LeftHand NETWORKS

Wide-area data services (WDS) Accelerating Remote Disaster Recovery

Reduce Replication Windows and transfer times leveraging your existing WAN

LeftHand Networks, Inc.
1688 Conestoga Street
Boulder, CO 80301
Phone:1.866.4.IPSANs
www.lefthandnetworks.com

Deploying Riverbed wide-area data services in a LeftHand iSCSI SAN – Remote Disaster Recovery Solution

OVERVIEW

The LeftHand SAN Network Storage Modules (NSMs) are the building blocks that are clustered together creating a highly available, virtualized storage pool using LeftHand's SAN/iQ engine. Advanced features supporting remote copy and snapshot capabilities provide robust disaster recovery for implementation at remote sites. A primary site and a secondary site tied together over a Wide Area TCP/IP Network (WAN) depict the most basic in remote disaster recovery (DR) configurations. Usually distributed across vast geographic locations, DR solutions introduce tradeoffs in recovery time/point objectives versus the cost of the WAN infrastructure required.

As method to achieving faster WAN-based replication while controlling bandwidth costs, Riverbed's Wide-Area Data Services (WDS) solutions enable application acceleration, WAN bandwidth optimizations and faster WAN backup in a single appliance. Deployed at both the remote and local site, the effect is as if bigger bandwidth with lower latency exists between those sites. As part of a joint LeftHand and Riverbed solution, our customers can typically see a **6.3x** increase in bandwidth performance and an **84%** of data reduction over the wire during SAN-to-SAN replication activities. WAN backup windows can be reduced by up to **8x**. The results for LeftHand Storage users are immediate: Reductions in remote backup windows which allow you to increase the frequency of the remote backups or increase the amount of data being remotely backed up; Accelerate restoration from a remote DR site after local disaster; Reduction in WAN costs and the ability to postpone WAN bandwidth upgrades. In addition, all other traffic between sites will also be accelerated with some protocols experiencing as much as a 100x increase in performance.

This document provides an overview of the configuration and performance characteristics in this tested solution so that your DR solutions can rely on a LeftHand SAN and Riverbed Steelhead WDS appliances.

CONFIGURATION

Storage

Site A and Site B both contain a three node LeftHand Networks' NSM 160 cluster. A null modem serial cable and terminal session (19200,8N1NoFlow) is used for basic configuration of the NSMs. Site A's NSMs are configured with unique names and Adaptive Load Balancing (ALB) network bonds at 10.20.80.191, 10.20.80.192 and 10.20.80.193. The LeftHand Networks Centralized Management Console (CMC) GUI is used to complete the configuration.

