
INVENTOR: SANCHAITA DATTA AND RAGULA BHASKAR
TITLE: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS
SUBJECT: AMENDMENT FILED VIA FAX

PAGE 1 OF: 10
FACSIMILE NUMBER: (703)872-9306
TRANSMITTED BY: SHEILA

IF YOU DO NOT RECEIVE THE COMPLETE DOCUMENT, PLEASE NOTIFY THE SENDER AS SOON AS POSSIBLE.

REMARKS:

PLEASE ACKNOWLEDGE RECEIPT

INCLUDED ARE: AMENDMENT (9 PAGES)
CERTIFICATE OF FAX



INTENDED SOLELY FOR THE ADDRESSEE ABOVE. IF YOU RECEIVE THIS MESSAGE AND ARE NOT THE AGENT OR EMPLOYEE OF THE ADDRESSEE, AND HAVE THEREFORE BEEN SENT OR RECEIVED THIS COMMUNICATION IN ERROR, YOU ARE ASKED NOT TO DISSEMINATE OR COPY ANY OF THE ATTACHED AND ARE TO NOTIFY THE SENDER IMMEDIATELY BY TELEPHONE. PLEASE ALSO RETURN THE ORIGINAL MESSAGE TO THE SENDER BY MAIL. THANK YOU.

a

MAY-18-04 TUE 02:46 PM

FAX NO.

RECEIVED
CENTRAL FAX CENTER

MAY 18 2004

P. 02

AJG
5/20/04
Debb

OFFICIAL

PATENT APPLICATION
ATTORNEY DOCKET NO. 3003.2.11A / 22973.NP

IN THE UNITED STATES PATENT & TRADEMARK OFFICE

<p>ART UNIT: 2663</p> <p>EXAMINER: Melvin C. Marcelo</p> <p>APPLICANT: Sanchaita Datta and Ragula Bhaskar</p> <p>SERIAL NO.: 10/361,837</p> <p>FILED: February 7, 2003</p> <p>FOR: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS</p>	<p style="text-align: center;">AMENDMENT</p> <div style="border: 1px solid black; padding: 5px;"> <p style="text-align: center;">CERTIFICATE OF FAX TRANSMISSION</p> <p>DATE OF FAXING: <u>May 18, 2004</u></p> <p>I hereby certify that this paper or fee (along with any paper or fee referred to as being attached or enclosed) is being faxed to the USPTO central fax number 703 872-9306 on the date indicated above.</p> <p style="text-align: center;"><i>Shella Halterman</i> Shella Halterman</p> </div>
---	--

Mail Stop Non-Fee Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Patent Official:

In response to the Office Action mailed February 25, 2004, Applicants and Assignee respectfully submit the following remarks and amendments to place the application in condition for allowance. This application has been granted "accelerated examination" status.

Change of Customer Number and Address

Please note that this application should now be associated with Customer Number 20,551, at the address for that Customer Number. The undersigned is already an attorney of record in this case; he has simply changed firms.

1

PAGE 2/10 * RCVD AT 5/18/2004 5:47:03 PM [Eastern Daylight Time] * SVR:USPTO-EFAX-1/4 * DNIS:8729306 * CSID: * DURATION (mm-ss):02:46

A

Claim Listing

Please cancel claims without prejudice, and amend claims, as shown below.

1-4. (canceled)

5. (currently amended) A controller which controls access to multiple independent disparate networks in a parallel network configuration, the disparate networks comprising at least one private network and at least one network based on the Internet, the controller comprising:

a site interface connecting the controller to a site;

at least two network interfaces which send packets toward the disparate networks; and

a packet path selector which selects between network interfaces according to at least: a

destination of the packet, an optional presence of alternate paths to that

destination, and at least one specified criterion for selecting between alternate

paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the packet

through the network interface that was selected by the packet path selector; and

~~The controller of claim 1,~~ wherein the packet path selector selects between network

interfaces according to a security criterion, thereby promoting use of multiple

disparate networks to carry different pieces of a given message so that

unauthorized interception of packets on fewer than all of the disparate networks

used to carry the message will not provide the total content of the message.

2. (currently amended) The controller of claim [1]¹2, wherein the controller sends packets out of sequence over the parallel disparate networks.

3. (original) The controller of claim 2, wherein the controller places an encrypted sequence number in at least some of the packets which are sent out of sequence.

8-10. (canceled)

a

4
11. (original) A controller which controls access to multiple networks in a parallel network configuration, suitable networks comprising Internet-based networks and private networks from at least one more provider, in combination, the controller comprising:

a site interface connecting the controller to a site;

at least two network interfaces which send packets toward the networks; and

A
a packet path selector which selects between network interfaces on a per-packet basis according to at least: a destination of the packet, an optional presence of alternate paths to that destination, and at least one specified criterion for selecting between alternate paths when such alternate paths are present;

wherein the controller receives a packet through the site interface and sends the packet through the network interface that was selected by the packet path selector.

5
12. (original) A method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of:

obtaining at least two known location address ranges which have associated networks; obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location;

receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location;

determining whether the destination address lies within a known location address range;

selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location;

forwarding the packet on the selected network path.

⁶
~~13~~. (original) The method of claim ⁵~~12~~, further comprising the step of modifying the packet destination address to lie within a known location address range associated with the selected network before the forwarding step.

⁷
~~14~~. (original) The method of claim ⁵~~12~~, wherein the forwarding step forwards the packet toward the Internet when the packet's destination address does not lie within any known location address range.

⁸
~~15~~. (original) The method of claim ⁵~~12~~, wherein the destination address identifies a destination location to which only a single associated network provides connectivity from the current location, and the forwarding step forwards the packet to that single associated network.

⁹
~~16~~. (original) The method of claim ⁵~~12~~, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

¹⁰
~~17~~. (original) The method of claim ⁵~~12~~, wherein repeated instances of the selecting step make network path selections on a per session basis.

¹¹
~~18~~. (original) The method of claim ⁵~~12~~, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.

¹²
~~19~~. (original) The method of claim ¹¹~~18~~, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance line loads by distributing packets between lines.

¹³
~~20~~. (original) The method of claim ¹¹~~18~~, wherein repeated instances of the selecting step select between network paths at least in part on the basis of a dynamic load-balancing criterion which tends to balance network loads by distributing packets between disparate networks.

a

¹⁴
~~21.~~ (original) The method of claim ⁵~~12~~, wherein the selecting step selects the network path at least in part on the basis of a reliability criterion.

¹⁵
~~22.~~ (original) The method of claim ⁵~~12~~, wherein the selecting step selects the network path at least in part on the basis of a security criterion.

¹⁶
~~23.~~ (currently amended) The method of claim [12] ⁶~~12~~, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a private network such that the modified packet destination address lies instead in a known location address range associated with a VPN.

¹⁷
~~24.~~ (currently amended) The method of claim [12] ⁶~~12~~, wherein the modifying step modifies a packet destination address which was in a known location address range associated with a VPN such that the modified packet destination address lies instead in a known location address range associated with a private network.

¹⁸
~~25.~~ (currently amended) The method of claim [12] ⁶~~12~~, wherein the modifying step modifies a packet destination address corresponding to one of: the Internet, a private network, thereby making the modified packet destination address correspond to the other of: the Internet, a private network.

¹⁹
~~26.~~ (canceled)

¹⁹
~~27.~~ (currently amended) A method for combining connections for access to parallel networks, the method comprising the steps of:

 sending a packet to a site interface of a controller, the controller comprising the site interface which receives packets, at least two network interfaces to parallel

networks, and a packet path selector which selects between the network interfaces on a per-session basis to promote load-balancing; and forwarding the packet through the network interface selected by the packet path selector; ~~The method of claim 26,~~ wherein the step of sending a packet to the controller site interface is repeated as multiple packets are sent, and the controller sends different packets of a given message to different parallel networks.

28-29. (canceled)

20
30.

(original) A method for combining connections for access to parallel networks, the method comprising the steps of:

- receiving at a first controller a packet which has a first site IP address as source address and a second site IP address as destination address;
- modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address; and
- forwarding the modified packet along a selected path toward the second site.

21
31.

(original) A method for combining connections for access to parallel networks, the method comprising the steps of:

- receiving at a first controller a packet which has a first VPN IP address as source address and a second VPN IP address as destination address;
- modifying the packet to have an IP address of the first controller as the source address and an IP address of a second controller as the destination address; and
- forwarding the modified packet along a selected path toward the second VPN.

32. (canceled)

a

AA
cancel

²²₃₃. (original) A computer storage medium having a configuration that represents data and instructions which will cause performance of a method for combining connections for access to multiple parallel disparate networks, the method comprising the steps of:

- obtaining at least two known location address rangos which have associated networks;
- obtaining topology information which specifies associated networks that provide, when working, connectivity between a current location and at least one destination location;
- receiving at the current location a packet which identifies a particular destination location by specifying a destination address for the destination location;
- determining whether the destination address lies within a known location address range;
- selecting a network path from among paths to disparate associated networks, said networks being in parallel at the current location, each of said networks specified in the topology information as capable of providing connectivity between the current location and the destination location;
- modifying the packet destination address to lie within a known location address range associated with the selected network if it does not already do so; and
- forwarding the packet on the selected network path.

²³₃₄. (original) The configured storage medium of claim ²²₃₃, wherein the selecting step selects the network path at least in part on the basis of a dynamic load-balancing criterion.

²⁴₃₅. (original) The configured storage medium of claim ²²₃₃, wherein repeated instances of the selecting step make network path selections on a packet-by-packet basis.

