

Network Computing®

SEPTEMBER 14, 2005 | WWW.NWC.COM For IT By IT

Sneak PREVIEWS

» Smart Steelhead Pounds Home Data

Riverbed's central and branch office appliance accelerates TCP applications and reduces file-transfer times BY SEAN DOHERTY

good

- Speeds all TCP applications
- Optimizes CIFS, MAPI, HTTP, SQL and FTP traffic
- Provides proxy file services
- Supports VLAN 802.1q, WCCP and Jumbo Frames

bad

- Lacks support for Cisco Interswitch Link (ISL)
- Autonegotiating 10/100/1000 NICs are hazardous to performance

To reduce costs, enterprises are removing local servers from branch offices and returning them to central locations. To help preserve the integrity of remote-office applications and file services, Riverbed Technology's Steelhead 2.0 delivers as if from the data center, and each time the same data passes through the Steelhead, it goes faster.

Riverbed sent two Steelhead Model 5010 appliances to our Syracuse University Real-World Labs®. I made homes for them in virtual data center and branch office environments and connected them using Shunra Software's Virtual Enterprise, which simulated a T1 link. In no time, the Steelheads were speeding file transfers (CIFS, FTP) and application services (MAPI, Microsoft SQL) between the data center and branch office, reducing transmissions that typically took minutes to seconds.

The Steelheads took up 3U in a standard rack and came with four Intel P4 Xeon processors (3.06 GHz) and 3,090 MB of RAM. Our units included six 250-GB SATA drives in a RAID 5 configuration with one hot spare. For network processing, the 5010s use 10/100/1000 autosensing NICs for LAN/WAN connec-

tivity and out-of-band management. The autosensing feature did not work with the switches in our test bed, though, and I had to configure the appliances' LAN/WAN ports for full-duplex 100-Mbps connections.

I labeled one Steelhead the "data center" and connected its LAN port to an SMC TigerSwitch 6752AL2, which was also connected to data center resources housed in a Dell 1650 computer and in a Rave Computer UltraSPARC III server. The Dell sported a Windows 2003 Server with Microsoft Virtual Server 2005 running a Domain controller, Exchange Server 2003, a Microsoft SQL 2000 Server and a file server. The Rave Computer ran Solaris



9 and was dedicated to serving HTTP and FTP files.

I labeled the other Steelhead "client center" and plugged it into an Extreme Networks Summit48 switch, attached to clients in a simulated branch office that made requests to servers in the data center. Traffic is optimized between the Steelheads in both directions, and whether clients were requesting or receiving data, the Steelheads compressed and sent the data over the WAN link without a hitch.

Smart, Speedy Transfers

In my tests, data traversing the WAN link for the first time had improved performance of 60 percent to 90 percent, depending on the file type. But the real story here is what happens after that first pass, when data passes through the Steelhead for a second time—that's where SDR (Scalable Data Referencing) comes in.

SDR lets one Steelhead appliance keep track of data passed through it and across the WAN to another Steelhead. So when a client sends or receives an original data request over the WAN and the data passes from one Steelhead to another, the receiving Steelhead saves the data on disk and returns reference pointers, down to the byte level, to the sending Steelhead.

The next time those same bytes are requested over the WAN, the sending Steelhead identifies the duplicate bytes and sends only the changed bytes.

The effect—a reduction in the amount of data traversing the WAN—lets clients retrieve data from a local Steelhead at lightning speed.

The results are amazing. In my tests, after data went through a Steelhead the second time, response time was so quick it was like requesting data from a local disk for small files—less than one second. For larger files, it was like requesting data from a remote disk on the LAN.

No Cache

Steelheads do not serve client requests from hard cache like the Peribit SM-250 (now owned by Juniper Networks and renamed the WXC application acceleration platform) does. If a back-end server

goes down, there's no original data for a Steelhead to send as duplicate data. If you need to preserve data in the event of a WAN outage, look to Riverbed's PFS (Proxy File Service).

I set up PFS for employees on the local Steelhead appliance and shut down a back-end server. These users could quickly access their stored files on the appliance, just as they could from a remote file share. In the background, the local Steelhead synchronized that data with a back-end file share using Riverbed's RCU (Remote Copy Utility) and SDR.

Next I saved a 593-KB file to the PFS directory on the local Steelhead. Then I reopened the file and added data, making it into a 601-KB file. After engaging manual synchronization, only 25 KB of data traversed the WAN link.

High-Speed TCP

Slow WAN links abound and some of these suffer from high latency (round-trip delay). If that's your situation, Riverbed has an answer in HSTCP (High Speed TCP). In this mode, two Steelheads can increase throughput and fill up a data pipe, maximizing WAN utilization.