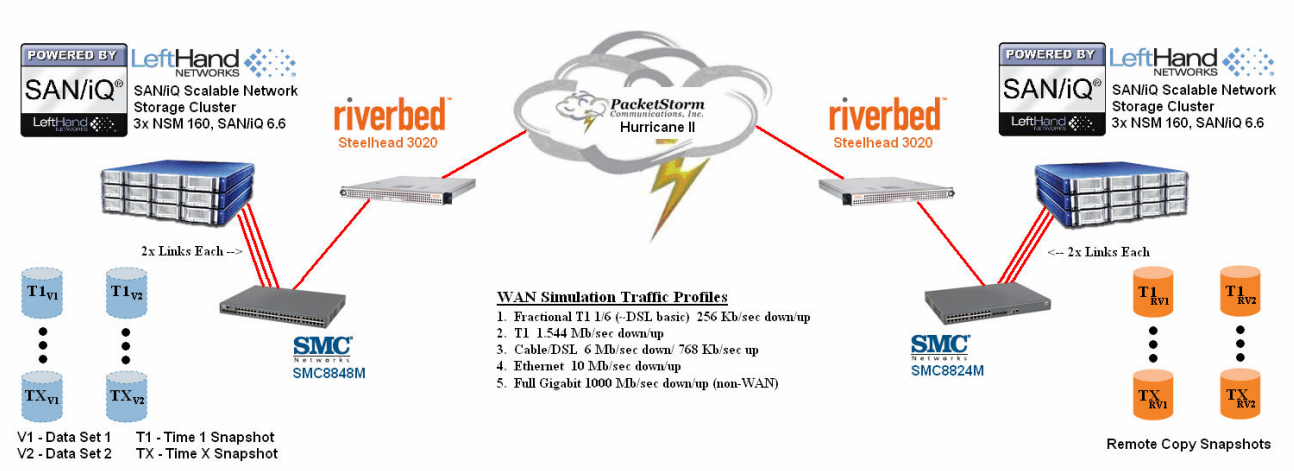

A Management Group "Site A" is created with these 3 NSMs. Each NSM is configured for RAID 0 with a Low Priority RAID Rebuild Rate. Through the Edit Management Group function, Local Bandwidth Priority is set to Maximum for these tests (Default is 4 MB/s). Simultaneous local application load requirements would need to be evaluated for optimal settings as this setting affects remote copy performance to the benefit of local application performance. "Cluster A" is created with these 3 Storage Modules. Cluster A will use a virtual IP address of 10.20.80.197. Two Volumes "VolA" and "VolB" are thin provisioned at 50GB each with 2-Way Replication and Auto Grow enabled. An Authentication Group "AG1" is created with the Initiator Node Name of a server for storing the original sourcing data. A Volume List "VL1" is also created containing volumes "VolA", "VolB" and the Authentication Group "AG1" thereby allowing access to the volumes by the initiator.

Site B's NSMs are configured with unique names and Adaptive Load Balancing (ALB) network bonds at 10.20.80.194, 10.20.80.195 and 10.20.80.196. A Management Group "Site B" is created with these 3 NSMs. Each NSM is configured for RAID 0 with a Low Priority RAID Rebuild Rate. Through the Edit Management Group function, Local Bandwidth Priority is set to Maximum for these tests (Default is 4 MB/s). Simultaneous local application load requirements would need to be evaluated for optimal settings as this setting affects remote copy performance to the benefit of local application performance. "Cluster B" is created with these 3 Storage Modules. Cluster B will use a virtual IP address of 10.20.80.198. Two Volumes "VolC" and "VolD" are thin provisioned at 50GB each with 2-Way Replication and Auto Grow enabled.

Wide-area Data Services (WDS) using Riverbed Steelhead appliances

Site A utilizes a Riverbed Steelhead 3020 managed at 10.20.80.201. Site B utilizes a Riverbed Steelhead 3020 managed at 10.20.80.202. A null modem serial cable and terminal session (9600,8N1HardwareFlow) is used for basic IP configuration. A web browser is used for additional configuration and monitoring. In order to segment each set of specific tests, data stores were cleaned on both units prior to each test run. The first data run after a data store clean is referenced by a cold run. A warm run will consist of any subsequent data throughput based upon Riverbed caching of the cold run and all subsequent warm runs.

Network Gigabit Switch

Site A utilizes an SMC 8848M Managed Gigabit Ethernet Switch at 10.20.80.19. Site B utilizes an SMC 8824M Managed Gigabit Ethernet Switch at 10.20.80.20. A null modem serial cable and terminal session (9600,8N1NoFlow) is used for basic IP configuration. A web browser is used for additional configuration. Default configurations were tested. In specific tests with Jumbo Frames, both switches must be configured with this enabled along with every device. To simplify the configuration across both simulated sites and to directly control the WAN simulations, no intermediate routing of the packets was tested. The testing subnet spans both sites. For maximum performance of switch infrastructure, gigabit switches should support Jumbo Frames and full mesh backplane throughput.

WAN Simulation

In between both sites, a PacketStorm Hurricane II WAN simulator is managed at 10.20.80.205. A local console allows basic configuration to enable browser mode access. Five WAN simulation profiles were configured: Gigabit, Ethernet 10Mb, Cable/DSL 6MB / 768Kb, T1 1.544Mb, and Fractional T1 256 Kb. For all WAN simulations except Gigabit, a 50ms latency delay was added in each direction.