Remarks

The Office Action stated that claims 11-22, 30, 31, and 33-35 were allowed. Those claims are repeated above in the original form allowed.

The Office Action also stated that claims 1-4, 8-10, 23-26, 28, 29 and 32 were rejected.

a

Of these, claims 1-4, 8-10, 26, 28, 29 and 32 are each canceled, although the right to pursue some or all of them in a subsequent application is reserved.

The remaining rejected claims (claims 23-25) were rejected merely as lacking antecedent basis, which was due to a typographical error. The correction suggested by the Office Action has been adopted above - these claims now depend from claim 13 instead of claim 12. The undersigned respectfully submits that the intended meaning in each of these claims was clear (mere objection would have sufficed), and that this correction should accordingly have no estoppel effect under *Festo*. The undersigned also thanks the Examiner for identifying this error in the recitation of dependency, since a certificate of correction may otherwise have been needed.

Conclusion

In light of the above, Applicants and Assignee respectfully submit that all pending claims are allowable. They request that the rejections be withdrawn, and that the claims be allowed and passed to issue. Their silence here does not signify agreement or acquiescence in the Office Action's assertions, and they reserve all arguments.

No claims were added. Claims 5 and 27 were amended to independent form. Claims 1-4, 8-10, 26, 28, 29, 32 were canceled, including independent claims 1, 10, 26, 32. Therefore, no additional fee is due. However, the Commissioner is authorized to charge any additional fee or to credit any overpayment in connection with this Amendment to Deposit Account No. 20-0100.

If any impediment to the allowance of these claims remains after entry of this Response, the Examiner is strongly encouraged to call John Ogilvie at 801-566-6633 so that such matters may be resolved as expeditiously as possible.

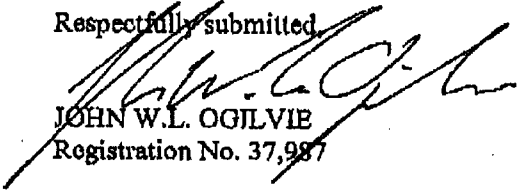
MAY--18-04 TUE 02:48 PM

FAX NO.

P. 10

DATED this 18th day of May, 2004.

Respectfully submitted,


JOHN W.L. OGILVIE
Registration No. 37,987

THORPE NORTH & WESTERN, LLP
Customer No. 20,551
P.O. Box 1219
Sandy, Utah 84091-1219
801-566-6633 (voice)
801-566-0750 (fax)

9

PAGE 10/10 * RCVD AT 5/18/2004 5:47:03 PM [Eastern Daylight Time] * SVR:USPTO-EFAX-1/4 * DNS:8729306 * CSID: * DURATION (mm-ss):02-46



Notice of Allowability	Application No.	Applicant(s)	
	10/361,837	DATTA ET AL.	
	Examiner	Art Unit	
	Melvin Marcelo	2663	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--
 All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. This communication is responsive to amendment filed 05-18-2004.
2. The allowed claim(s) is/are 5-7,11-25,27,30,31 and 33-35, renumbered as 1-24, respectively.
3. The drawings filed on 07 February 2003 are accepted by the Examiner.
4. Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 - a) All b) Some* c) None of the:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

* Certified copies not received: _____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.
6. CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
 - (a) including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
 - 1) hereto or 2) to Paper No./Mail Date _____.
 - (b) including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).
7. DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

- | | |
|--|---|
| 1. <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 5. <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| 2. <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 6. <input type="checkbox"/> Interview Summary (PTO-413),
Paper No./Mail Date _____ |
| 3. <input checked="" type="checkbox"/> Information Disclosure Statements (PTO-1449 or PTO/SB/08),
Paper No./Mail Date <u>06-05-03</u> | 7. <input type="checkbox"/> Examiner's Amendment/Comment |
| 4. <input type="checkbox"/> Examiner's Comment Regarding Requirement for Deposit
of Biological Material | 8. <input type="checkbox"/> Examiner's Statement of Reasons for Allowance |
| | 9. <input type="checkbox"/> Other _____ |

Melvin Marcelo
 Melvin Marcelo
 Primary Examiner
 Art Unit: 2663

Notice of References Cited

Application/Control No. 10/361,837	Applicant(s)/Patent Under Reexamination DATTA ET AL.	
Examiner Melvin Marcelo	Art Unit 2663	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-6,665,702 B1	12-2003	Zisapel et al.	718/105
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
*	U	'Radware announces LinkProof: The first IP Load Balancing Solution for networks with multiple ISP connection', Press Release, published Oct. 7, 1999.
*	V	'Radware Balances the Network', Internet Traffic Management Center, published Jan. 1, 2000.
*	W	'Global Product Spotlight: Radware Linkproof', NetworkMagazine.com, published Dec. 1, 1999.
*	X	'Radware Seeks Solutions to Easy-Access Problems', South China Morning Post, published Dec. 7, 1999.

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
 Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Radware announces LinkProof: The first IP Load Balancing Solution for networks with multiple ISP connection

(J)

Mahwah, NJ; October 7, 1999. To ensure 7x24 availability many enterprises, e-commerce sites and regional ISPs are utilizing multiple Internet router connections. The LinkProof by Radware (Nasdaq: RDWR) is the first technology designed to intelligently load balance IP traffic between these "multi-homed" sites, creating redundancy and eliminating single points of failure.

Deploying independent router connections to two or more ISPs creates these multi-homed sites. This diversity ensures 7x24 availability and an uninterrupted packet delivery to and from the enterprise in the event one or more ISP connection falls. While this adds redundancy, it also creates configuration complexity that may necessitate intricate routing protocols such as BGP (Border Gateway Protocol) and/or coordination between the contracted ISPs.

LinkProof removes this complexity by taking responsibility for the packet delivery through a healthy ISP connection. Sitting logically between the enterprise network and a farm of Internet routers, the LinkProof verifies ISP health and intelligently load balances all inbound and outbound traffic. In addition, it performs Smart NAT to ensure the uninterrupted packet delivery to and from the enterprise network. Smart NAT allows the LinkProof to perform network address translation according to the ISP connection selected to carry the session to the Internet. For example, if the LinkProof chooses ISP_1 for outbound session delivery, then the translated source address will belong to the ISP_1 IP address pool for the inbound response.

Internet traffic is optimized by the LinkProof through intelligent load balancing based on the current session and/or load per verified ISP connection. Additionally, network proximity is measured to determine the closest and fastest route. Network proximity is calculated in both router hops and round trip latency. This allows multi-homed sites to transmit information through a fast, healthy route.

LinkProof also uses proximity detection to perform inbound traffic management. For Internet users attempting to access a resource on the enterprise network (such as a Web server), the LinkProof uses DNS to assure the most optimal ISP connection. This feature allows the LinkProof to consistently use the best and quickest path to satisfy user requests for information.

The LinkProof continuously monitors the health of all routers in the farm, and periodically checks each router path and the health of user defined nodes beyond the router. This monitoring allows the LinkProof to continually send sessions through healthy routers on a healthy Internet path.

About Radware

Radware develops, manufactures and markets products that manage and direct Internet traffic among network resources to enable continuous access to Web sites and other services, applications and content based on the Internet protocol. Radware offers a broad range of Internet traffic management solutions to service providers, e-commerce businesses and corporate enterprises that require uninterrupted availability and optimal

performance of IP-based applications that are critical to their business. Radware's internet traffic management solutions enable its customers to manage their network infrastructure to bypass system failures and to scale their network infrastructure to accommodate increasing IP traffic. Radware's products improve the productivity of network infrastructure by distributing traffic within a network to optimize the use of available network resources. Radware's products can be deployed either as independent solutions to address specific application needs at a particular location within a network or as an end-to-end integrated solution to manage traffic throughout a network.

This press release contains forward-looking statements that are subject to risks and uncertainties. Factors that could cause actual results to differ materially from these forward-looking statements include, but are not limited to, general business conditions in the internet traffic management industry, changes in demand for internet traffic management products, the timing and amount or cancellation of orders and other risks detailed from time to time in Radware's filings with the Securities and Exchange Commission, including Radware's Form F-1.

Radware Balances the Network

Internet Traffic Management Center, January 1, 2000.



By Peter Christy

One of the absolutely thrilling parts of our job is being exposed to the continuing innovation of in the industry. We love watching the process of application invention - new ideas, seemingly out of the blue, that redefine "common knowledge" on what the product category is good for.

In the past, Alteon had some of our favorite inventions: cache redirection and balancing was certainly a good idea, and they invented a particularly cute DNS request capture application. This time we focus on Radware with LinkProof - their invention for balancing and managing multi-homed connections out to the Internet.

Multi-homing is a simple concept. You want to have multiple connections to the Internet, provided by multiple ISPs. But multi-homing quickly gets very complicated, is difficult to configure, and is certainly not something you would want to reconfigure casually. Radware looked at this problem and developed an innovative application of traffic management.

For this discussion, let's assume a fairly simple multi-homing configuration: a branch office LAN connected to the Internet through two different ISPs. The obvious application of traffic management is simple life testing of the two links, and assuring that no traffic is sent to an ISP if a link is down. And you can imagine how a traffic manager could look at the load on the two links and balance it suitably.

But Radware goes well beyond this, using their DNS technology to determine which of the ISPs is the better path for specific traffic, and then routing traffic accordingly. This is clearly an innovative and clever use of traffic management, and certainly one we had never come close to imagining before. (See Radware's [white paper](#) for more interesting details.)