I easily set up HSTCP over the Web interface. I checked a box for each of the Steelheads and increased the LAN buffer size to 1 MB. Then I increased the queue size on each of the WAN interfaces to 1,250 KB and set a rule to pass data through the Steelheads unmodified, without compression or SDR. Once this was set, I fired up a client with FTP to transfer files from the Rave Computer. I increased throughput from approximately 183.5 KBps without HSTCP to a maximum of 7,533.56 KBps with HSTCP, before the server hit its CPU limits. I tried the Dell 1650, a less CPU-bound machine, and was able to transfer 10,956.55 KBps.

HSTCP is for big pipes with high latencies and beefy servers. With HSTCP enabled, you won't get the benefits of compression or SDR to reduce bandwidth utilization. But if this fits your scenario, a couple of Steelheads will bring your data to the data center fast and possibly complete that remote, overnight backup on time before your users get to work.



RIVERBED TECHNOLOGY, INC.

501 2nd Street, Suite 410
San Francisco, CA 94107
+1 415 247 8800
www.riverbed.com

FTP

File size (bytes)	File type	No service Straight through (min:sec)	Service enabled	
			First pass (min:sec)	Second pass (min:sec)
1,132,544	XLS	6.00	1.97	0.25
2,489,344	DOC	14.00	11.49	0.40
5,562,921	MSI	30.00	15.84	0.63
11,338,499	MSI	61.00	47.58	1.17
25,067,348	MSI	131.00	66.01	2.54
281,769,984	TXT	1,494.71	122.02	31.24

Configuration: WAN, T1 100-ms delay, 29-KB queue

CIFS

File size (bytes)	File type	No service Straight through (min:sec)	Service enabled	
			First pass (min:sec)	Second pass (min:sec)
1,132,544	XLS	0.11	0.01	<0.01
2,489,344	DOC	0.23	0.03	<0.01
5,562,921	MSI	0.49	0.16	0.01
11,338,499	MSI	1.41	0.47	0.01
25,067,348	MSI	3.42	1.07	0.03
281,769,984	TXT	41.44	2.09	0.33

Configuration: WAN, T1 100-ms delay, 29-KB queue

Open Remote File

File size (bytes)	File type	No service One pass (min:sec)	First pass (min:sec)	Service enabled	
				Second pass file unmodified (min:sec)	Third pass file unmodified (min:sec)
1,136,640	DOC	0.16	0.12	0.03	0.03
1,201,152	PDF	0.19	0.06	0.02	PDF file not modified
2,489,344	DOC	0.22	0.12	0.03	0.04
6,572,000	DOC	0.30	0.17	0.03	0.05
6,448,121	PPT	1.41	0.21	0.15	0.13

Configuration: WAN, T1 100-ms delay, 29-KB queue

High-Speed TCP

File size (bytes)	File type	Service enabled		Service enabled	
		Time (sec)	Speed (KBps)	Time (sec)	Speed (KBps)
5,562,921	MSI	30	181	1.75	3,176.68
11,338,499	MSI	61	182	2.52	4,510.02
25,067,348	MSI	131	187	3.74	6,709.8
281,769,487	TXT	1,494.71	184	37.11	7,533.56

Configuration: WAN emulation, 100 Mbps, 100-ms delay, 1,250-KB queue

Exchange Client 2000/2003

File size (bytes)	File type	Exchange 2000 Client			Exchange 2003 Client		
		No service (min:sec)	Service Enabled First pass (min:sec)	Service Enabled Second pass (min:sec)	No service (min:sec)	Service Enabled First pass (min:sec)	Service Enabled Second pass (min:sec)
1,136,640	DOC	0.24	0.06	0.02	0.15	0.03	0.01
1,201,152	PDF	0.28	0.10	0.03	0.12	0.04	0.01
2,489,344	DOC	0.49	0.13	0.02	0.29	0.03	0.01
6,572,000	DOC	2.11	0.17	0.03	0.46	0.08	0.02
6,448,121	PPT	3.10	0.40	0.04	1.38	0.37	0.20

Default configuration: WAN, T1 100-ms delay, 29-KB queue Configuration: no-cache mode, WAN, T1 100-ms delay, 29-KB queue

SQL Queries

	No service (min:sec)	Service enabled	
		First pass (min:sec)	Second pass (min:sec)
SQL Query 2	0.18	0.04	0.01
SQL Query 3	0.43	0.04	0.01
SQL Query 4	1.26	0.07	0.02
SQL Query 5	4.19	0.10	0.06

Configuration: WAN T1 100-ms delay, 29-KB queue