Public Data Sets

Sequencing Data sets from ftp://ftp.ncbi.nih.gov/repository/dbEST were utilized for both compressed and non-compressed data set runs. Both 32GB and 5-8GB profile data sets consist of this compressed data. Non-compressed data is represented by the same data but de-compressed. Each dataset represents original data, 10% change rate, and 25% change rate.

Private Exchange Data Sets

The LeftHand Networks' live Corporate Exchange data set was utilized to profile a real world application. Several Information Store Volumes and Log Volumes along with five 24 hour snapshots for each volume were also profiled. This corporate Exchange dataset represents 200 real world users over a 5 day work week period of snapshots.

KEY TECHNOLOGIES

LeftHand's key technologies specifically demonstrated in this paper include SAN/iQ Network RAID, SAN/iQ Remote Copy, clustering performance and capacity scaling. Note that only three storage modules were used at each site. Specific RAID configurations, storage modules and number of modules in each cluster could be altered to greatly enhance performance. In performing remote copies of point in time snapshots, entire data set contents were transferred to provide consistent data sets through each WAN simulation. As LeftHand Remote Copies are efficient in network usage and only copy unique data changes, real-world implementations of remote copies would decrease the amount of data transferred over the WAN than what is simulated in these tests.

Riverbed's key technologies are contained within the Riverbed Optimization System (RiOS). The main technologies demonstrated here include Data Streamlining and Transport Streamlining. Data streamlining segments, indexes and stores the segmented data with associated references on disks on both sides of the WAN. Once the data is indexed, current data is compared to what is already on the Riverbed disk. If the data has been seen before, a reference key is sent over the WAN instead of the data thereby reducing the amount of traffic sent over the WAN. One small reference key can potentially represent megabytes of actual data. Transport Streamlining improves the transport protocols moving the data by adjusting the transport parameters to better match the conditions of the network. Typically Riverbed can reduce bandwidth utilization by 60% to 95% and reduce packet roundtrips by 65% to 98%, resulting in much faster replication.

CREATING THE DATA SET

Compressed and Uncompressed Data

The server for storing the original sourcing data (iSCSI Initiator) logs in to VolA. Copy 32 GB of compressed data from the public data sets to Vol1. Logout from VolA. Perform a local snapshot from the LeftHand Centralized Management Console (CMC). Remount VolA. Delete 10% of the data set and replace its contents with new data from the public data set. (10% Change Rate) Logout from VolA. Perform a local snapshot from the CMC. Remount VolA. Delete 25% of the data set and replace its contents with new data from the public data set. (25% Change Rate) Logout from VolA. Perform a local snapshot from the CMC. Repeat with VolB and 32GB of decompressed data.

Microsoft Exchange Corporate Exchange Data

The LeftHand Networks' Corporate Exchange Data set consists of 200 Users spread over 5 Volumes with unique Information Stores and 3 Volumes with unique Log files. This live data was snapshot over a 5 day cycle Monday thru Friday. No specific profiling of data or its change rates will be stated. The important benchmark is a consistency point in reference of results to a real world data set application.

METHODOLOGY

From the CMC perform a remote snapshot one at a time to SiteB for each snapshot on SiteA (VolA and VolB) and note time of completion and data throughput. Repeat with Jumbo Frames enabled for all devices. Repeat with all 5 WAN Profiles editing the remote bandwidth setting appropriately. Repeat with both Riverbed Steelhead 3020s enabled ensuring that after each set of snapshots, the data store is cleaned on both units. As the data store contains all of the previous acceleration lookup data previously tested, a COLD run is defined as clean store and acceleration is only based upon current data set. A WARM run is defined as the acceleration data store pre-populated with previous test run data. Due to the patented technology approaches by Riverbed, it is noted that larger data sets and extended runs improve the throughput performance.

The LeftHand Networks' Corporate Exchange Data set was remotely mirrored to SiteA as 8 unique volumes each with 5 snapshots. From the CMC perform a remote snapshot one at a time to SiteB for each snapshot on SiteA (all 8 volumes) and note time of completion and data throughput. Repeat with Jumbo Frames enabled for all devices. Repeat with all 5 WAN Profiles. Repeat with both Riverbed Steelhead 3020s enabled ensuring that after each set of snapshots, the data store is cleaned on both units.