This kind of innovation is particularly important given a question we are regularly asked: "Won't the traffic management product category disappear over time as the functionality migrates into conventional routers and switches?" The answer we give is "Yes, if a traffic management company invents nothing new, then over time the value of that product will diminish." But we strongly feel that this is the wrong way to look at traffic management. In the server room, we see traffic management systems in effect becoming the operating system of the clustered computers that are serving out a return to centralized information systems. That's a big deal and a big future. And at the global level, we see the DNS solutions evolving into fairly full-blown content-directed routing schemes (as in the Akamai network), and that's also a very big deal. So the future of traffic management lies in innovation, and it's a significant and exciting future, if an unknown one.

https://www.radware.com/content/company/press/presscov/default.asp?_v=Read&document=2622

3/29/2004



3

Global Product Spotlight: Radware Linkproof



NetworkMagazine.com, December 1, 1999.

Radware's new load balancer maximizes backup Internet links

By David Greenfield

What's an easy way to strengthen an Internet hookup? Add a link to another upstream Internet provider. That might be smart planning, but it doesn't make for great accounting. Backup links sit idle most of the time, which means ISPs pay full tariffs for rarely used lines.

Radware (www.radware.com) thinks there's a better business solution. Its new LinkProof is the first load balancer to make running parallel links to the Internet easy and cost-effective. For starters, LinkProof optimally distributes traffic across multiple access lines. What's more, if a line or router fails, LinkProof rolls the traffic over to the backup connections.

That might not sound like such a big deal. After all, tweaking the Border Gateway Protocol 4 (BGP4) routing protocol can yield similar benefits. But not everyone runs BGP4, and those who do spend considerable time and expertise configuring the protocol. Finally, while BGP4 will switch to a backup link, the protocol won't let you weight your traffic distribution to maximize your connections. LinkProof will do all of that, and it doesn't require a Ph.D. to deploy.

Or so says Radware. Although there are plenty of users briefed on the product, nobody has tested it. What's more, none of these users are the second-tier ISPs that are supposed to adopt the product. Finally, because LinkProof only works with links on its subnet, the box can't distribute traffic loads across lines on other networks or offices.

Still, that's not stopping some major networkers from getting excited about the product. "On paper at least, LinkProof sounds like just what we want," says George Kurian, consultant of architecture and technology planning at Pacific Corp., a utility company in Portland, OR. Pacific currently runs its Internet access out of Portland, while paying for a backup link out of Salt Lake City, UT.

LinkProof, a modular box with two Ethernet or Fast Ethernet ports, sits between the firewall protecting the corporate backbone and in front of the routers connected to the Internet. At install-time, the network manager assigns a weight to each link that indicates the speed or cost of each line.

The rest of the configuration depends on the particular application. When load balancing incoming traffic, as is common with an e-commerce site, the LinkProof appears as the default DNS server. DNS queries from users looking to access the site are sent to LinkProof. It has IP addresses that are associated with each of the ISPs' links. LinkProof determines the optimum link based on latency and packet loss and then responds with the appropriate destination IP address.

When balancing outgoing Internet traffic, LinkProof is defined as the default router. It receives all outgoing packets and determines

https://www.radware.com/content/company/press/presscov/default.asp?_v=Read&document=2627

3/29/2004

the optimal link. LinkProof then changes the packet's source address to an address associated with an ISP's line and forwards the packet to the appropriate router.

So what happens in the event of a failure? LinkProof constantly monitors the health of each connection by testing the availability of up to 10 IP addresses along the path. If the address doesn't respond after some user-defined period of time, the traffic is directed to the alternative link. By default, the switch time is two seconds.

The key in both cases is Smart Network Address Translation (SmartNAT), which is the ability to reply with an IP address specific to a link. With SmartNAT, LinkProof insures that the client's responses return along the same link as the outgoing request. This enables LinkProof to account for traffic flowing in both directions when making a load-balancing decision. "Without SmartNAT, you don't get real load balancing," says Kurian.

Radware certainly isn't the only vendor in the load-balancing market. A number of other companies—including Alteon WebSystems, Foundry Networks, and F5 Networks—deliver products that distribute traffic across Web sites and firewalls.

However, they stumble when it comes to delivering SmartNAT capabilities. Alteon is close, but the implementation is too cumbersome, says Kurian. Foundry and F5 don't offer products with SmartNAT today. F5 says it will add the SmartNAT feature in the next release of BIG/Ip, which is expected to ship in December 1999. Foundry has not announced plans for releasing SmartNAT.

(7)

Radware Seeks Solutions to Easy-Access Problems

South China Morning Post, December 7, 1999.

INTERNET **South China Morning Post** EDITION

By Veronique Saunier

Continuous access to Web sites is at the core of every product developed by Radware, a small Israeli company that claims to be the second-largest vendor of Internet traffic-management solutions. "The Internet is cruel. For a company cashing on e-commerce, a down time of even one minute means lost business and lost customers who may never come back," said Yaron Danelli, Radware's vice-president of sales for Asia Pacific.

"Yet the Internet is vulnerable. Everything from traffic overload to a pulled Ethernet cable can make a Web server unavailable."

Maintaining Web sites to keep them up and running continuously has become a business in itself for many companies, including France Telecom Hebergement - the host of the prestigious Presidency of Republic of France site - or Spirit IP Web hosting.

These carriers guarantee 100 per cent availability and offer their customers financial compensation if their sites are down for even a few seconds.

The way they keep their promises without bankruptcy is by making every machine and circuit of the network redundant by ensuring if one machine breaks the other still operates.

They also place so-called "load balancers" at strategic points of the network to make the Internet as fluid and fast as possible.

Although the concept of load balancing is quite simple - it directs Internet traffic to the server that is less busy - Radware claims it pioneered the concept and has been perfecting it since the launch of Web Server Director (WSD) four years ago.

WSD won Radware top honours from several United States technical magazines for its management, configuration, and ability to act as both primary and secondary load balancer at once.

Follow-up products include WSD Pro, which supports multiple networks, WSD DS which dispatches traffic to the nearest server in

https://www.radware.com/content/company/press/presscov/default.asp?_v=Read&document=2626

3/29/2004

the case of distributed sites, and Cache Server Director which intercepts Web users' requests and directs them to the most available cache server.

High availability also has become a critical component of firewalls deployed across enterprise networks to provide secure connectivity for Internet, and Intranet and extranet communications.

Last year Radware launched FireProof, which load balances data to the best available firewall of the network.

According to Mr. Daniel, many Cisco Systems' firewalls are load balanced by Radware's FireProof.

"While we were installing FireProof, we realised many companies wished to have multiple connections to the Internet instead of relying on one single ISP but were not ready to go through the hassles," Sharon Trachtman, vice-president marketing, said.

Using multiple ISPs adds redundancy abilities but necessitates complex configuration and routing protocols as well as close coordination between the contracted ISPs.

To make things easier, Radware designed a dedicated product that determined the closest, fastest and healthiest route for incoming and out-bound IP traffic between different ISPs.

That product, LinkProof, was launched globally last month.

Such responsiveness is the key to Radware's success.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

23484 7590 05/26/2004
JOHN W. L. OGILVIE
1320 EAST LAIRD AVENUE
SALT LAKE CITY, UT 84105

EXAMINER
MARCELO, MELVIN C

ART UNIT PAPER NUMBER
2663 10

DATE MAILED: 05/26/2004

Table with 5 columns: APPLICATION NO., FILING DATE, FIRST NAMED INVENTOR, ATTORNEY DOCKET NO., CONFIRMATION NO. Values: 10/361,837, 02/07/2003, Sanchaita Datta, 3003.2.11A, 3645

TITLE OF INVENTION: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS

Table with 6 columns: APPLN TYPE, SMALL ENTITY, ISSUE FEE, PUBLICATION FEE, TOTAL FEE(S) DUE, DATE DUE. Values: nonprovisional, YES, \$665, \$300, \$965, 08/26/2004

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 1.51. THE ISSUE FEE DUE INDICATED ABOVE REFLECTS A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE APPLIED IN THIS APPLICATION. THE PTOL-85B (OR AN EQUIVALENT) MUST BE RETURNED WITHIN THIS PERIOD EVEN IF NO FEE IS DUE OR THE APPLICATION WILL BE REGARDED AS ABANDONED.

HOW TO REPLY TO THIS NOTICE:

I. Review the SMALL ENTITY status shown above.

If the SMALL ENTITY is shown as YES, verify your current SMALL ENTITY status:

A. If the status is the same, pay the TOTAL FEE(S) DUE shown above.

B. If the status is changed, pay the PUBLICATION FEE (if required) and twice the amount of the ISSUE FEE shown above and notify the United States Patent and Trademark Office of the change in status, or

If the SMALL ENTITY is shown as NO:

A. Pay TOTAL FEE(S) DUE shown above, or

B. If applicant claimed SMALL ENTITY status before, or is now claiming SMALL ENTITY status, check the box below and enclose the PUBLICATION FEE and 1/2 the ISSUE FEE shown above.

[] Applicant claims SMALL ENTITY status. See 37 CFR 1.27.

II. PART B - FEE(S) TRANSMITTAL should be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). Even if the fee(s) have already been paid, Part B - Fee(s) Transmittal should be completed and returned. If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Utility patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees. It is patentee's responsibility to ensure timely payment of maintenance fees when due.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: **Mail** **Mail Stop ISSUE FEE**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax **(703) 746-4000**

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 4 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Legibly mark-up with any corrections or use Block 1)

23484 7590 05/26/2004

JOHN W. L. OGILVIE
1320 EAST LAIRD AVENUE
SALT LAKE CITY, UT 84105

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission
 I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO, on the date indicated below.

_____ (Depositor's name)
_____ (Signature)
_____ (Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A	3645

TITLE OF INVENTION: **TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS**

APPLN TYPE	SMALL ENTITY	ISSUE FEE	PUBLICATION FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	YES	\$665	\$300	\$965	08/26/2004

EXAMINER	ART UNIT	CLASS-SUBCLASS
MARCELO, MELVIN C	2663	370-401000

<p>1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).</p> <p><input type="checkbox"/> Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.</p> <p><input type="checkbox"/> "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. Use of a Customer Number is required.</p>	<p>2. For printing on the patent front page, list (1) the names of up to 3 registered patent attorneys or agents OR, alternatively, (2) the name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.</p> <p>1 _____</p> <p>2 _____</p> <p>3 _____</p>
--	---

3. ASSIGNEE: NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. Inclusion of assignee data is only appropriate when an assignment has been previously submitted to the USPTO or is being submitted under separate cover. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE _____ (B) RESIDENCE: (CITY and STATE OR COUNTRY) _____

Please check the appropriate assignee category or categories (will not be printed on the patent); individual corporation or other private group entity government

<p>4a. The following fee(s) are enclosed:</p> <p><input type="checkbox"/> Issue Fee</p> <p><input type="checkbox"/> Publication Fee</p> <p><input type="checkbox"/> Advance Order - # of Copies _____</p>	<p>4b. Payment of Fee(s):</p> <p><input type="checkbox"/> A check in the amount of the fee(s) is enclosed.</p> <p><input type="checkbox"/> Payment by credit card. Form PTO-2038 is attached.</p> <p><input type="checkbox"/> The Director is hereby authorized by charge the required fee(s), or credit any overpayment, to Deposit Account Number _____ (enclose an extra copy of this form).</p>
---	---

Director for Patents is requested to apply the Issue Fee and Publication Fee (if any) or to re-apply any previously paid issue fee to the application identified above.

(Authorized Signature) _____ (Date) _____

NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

This collection of information is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, Alexandria, Virginia 22313-1450.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMIT THIS FORM WITH FEE(S)



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A	3645
23484	7590	05/26/2004	EXAMINER	
JOHN W. L. OGILVIE 1320 EAST LAIRD AVENUE SALT LAKE CITY, UT 84105			MARCELO, MELVIN C	
			ART UNIT	PAPER NUMBER
			2663	

DATE MAILED: 05/26/2004

10

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment to date is 0 day(s). If the issue fee is paid on the date that is three months after the mailing date of this notice and the patent issues on the Tuesday before the date that is 28 weeks (six and a half months) after the mailing date of this notice, the Patent Term Adjustment will be 0 day(s).

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) system (<http://pair.uspto.gov>).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (703) 305-1383. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at (703) 305-8283.

JUN-07-04 MON 09:48 AM

FAX NO.

RECEIVED
CENTRAL FAX CENTER
JUN 07 2004

TT # 11
7-8-04
TC

FACSIMILE MEMORANDUM

THORPE NORTH & WESTERN L.L.P.
SINCE 1979
THE TEAM APPROACH TO PREMIER PERFORMANCE ®

P.O. Box 1219
SANDY, UTAH 84091-1219
USA
TELEPHONE 1.801.566.8633
FACSIMILE 1.801.566.8673
FACSIMILE 1.801.566.0750

OFFICIAL

DATE: JUNE 7, 2004
TO: COMMISSIONER FOR PATENTS
FROM: JOHN W. L. OOLVIE
DOCKET NUMBER: MULTIPLE DOCKET NOS. FOR FATPIPE SYSTEMS AKA RAQUA SYSTEMS DEVELOPMENT COMPANY

PAGE 1 OF: 4
FACSIMILE NUMBER: (703)872-9308
TRANSMITTED BY: SHEILA

SUBJECT:
SUBSTITUTE POWER OF ATTORNEY AND CHANGE OF ADDRESS FOR CORRESPONDENCE

IF YOU DO NOT RECEIVE THE COMPLETE DOCUMENT, PLEASE NOTIFY THE SENDER AS SOON AS POSSIBLE.

REMARKS:

PLEASE ACKNOWLEDGE RECEIPT



THE PAGES THAT FOLLOW MAY CONTAIN SENSITIVE, PRIVILEGED OR CONFIDENTIAL INFORMATION INTENDED SOLELY FOR THE ADDRESSEE NAMED ABOVE. IF YOU RECEIVE THIS MESSAGE AND ARE NOT THE AGENT OR EMPLOYEE OF THE ADDRESSEE, AND HAVE THEREFORE BEEN SENT OR RECEIVED THIS COMMUNICATION IN ERROR, YOU ARE ASKED NOT TO DISSEMINATE OR COPY ANY OF THE ATTACHED AND ARE TO NOTIFY THE SENDER IMMEDIATELY BY TELEPHONE. PLEASE ALSO RETURN THE ORIGINAL MESSAGE TO THE SENDER BY MAIL. THANK YOU.

PAGE 14 * RCVD AT 6/7/2004 12:40:39 PM [Eastern Daylight Time] * SVR:USPTO-EFAXF-18 * CHG:3729308 * CSID: * DURATION (mm-ss):01-24

JUN-07-04 MON 09:48 AM

FAX NO.

P. 02

RECEIVED
CENTRAL FAX CENTER

JUN 07 2004

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

**REVOCATION OF POWER OF ATTORNEY
WITH NEW POWER OF ATTORNEY AND
CHANGE OF CORRESPONDENCE ADDRESS**

CERTIFICATE OF DEPOSIT DATE OF DEPOSIT: <u>6/7/04</u> I hereby certify that this paper or fee (along with any paper or fee referred to as being attached or enclosed) is being sent via facsimile No. (703)872-9306 to the Commissioner for Patents, on the date indicated above. <i>Shelia Halterman</i> Shelia Halterman	OFFICIAL
---	----------

**SUBSTITUTE POWER OF ATTORNEY AND
CHANGE OF ADDRESS FOR CORRESPONDENCE**

FatPipe Systems aka Ragula Systems Development Company, a corporation of the State of Utah, whose address is 4455 South 700 East, Suite 100, Salt Lake City, Utah, 84107, hereby appoints as its attorneys and/or patent agents the law firm of THORPE NORTH & WESTERN, LLP, having a business address of 8180 South 700 East, Suite 200, Sandy, Utah 84070, and VAUGHN W. NORTH, Registration No. 27,930; M. WAYNE WESTERN, Registration No. 22,788; CLIFTON W. THOMPSON, Registration No. 36,947; GARRON M. HOBSON, Registration No. 41,073; PETER M. DE JONGE, Registration No. 47,521; WEILI CHENG, Registration No. 44,609; DAVID R. MCKINNEY, Registration No. 42,868; STEVE M. FERRY, Registration No. 45,357; GARY P. OAKESON, Registration No. 44,266; DAVID W. OSBORNE, Registration No. 44,989; JASON R. JONES, Registration No. 51,008; BRIK S. ERICKSEN, Registration No. 48,954; JOHN W.L. OGILVIE, Registration No. 37,987; and CHRISTOPHER L. JOHNSON, Registration No. 46,809; all with full power of substitution and revocation, to prosecute applications and to transact all business in the Patent and Trademark Office connected with regard to the following:

013765 / 0738

Revocation Of Power Of Attorney
 With New Power Of Attorney And
 Change Of Correspondence Address
 Page 2

<u>DOCKET NO.</u>	<u>TITLE</u>	<u>SERIAL/PATENT NO.</u>
22807.NP (formerly 3003.2.1A)	System and Method for Transmitting a User's Data Packets Concurrently Over Different Telephone Lines Between Two Computer Networks	Pat. No. 6,253,247
22809 (formerly 3003.2.3)	Combining Routers to Increase Concurrency and Redundancy in External Network Access	Pat. No. 6,295,276
22814.NP (formerly 3003.2.8A)	Combining Routers to Increase Concurrency and Redundancy in External Network Access	Pat. No. 6,493,341
22814.CIP (formerly 3003.2.8B)	Combining Routers to Increase Concurrency and Redundancy in External Network Access	Serial No. 10/263,497
22973.NP (formerly 3003.2.11A)	Tools and Techniques for Directing Packets Over Disparate Networks	<u>Serial No. 10/361,837</u>
22972.NP (formerly 3003.2.10B)	Domain Name Resolution Making IP Address Selections in Response to Connection Status When Multiple Connections are Present	Serial No. 10/034,190
22971.NP (formerly 3003.2.9A)	Combining Connections for Parallel Access to Multiple Frame Relay and Other Private Networks.	Serial No. 10/034,197

JUN-07-04 MON 09:49 AM

FAX NO.

P. 04

Revocation Of Power Of Attorney
With New Power Of Attorney And
Change Of Correspondence Address
Page 3

All correspondence concerning this application should be directed to:

John W. L. Ogilvie
THORPE NORTH & WESTERN, LLP
Customer No. 20,551
P.O. Box 1219
Sandy, Utah 84091-1219
Telephone: (801) 566-6633
Facsimile: (801) 566-0750

All previous powers of attorney with regard to these matters are hereby revoked.

Dated this 11 day of May, 2004 at Salt Lake City, UT (City, State).

FatPipe Systems
Aka: Ragula Systems Development Company

By: Sanchaita Datla
Name: SANCHAITA DATTA
Title: V.P.

Transaction History Date 2004-06-14

Date information retrieved from USPTO Patent

Application Information Retrieval (PAIR)

system records at www.uspto.gov

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: Mail

Mail Stop ISSUE FEE
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (703) 746-4000

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 4 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Lightly mark-up with any corrections or use Block 1)

23484 7590 03/26/2004

JOHN W. L. OGILVIE
1320 EAST LAIRD AVENUE
SALT LAKE CITY, UT 84105



Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO, on the date indicated below.

Sheila Halterman (Depositor's name)
Sheila Halterman (Signature)
June 10, 2004 (Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A	3645

TITLE OF INVENTION: TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS 22973.NP

APPLN. TYPE	SMALL ENTITY	ISSUE FEE	PUBLICATION FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	YES	\$665	\$300	\$965	08/26/2004

EXAMINER	ART UNIT	CLASS-SUBCLASS
MARCELO, MELVIN C	2663	370-401000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).
 Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.
 "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. Use of a Customer Number is required.

2. For printing on the patent front page, list (1) the names of up to 3 registered patent attorneys or agents OR, alternatively, (2) the name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.
 1. Thorpe
 2. North &
 3. Western LLP

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)
 PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. Inclusion of assignee data is only appropriate when an assignment has been previously submitted to the USPTO or is being submitted under separate cover. Completion of this form is NOT a substitute for filing an assignment.
 (A) NAME OF ASSIGNEE: Ragula Systems
 (B) RESIDENCE: (CITY and STATE OR COUNTRY) Salt Lake City, Utah

Please check the appropriate assignee category or categories (will not be printed on the patent); individual corporation or other private group entity government

4a. The following fee(s) are enclosed:
 Issue Fee
 Publication Fee
 Advance Order - # of Copies 10

4b. Payment of Fee(s):
 A check in the amount of the fee(s) is enclosed.
 Payment by credit card. Form PTO-2038 is attached.
 The Director is hereby authorized by charge the required fee(s), or credit lay overpayment, to Deposit Account Number 20-0100 (enclose an extra copy of this form). deficiencies only

Director for Patents is requested to apply the Issue Fee and Publication Fee (if any) or to re-apply any previously paid issue fee to the application identified above.

(Authorized Signature) [Signature] (Date) 10 June 2004

NOTE: The Issue Fee and Publication Fee (if required) will not be accepted from anyone other than the applicant; a registered attorney or agent; or the assignee or other party in interest as shown by the records of the United States Patent and Trademark Office.

This collection of information is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, Alexandria, Virginia 22313-1450.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

06/16/2004 JBA/LMM 00000120 10361837
 01 FC:2501 665.00 DP
 02 FC:1504 300.00 DP
 03 FE:8001 30.00 DP

TRANSMIT THIS FORM WITH FEE(S)



TT #

Approved for use through 11/30/2005. OMB 0851-0035
 U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**CHANGE OF
 CORRESPONDENCE ADDRESS
 Application**

Address to:
 Commissioner for Patents
 P.O. Box 1450
 Alexandria, VA 22313-1450.

Application Number	10/361,837
Filing Date	February 7, 2003
First Named Inventor	Sanchata Delta
Art Unit	2663
Examiner Name	Marcelo, Melvin C.
Attorney Docket Number	22973.NP

Please change the Correspondence Address for the above-identified patent application to:

Customer Number :

OR

Firm or Individual Name

Address

Address

City State Zip

Country

Telephone Fax

This form cannot be used to change the data associated with a Customer Number. To change the data associated with an existing Customer Number use "Request for Customer Number Data Change" (PTO/SB/124).

I am the:

- Applicant/Inventor
- Assignee of record of the entire interest. Statement under 37 CFR 3.73(b) is enclosed. (Form PTO/SB/96).
- Attorney or Agent of record. Registration Number 37,987
- Registered practitioner named in the application transmittal letter in an application without an executed oath or declaration. See 37 CFR 1.33(a)(1). Registration Number _____

Typed or Printed Name John W. L. Ogilvie

Signature

Date June 10, 2004 Telephone (801)566-6633

NOTE: Signatures of all the inventors or assignee of record of the entire interest or their representative(s) are required. Submit multiple forms if more than one signature is required (see below).

Total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.


If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.



IN THE UNITED STATES PATENT & TRADEMARKS OFFICE

ART UNIT: 2663
EXAMINER: Melvin C. Marcelo
APPLICANT: Sanchaita Datta
SERIAL NO.: 10/361,837
FILED: February 7, 2003
CONFRM. NO.: 3645
FOR: TOOLS AND TECHNIQUES FOR
DIRECTING PACKETS OVER
DISPARATE NETWORKS
ATTORNEY DOCKET NO. 22973.NP

**REQUEST FOR
DESIGNATION OF PATENT
DRAWING**



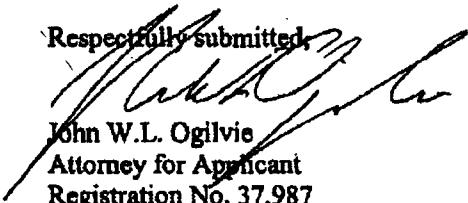
Mail Stop Issue Fee
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

We have today paid the issue fee in the above-referenced application and respectfully request that Figure 7 be designated when the patent is published.

Dated this 10th day of June, 2004.

Respectfully submitted,


John W.L. Ogilvie
Attorney for Applicant
Registration No. 37,987

THORPE NORTH & WESTERN, LLP
Customer No. 20,551
P.O. Box 1219
Sandy, Utah 84091-1219
Telephone: (801) 566-6633

JWO/sbh



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371 (c) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A

CONFIRMATION NO. 3645



20551
THORPE NORTH & WESTERN, LLP.
8180 SOUTH 700 EAST, SUITE 200
P.O. BOX 1219
SANDY, UT 84070

Date Mailed: 07/09/2004

NOTICE OF ACCEPTANCE OF POWER OF ATTORNEY

This is in response to the Power of Attorney filed 06/07/2004.

The Power of Attorney in this application is accepted. Correspondence in this application will be mailed to the above address as provided by 37 CFR 1.33.

Timothy Callwell

Fw

JANA ROBBINS
OPFD 0-

OFFICE COPY



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371 (c) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
10/361,837	02/07/2003	Sanchaita Datta	3003.2.11A

23484
JOHN W. L. OGILVIE
1320 EAST LAIRD AVENUE
SALT LAKE CITY, UT 84105

CONFIRMATION NO. 3645



#12

Date Mailed: 07/09/2004

NOTICE REGARDING CHANGE OF POWER OF ATTORNEY

This is in response to the Power of Attorney filed 06/07/2004.

- The Power of Attorney to you in this application has been revoked by the assignee who has intervened as provided by 37 CFR 3.71. Future correspondence will be mailed to the new address of record(37 CFR 1.33).

Timothy Caldwell
For JANA ROBBINS
OPPD 0-

OFFICE COPY

File History Content Report

The following content is missing from the original file history record obtained from the United States Patent and Trademark Office. No additional information is available.

Document Date - 2004-08-10

Document Title - USPTO Grant

Table of Contents

1. US6775235B2 Tools and techniques for directing packets over disparate networks
-

Family 1/1

12 record(s) per family, collapsed by 9 record(s)

Record 1/9 US6295276B1 Combining routers to increase concurrency and redundancy in external network access

Publication Number:

US6295276B1 20010925

Title:

Combining routers to increase concurrency and redundancy in external network access

Title - DWPI:

Routers combining method for connecting local area network to Internet, involves selecting router and responding its address which does not necessarily include Internet protocol address

Priority Number:

US1999476646A

Priority Date:

1999-12-31

Application Number:

US1999476646A

Application Date:

1999-12-31

Publication Date:

2001-09-25

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
G06F001300	G	G06	G06F	G06F0013	G06F001300
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001266	H	H04	H04L	H04L0012	H04L001266

Assignee/Applicant:

Ragula Systems, Salt Lake City, UT

JP F Terms:

JP FI Codes:

Assignee - Original:

Ragula Systems

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 29/12018	-	20130101	EP
Current	H04L 29/12009		20130101	EP
Current	H04L 29/12783		20130101	EP
Current	H04L 45/00		20130101	EP
Current	H04L 45/22		20130101	EP
Current	H04L 45/58		20130101	EP
Current	H04L 61/10		20130101	EP
Current	H04L 61/35		20130101	EP
Current	H04L 69/16		20130101	EP
Current	H04L 69/163		20130101	EP

ECLA:

H04L002912A1 | H04L002906J | H04L002906J7 | H04L002912A | H04L002912A6 |
H04L004500 | H04L004522 | H04L004558 | H04L006110 | H04L006135

Abstract:

Methods, configured storage media, and systems are provided for increasing bandwidth between a local area network ("LAN") and other networks by using multiple routers on the given LAN. Data packets are multiplexed between the routers using a novel variation on the standard address resolution protocol, and other components. On receiving data destined for an external network, a controller or gateway computer will direct the data to the appropriate router. In addition to providing higher speed connections, the invention provides better fault tolerance in the form of redundant connections from the originating LAN to a wide area network such as the Internet.

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2012-11-05	FPAY	+
Description: FEE PAYMENT		
2009-06-18	AS	-
Description: ASSIGNMENT FATPIPE NETWORKS INDIA LIMITED, INDIA ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS); REEL/FRAME:022835/0994 2009-06-18		
2008-09-29	FPAY	+
Description: FEE PAYMENT		
2005-03-25	FPAY	+
Description: FEE PAYMENT		
2001-03-05	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS, UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:011595/0460 2001-01-19		
1999-12-31	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS, UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:010494/0173 1999-12-30		

Post-Issuance (US):

Reassignment (US) Table:

Assignee	Assignor	Date Signed	Reel/Frame	Date
FATPIPE NETWORKS INDIA LIMITED,CHENNAI, STATE OF TAMILNADU,IN	DATTA, SANCHAITA	2009-06-18	022835/0994	2009-06-18
	BHASKAR, RAGULA	2009-06-18		
	RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS)	2009-06-18		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN OGILVIE 2552 SOUTH WILSHIRE CIRCLE SALT LAKE CITY, UT 84109				
RAGULA SYSTEMS,SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2001-01-19	011595/0460	2001-03-05
	BHASKAR, RAGULA	2001-01-19		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN W. L. OGILVIE COMPUTER LAW+ 1211 EAST YALE AVENUE SALT LAKE CITY, UTAH 84105				
RAGULA SYSTEMS,SALT LAKE CITY,UT,US	DATTA, SANCHAITA	1999-12-30	010494/0173	1999-12-31
	BHASKAR, RAGULA	1999-12-30		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN W.L. OGILVIE COMPUTER LAW ++ 8 EAST BROADWAY SUITE 725 SALT LAKE CITY, UT 84111				

Maintenance Status (US):


Litigation (US):

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Record 2/9 US20020087724A1 Combining connections for parallel access to multiple frame relay and other private networks

Publication Number:

US20020087724A1 20020704

Title:

Combining connections for parallel access to multiple frame relay and other private networks

Title - DWPI:

Independent private networks access controller receives data packets through site interface and sends them to private network through selected private network interface

Priority Number:

US2000259269P

Priority Date:

2000-12-29

Application Number:

US200134197A

Application Date:

2001-12-28

Publication Date:

2002-07-04

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002908	H	H04	H04L	H04L0029	H04L002908

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002908	H	H04	H04L	H04L0029	H04L002908
G06F0015173	G	G06	G06F	G06F0015	G06F0015173

Assignee/Applicant:

Ragula Systems d/b/a FatPipe Networks

JP F Terms:**JP FI Codes:****Assignee - Original:**

Ragula Systems d/b/a FatPipe Networks

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 12/5692	H04L 63/0272	20130101	EP
Current	H04L 45/24	H04L 63/18	20130101	EP
Current	H04L 47/125	H04L 67/1002	20130101	EP
Current		H04L 2029/06054	20130101	EP

ECLA:

H04L001256F1 | H04L004524 | H04L004712B

Abstract:

Methods, configured storage media, and systems are provided for communications using two or more frame relay or point-to-point networks in parallel to provide load balancing across network connections, greater reliability, and/or increased security. A controller provides access to two or more private networks in parallel, through direct or indirect network interfaces. When one attached network fails, the failure is sensed by the controller and traffic is routed through one or more other private networks. When all attached networks are operating, the controller preferably balances the load between them.

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2009-06-18	AS	-
Description: ASSIGNMENT FATPIPE NETWORKS INDIA LIMITED, INDIA ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS); REEL/FRAME:022835/0994 2009-06-18		
2002-01-25	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS (FATPIPE NETWORKS), UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:012338/0729 2002-01-24		

2002-01-25	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS (FATPIPE NETWORKS), UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:012338/0707 2002-01-24		
2002-01-25	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS (FATPIPE NETWORKS), UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:012338/0724 2002-01-24		

Post-Issuance (US):

Reassignment (US) Table:

Assignee	Assignor	Date Signed	Reel/Frame	Date
FATPIPE NETWORKS INDIA LIMITED,CHENNAI, STATE OF TAMILNADU,IN	DATTA, SANCHAITA	2009-06-18	022835/0994	2009-06-18
	BHASKAR, RAGULA	2009-06-18		
	RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS)	2009-06-18		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN OGILVIE 2552 SOUTH WILSHIRE CIRCLE SALT LAKE CITY, UT 84109				
RAGULA SYSTEMS (FATPIPE NETWORKS),SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2002-01-24	012338/0729	2002-01-25
	BHASKAR, RAGULA	2002-01-24		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: COMPUTER LAW++ JOHN W.L. OGILVIE 1211 EAST YALE AVENUE SALT LAKE CITY, UT 84105				
RAGULA SYSTEMS (FATPIPE NETWORKS),SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2002-01-24	012338/0724	2002-01-25
	BHASKAR, RAGULA	2002-01-24		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: COMPUTER LAW++ JOHN W.L. OGILVIE 1211 EAST YALE AVENUE SALT LAKE CITY, UT 84105				
RAGULA SYSTEMS (FATPIPE NETWORKS),SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2002-01-24	012338/0707	2002-01-25
	BHASKAR, RAGULA	2002-01-24		

Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).

Correspondent: COMPUTER LAW++ JOHN W.L. OGIVIE 1211 EAST YALE AVENUE SALT LAKE CITY, UT 84105

Maintenance Status (US):


Litigation (US):

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Record 3/9 US6493341B1 Combining routers to increase concurrency and redundancy in external network access

Publication Number:

US6493341B1 20021210

Title:

Combining routers to increase concurrency and redundancy in external network access

Title - DWPI:

Routers combining method for local area network, involves selecting specific router such that consequent use of selected router increases concurrent operation of several other identified routers

Priority Number:

US1999174114P | US1999476646A

Priority Date:

1999-12-31 | 1999-12-31

Application Number:

US2000751590A

Application Date:

2000-12-29

Publication Date:

2002-12-10

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912
G06F001300	G	G06	G06F	G06F0013	G06F001300
H04L001246	H	H04	H04L	H04L0012	H04L001246

Assignee/Applicant:

Ragula Systems, Salt Lake City, UT

JP F Terms:

JP FI Codes:

Assignee - Original:

Ragula Systems

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 29/12009	-	20130101	EP
Current	H04L 12/5692		20130101	EP
Current	H04L 29/12018		20130101	EP
Current	H04L 29/12783		20130101	EP
Current	H04L 45/00		20130101	EP
Current	H04L 45/04		20130101	EP
Current	H04L 45/22		20130101	EP
Current	H04L 45/28		20130101	EP
Current	H04L 45/58		20130101	EP
Current	H04L 47/125		20130101	EP
Current	H04L 61/10		20130101	EP
Current	H04L 61/35		20130101	EP
Current	H04L 69/16		20130101	EP
Current	H04L 69/163		20130101	EP

ECLA:

H04L002912A | H04L001256F1 | H04L002906J | H04L002906J7 | H04L002912A1 | H04L002912A6 | H04L004500 | H04L004504 | H04L004522 | H04L004528 | H04L004558 | H04L004712B | H04L006110 | H04L006135

Abstract:

Methods, configured storage media, and systems are provided for increasing bandwidth between a local area network ("LAN") and other networks by using multiple routers on the given LAN. Data packets are multiplexed between the routers using a novel variation on the standard SYN packet synchronization protocol, and other components. On receiving data destined for an external network, a controller or gateway computer will direct the data to the appropriate router. In addition to providing higher speed connections, the invention provides better fault tolerance in the form of redundant connections from the originating LAN to a wide area network such as the Internet.

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2014-07-29	SULP	+

Description: SURCHARGE FOR LATE PAYMENT		
2014-07-29	FPAY	+
Description: FEE PAYMENT		
2014-07-18	REMI	-
Description: MAINTENANCE FEE REMINDER MAILED		
2009-12-17	FPAY	+
Description: FEE PAYMENT		
2009-06-18	AS	-
Description: ASSIGNMENT FATPIPE NETWORKS INDIA LIMITED, INDIA ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS); REEL/FRAME:022835/0994 2009-06-18		
2005-12-12	FPAY	+
Description: FEE PAYMENT		
2001-03-05	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS, UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:011595/0460 2001-01-19		

Post-Issuance (US):

Reassignment (US) Table:

Assignee	Assignor	Date Signed	Reel/Frame	Date
FATPIPE NETWORKS INDIA LIMITED,CHENNAI, STATE OF TAMILNADU,IN	DATTA, SANCHAITA	2009-06-18	022835/0994	2009-06-18
	BHASKAR, RAGULA	2009-06-18		
	RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS)	2009-06-18		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN OGILVIE 2552 SOUTH WILSHIRE CIRCLE SALT LAKE CITY, UT 84109				
RAGULA SYSTEMS,SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2001-01-19	011595/0460	2001-03-05

Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).

Corresponent: JOHN W. L. OGILVIE COMPUTER LAW+ 1211 EAST YALE AVENUE SALT LAKE CITY, UTAH 84105

Maintenance Status (US):


Litigation (US):

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Record 4/9 WO2003067817A1 TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS | OUTILS ET TECHNIQUES POUR DIRIGER DES PAQUETS SUR DES RESEAUX DISPARATES

Publication Number:

WO2003067817A1 20030814

Title:

TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS | OUTILS ET TECHNIQUES POUR DIRIGER DES PAQUETS SUR DES RESEAUX DISPARATES

Title - DWPI:

Controller for combining routers for connecting local area network to internet, provides modified synchronization requests containing physical and IP addresses of selected router

Priority Number:

US2002355509P | US2002263497A

Priority Date:

2002-02-08 | 2002-10-02

Application Number:

WO2003US3988A

Application Date:

2003-02-07

Publication Date:

2003-08-14

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001254	H	H04	H04L	H04L0012	H04L001254
H04L0012703	H	H04	H04L	H04L0012	H04L0012703
H04L0012707	H	H04	H04L	H04L0012	H04L0012707
H04L0012715	H	H04	H04L	H04L0012	H04L0012715
H04L0012803	H	H04	H04L	H04L0012	H04L0012803
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002908	H	H04	H04L	H04L0029	H04L002908

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group -	Subgroup - DWPI
------------	----------------	--------------	-----------------	---------------	-----------------

				DWPI	
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912
H04L001226	H	H04	H04L	H04L0012	H04L001226
G01R003108	G	G01	G01R	G01R0031	G01R003108
H04L0012288	H	H04	H04L	H04L0012	H04L0012288

Assignee/Applicant:

RAGULA SYSTEMS (d/b/a/ FatPipe Networks) (a Utah corporation),US

JP F Terms:

JP FI Codes:

Assignee - Original:

RAGULA SYSTEMS (d/b/a/ FatPipe Networks) (a Utah corporation)

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 63/0272	H04L 67/1002	20130101	EP
Current	H04L 12/4641	H04L 2029/06054	20130101	EP
Current	H04L 12/5692		20130101	EP
Current	H04L 45/04		20130101	EP
Current	H04L 45/22		20130101	EP
Current	H04L 45/28		20130101	EP
Current	H04L 47/125		20130101	EP
Current	H04L 67/1008		20130101	EP
Current	H04L 67/101		20130101	EP
Current	H04L 67/1017		20130101	EP
Current	H04L 67/1034		20130101	EP
Current	H04L 69/16		20130101	EP
Current	H04L 69/163		20130101	EP

ECLA:

H04L006302C | H04L001246V | H04L001256F1 | H04L002906J | H04L002906J7 |
H04L002908N9A1B | H04L002908N9A1C | H04L002908N9A1F | H04L002908N9A11 |
H04L004504 | H04L004522 | H04L004528 | H04L004712B | T04L002906C8A | T04L002908N9A

Abstract:

Methods, configured storage media, and systems are provided for communications using two or more disparate networks (106, 500) in parallel to provide load balancing across network connections, greater reliability, and/or increased security. A controller (602) provides access to two or more disparate networks in parallel, through direct or indirect network interfaces (706). When one attached network fails, the failure is sensed by the controller and traffic is routed through one or more other disparate networks. When all attached disparate networks are operating, one controller preferably balances (910) the load between them.

La présente invention concerne des procédés, des supports de stockage configurés et des systèmes permettant d'établir des communications à l'aide d'au moins deux réseaux disparates (106, 500) en parallèle afin d'obtenir un équilibrage des charges à travers des connexions réseau, une plus grande fiabilité, et/ou une sécurité augmentée. Un contrôleur (602) fournit l'accès à au moins deux réseaux disparates en parallèle, à travers des interfaces réseau (706) directes ou indirectes. Lorsqu'un réseau associé est défaillant, la défaillance est détectée par le contrôleur et le trafic est acheminé à travers un ou plusieurs autres réseaux disparates. Lorsque tous les réseaux disparates associés fonctionnent, de préférence, un contrôleur équilibre (910) la charge entre eux.

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2006-06-15	WWW	-
Description: WIPO INFORMATION: WITHDRAWN IN NATIONAL OFFICE		
2006-06-15	NENP	-
Description: NON-ENTRY INTO THE NATIONAL PHASE IN: JP		
2005-03-16	122	-
Description: EP: PCT APP. NOT ENT. EUROP. PHASE		
2003-12-04	DFPE	-
Description: REQUEST FOR PRELIMINARY EXAMINATION FILED PRIOR TO EXPIRATION OF 19TH MONTH FROM PRIORITY DATE (PCT APPLICATION FILED BEFORE 20040101)		
2003-10-08	121	-
Description: EP: THE EPO HAS BEEN INFORMED BY WIPO THAT EP WAS DESIGNATED IN THIS APPLICATION		
2003-08-14	AL	+
Description: DESIGNATED COUNTRIES FOR REGIONAL PATENTS WO 03067817 A1 GH; GM; KE; LS; MW; MZ; SD; SL; SZ; TZ; UG; ZM; ZW; AM; AZ; BY; KG; KZ; MD; RU; TJ; TM; AT; BE; BG; CH; CY; CZ; DE; DK; EE; ES; FI; FR; GB; GR;		

HU; IE; IT; LU; MC; NL; PT; SE; SI; SK; TR; BF; BJ; CF; CG; CI; CM; GA; GN; GQ; GW; ML; MR; NE; SN; TD; TG

2003-08-14

AK

+

Description: DESIGNATED STATES WO 03067817 A1 AE; AG; AL; AM; AT; AU; AZ; BA; BB; BG; BR; BY; BZ; CA; CH; CN; CO; CR; CU; CZ; DE; DK; DM; DZ; EC; EE; ES; FI; GB; GD; GE; GH; GM; HR; HU; ID; IL; IN; IS; JP; KE; KG; KP; KR; KZ; LC; LK; LR; LS; LT; LU; LV; MA; MD; MG; MK; MN; MW; MX; MZ; NO; NZ; OM; PH; PL; PT; RO; RU; SC; SD; SE; SG; SK; SL; TJ; TM; TN; TR; TT; TZ; UA; UG; UZ; VC; VN; YU; ZA; ZM; ZW

Post-Issuance (US):

Reassignment (US) Table:

Maintenance Status (US):


Litigation (US):

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Record 5/9 AU2003219731A1 TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS

Publication Number:

AU2003219731A1 20030902

Title:

TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS

Title - DWPI:

Controller for combining routers for connecting local area network to internet, provides modified synchronization requests containing physical and IP addresses of selected router

Priority Number:

US2002355509P | US2002263497A | WO2003US3988A

Priority Date:

2002-02-08 | 2002-10-02 | 2003-02-07

Application Number:

AU2003219731A

Application Date:

2003-02-07

Publication Date:

2003-09-02

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001254	H	H04	H04L	H04L0012	H04L001254
H04L0012703	H	H04	H04L	H04L0012	H04L0012703
H04L0012707	H	H04	H04L	H04L0012	H04L0012707
H04L0012715	H	H04	H04L	H04L0012	H04L0012715
H04L0012803	H	H04	H04L	H04L0012	H04L0012803
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002908	H	H04	H04L	H04L0029	H04L002908

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
H04L001228	H	H04	H04L	H04L0012	H04L001228

H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912
H04L001226	H	H04	H04L	H04L0012	H04L001226
G01R003108	G	G01	G01R	G01R0031	G01R003108
H04L0012288	H	H04	H04L	H04L0012	H04L0012288

Assignee/Applicant:

RAGULA SYSTEMS D B A FATPIPE N

JP F Terms:

JP FI Codes:

Assignee - Original:

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 63/0272	H04L 67/1002	20130101	EP
Current	H04L 12/4641	H04L 2029/06054	20130101	EP
Current	H04L 12/5692		20130101	EP
Current	H04L 45/04		20130101	EP
Current	H04L 45/22		20130101	EP
Current	H04L 45/28		20130101	EP
Current	H04L 47/125		20130101	EP
Current	H04L 67/1008		20130101	EP
Current	H04L 67/101		20130101	EP
Current	H04L 67/1017		20130101	EP
Current	H04L 67/1034		20130101	EP
Current	H04L 69/16		20130101	EP
Current	H04L 69/163		20130101	EP

ECLA:

H04L006302C | H04L001246V | H04L001256F1 | H04L002906J | H04L002906J7 |
H04L002908N9A1B | H04L002908N9A1C | H04L002908N9A1F | H04L002908N9A11 |
H04L004504 | H04L004522 | H04L004528 | H04L004712B | T04L002906C8A | T04L002908N9A

Abstract:

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2004-11-18	MK6	-
Description: APPLICATION LAPSED SECTION 142(2)(F)/REG. 8.3(3) - PCT APPLIC. NOT ENTERING NATIONAL PHASE		

Post-Issuance (US):

Reassignment (US) Table:

Maintenance Status (US):


Litigation (US):

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Record 6/9 US6775235B2 Tools and techniques for directing packets over disparate networks

Publication Number:

US6775235B2 20040810
 US20030147408A1 20030807

Title:

Tools and techniques for directing packets over disparate networks

Title - DWPI:

Disparate network access controller e.g. for frame relay network, has packet path selector to select network interface through which received packet is sent to virtual private network

Priority Number:

US2000259269P | US200134197A | US2002355509P

Priority Date:

2000-12-29 | 2001-12-28 | 2002-02-08

Application Number:

US2003361837A

Application Date:

2003-02-07

Publication Date:

2004-08-10

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002908	H	H04	H04L	H04L0029	H04L002908

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L001264	H	H04	H04L	H04L0012	H04L001264

Assignee/Applicant:

Ragula Systems, Salt Lake City, UT

JP F Terms:**JP FI Codes:****Assignee - Original:**

Ragula Systems

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 63/0272	H04L 63/18	20130101	EP
Current	H04L 12/2854	H04L 67/1002	20130101	EP
Current	H04L 12/2856	H04L 2029/06054	20130101	EP
Current	H04L 12/2859		20130101	EP
Current	H04L 12/2898		20130101	EP
Current	H04L 12/5692		20130101	EP
Current	H04L 43/0811		20130101	EP
Current	H04L 45/04		20130101	EP
Current	H04L 45/22		20130101	EP
Current	H04L 45/24		20130101	EP
Current	H04L 45/28		20130101	EP
Current	H04L 47/125		20130101	EP
Current	H04L 63/10		20130101	EP
Current	H04L 63/1466		20130101	EP
Current	H04L 69/16		20130101	EP

ECLA:

T04L002906C8A | T04L002908N9A | T04L006302C | T04L006318

Abstract:

Methods, configured storage media, and systems are provided for communications using two or more disparate networks in parallel to provide load balancing across network connections, greater reliability, and/or increased security. A controller provides access to two or more disparate networks in parallel, through direct or indirect network interfaces. When one attached network fails, the failure is sensed by the controller and traffic is routed through one or more other disparate networks. When all attached disparate networks are operating, one controller preferably balances the load between them.

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2011-08-19	FPAY	+

Description: FEE PAYMENT		
2009-06-18	AS	-
Description: ASSIGNMENT FATPIPE NETWORKS INDIA LIMITED, INDIA ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS); REEL/FRAME:022835/0994 2009-06-18		
2007-08-18	FPAY	+
Description: FEE PAYMENT		
2003-04-01	AS	-
Description: ASSIGNMENT SHAREWAVE, INC., TEXAS ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNOR:GUBBI, RAJUGOPAL R.; REEL/FRAME:013911/0312 1998-09-02		
2003-02-07	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS (FATPIPE NETWORK), UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:013765/0738 2003-02-05		

Post-Issuance (US):

Reassignment (US) Table:

Assignee	Assignor	Date Signed	Reel/Frame	Date
FATPIPE NETWORKS INDIA LIMITED,CHENNAI, STATE OF TAMILNADU,IN	DATTA, SANCHAITA	2009-06-18	022835/0994	2009-06-18
	BHASKAR, RAGULA	2009-06-18		
	RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS)	2009-06-18		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN OGILVIE 2552 SOUTH WILSHIRE CIRCLE SALT LAKE CITY, UT 84109				
RAGULA SYSTEMS (FATPIPE NETWORK),SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2003-02-05	013765/0738	2003-02-07
	BHASKAR, RAGULA	2003-02-05		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN W.L. OGILVIE 1211 EAST YALE AVENUE SALT LAKE CITY, UT 84105				
SHAREWAVE INC.,AUSTIN,TX,US	GUBBI, RAJUGOPAL R.	1998-09-02	013911/0312	2003-04-01

Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).

Correspondent: WINSTEAD SECHREST & MINICK P.C. JAMES J. MURPHY P.O. BOX 50784 1201 MAIN STREET DALLAS, TX 75250-0784

Maintenance Status (US):

Litigation (US):


2015-05-06 2015 Fatpipe, Inc. Talari Networks, Inc. E.D. Texas 6:15cv00458

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Record 7/9 US7269143B2 Combining routers to increase concurrency and redundancy in external network access

Publication Number:

US7269143B2 20070911
 US20030031180A1 20030213

Title:

Combining routers to increase concurrency and redundancy in external network access

Title - DWPI:

Controller for combining routers for connecting local area network to internet, provides modified synchronization requests containing physical and IP addresses of selected router

Priority Number:

US1999174114P | US1999476646A | US2000751590A

Priority Date:

1999-12-31 | 1999-12-31 | 2000-12-29

Application Number:

US2002263497A

Application Date:

2002-10-02

Publication Date:

2007-09-11

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002912	H	H04	H04L	H04L0029	H04L002912

H04L001226	H	H04	H04L	H04L0012	H04L001226
G01R003108	G	G01	G01R	G01R0031	G01R003108
H04L0012288	H	H04	H04L	H04L0012	H04L0012288

Assignee/Applicant:

Ragula Systems (FatPipe Networks), Salt Lake City, UT, US

JP F Terms:

JP FI Codes:

Assignee - Original:

Ragula Systems (FatPipe Networks)

Any CPC Table:

Type	Invention	Additional	Version	Office
Current	H04L 45/00	-	20130101	EP
Current	H04L 12/5692		20130101	EP
Current	H04L 29/12009		20130101	EP
Current	H04L 29/12018		20130101	EP
Current	H04L 29/12783		20130101	EP
Current	H04L 45/04		20130101	EP
Current	H04L 45/22		20130101	EP
Current	H04L 45/28		20130101	EP
Current	H04L 45/58		20130101	EP
Current	H04L 47/125		20130101	EP
Current	H04L 61/10		20130101	EP
Current	H04L 61/35		20130101	EP
Current	H04L 69/16		20130101	EP
Current	H04L 69/163		20130101	EP

ECLA:

H04L004500 | H04L001256F1 | H04L002906J | H04L002906J7 | H04L002912A |
H04L002912A1 | H04L002912A6 | H04L004504 | H04L004522 | H04L004528 | H04L004558 |
H04L004712B | H04L006110 | H04L006135

Abstract:

A controller is provided for increasing bandwidth between a local area network ("LAN") and other networks by using multiple routers on the given LAN. Data packets are multiplexed between the routers using a novel variation on the standard SYN packet synchronization protocol, and other components. On receiving data destined for an external network, the controller or gateway computer will direct the data to the appropriate router. In addition to providing higher speed connections, the invention provides better fault tolerance in the form of redundant connections from the originating LAN to a wide area network such as the Internet.

Language of Publication:

EN

INPADOC Legal Status Table:

Gazette Date	Code	INPADOC Legal Status Impact
2014-09-12	FPAY	+
Description: FEE PAYMENT		
2010-09-17	FPAY	+
Description: FEE PAYMENT		
2009-06-18	AS	-
Description: ASSIGNMENT FATPIPE NETWORKS INDIA LIMITED, INDIA ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS); REEL/FRAME:022835/0994 2009-06-18		
2002-10-02	AS	-
Description: ASSIGNMENT RAGULA SYSTEMS (FATPIPE NETWORKS), UTAH ASSIGNMENT OF ASSIGNORS INTEREST; ASSIGNORS:DATTA, SANCHAITA; BHASKAR, RAGULA; REEL/FRAME:013365/0351 2002-09-23		

Post-Issuance (US):

Reassignment (US) Table:

Assignee	Assignor	Date Signed	Reel/Frame	Date
FATPIPE NETWORKS INDIA LIMITED,CHENNAI, STATE OF TAMILNADU,IN	DATTA, SANCHAITA	2009-06-18	022835/0994	2009-06-18
	BHASKAR, RAGULA	2009-06-18		
	RAGULA SYSTEMS (D/B/A/ FATPIPE NETWORKS)	2009-06-18		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN OGILVIE 2552 SOUTH WILSHIRE CIRCLE SALT LAKE CITY, UT 84109				
RAGULA SYSTEMS (FATPIPE NETWORKS),SALT LAKE CITY,UT,US	DATTA, SANCHAITA	2002-09-23	013365/0351	2002-10-02
	BHASKAR, RAGULA	2002-09-23		
Conveyance: ASSIGNMENT OF ASSIGNORS INTEREST (SEE DOCUMENT FOR DETAILS).				
Corresponent: JOHN W.L. OGILVIE 1211 EAST YALE AVENUE COMPUTER LAW SALT LAKE CITY, UT 84105				

Maintenance Status (US):

Litigation (US):


2009-03-02 2009 Fatpipe Networks India Xroads networks Utah 2:09cv186

Opposition (EP):

License (EP):

EPO Procedural Status:

Front Page Drawing:



Publication Number:

US7406048B2 20080729
 US20050008017A1 20050113

Title:

Tools and techniques for directing packets over disparate networks

Title - DWPI:

Disparate network e.g. virtual private network, access controller, has packet path selector selecting path between network interfaces using two known location address ranges associated with disparate networks

Priority Number:

US2000259269P | US200134197A | US2002355509P | US2003361837A

Priority Date:

2000-12-29 | 2001-12-28 | 2002-02-08 | 2003-02-07

Application Number:

US2004911846A

Application Date:

2004-08-03

Publication Date:

2008-07-29

IPC Class Table:

IPC	Section	Class	Subclass	Class Group	Subgroup
H04L001264	H	H04	H04L	H04L0012	H04L001264
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256
H04L002906	H	H04	H04L	H04L0029	H04L002906
H04L002908	H	H04	H04L	H04L0029	H04L002908

IPC Class Table - DWPI:

IPC - DWPI	Section - DWPI	Class - DWPI	Subclass - DWPI	Class Group - DWPI	Subgroup - DWPI
H04L001228	H	H04	H04L	H04L0012	H04L001228
H04L001246	H	H04	H04L	H04L0012	H04L001246
H04L001256	H	H04	H04L	H04L0012	H04L001256

USPTO Maintenance Report

Patent Bibliographic Data			05/15/2015 03:17 PM		
Patent Number:	6775235		Application Number:	10361837	
Issue Date:	08/10/2004		Filing Date:	02/07/2003	
Title:	TOOLS AND TECHNIQUES FOR DIRECTING PACKETS OVER DISPARATE NETWORKS				
Status:	12th year fee window opens: 08/10/2015			Entity:	SMALL
Window Opens:	N/A	Surcharge Date:	N/A	Expiration:	N/A
Fee Amt Due:	Window not open	Surchg Amt Due:	Window not open	Total Amt Due:	Window not open
Fee Code:					
Surcharge Fee Code:					
Most recent events (up to 7):	08/19/2011 08/18/2007 12/19/2005 12/19/2005	Payment of Maintenance Fee, 8th Yr, Small Entity. Payment of Maintenance Fee, 4th Yr, Small Entity. Payor Number Assigned. Payer Number De-assigned. --- End of Maintenance History ---			
Address for fee purposes:	OGILVIE LAW FIRM 2552 South Wilshire Circle SALT LAKE CITY UT 84109				