MEASUREMENTS

LeftHand Remote Copy Snapshot Performance Over 5 WAN profiles without Riverbed

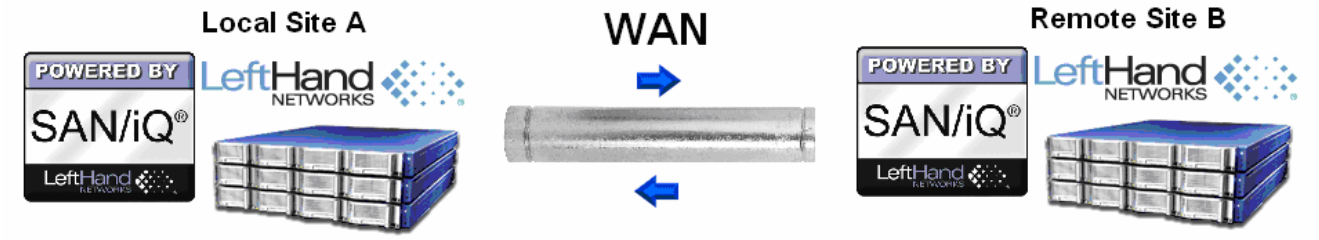

WAN Profiles	Theoretical DOWN MB/s	Theoretical UP MB/s	Actual LeftHand MB/S	Remote Copy #Hours/32GB	Remote Copy #Hours/5GB
Gigabit	125	125	21.58	0.4	0.1
10Mb/s	1.25	1.25	1.17	7.8	1.3
6Mb/s↓768Kb/s↑	0.75	0.096	0.62	14.7	2.5
1.544 Mb/s	0.193	0.193	0.16	*57	9
256Kb/s	0.032	0.032	0.03	*303	*47

* Test aborted after 1 hour to determine trend. Estimate Time for completion based upon trend.

The first set of tests compare theoretical maximums of the WAN bandwidth to an actual LeftHand 3 node cluster of NSM 160s at each site. Note that certain tests were aborted due to the length of time required to complete the expected test profile. Instead, a data run of 1 hour was used to generate a throughput trend in which an expected time for completion was extrapolated.

All WAN profiles less than Gigabit employ a 50 ms latency delay in each direction under best case no packet loss scenarios. It is worth noting that for all tests, remote copies with Jumbo Frames enabled slowed throughput by up to 18%. Therefore performance results only depict the default Ethernet 1500 Maximum Transmission Unit (MTU) byte size. Note that generally, local SAN throughput performance increases about 10% by enabling Jumbo Frames.

LeftHand Remote Copy Snapshot Performance Over 5 WAN profiles

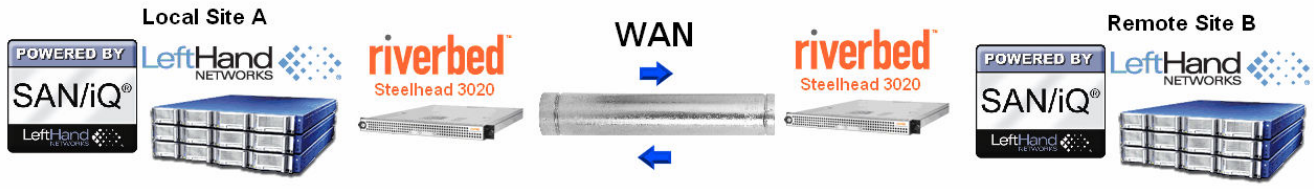

	Actual LeftHand Riverbed					
	COLD	WARM	Remote Copy COLD		Remote Copy WARM	
WAN Profiles	MB/S	MB/S	#Hours/32GB	#Hours/5GB	#Hours/32GB	#Hours/5GB
10Mb/s	2.2	6.63	4.1	0.8	1.4	0.2
6Mb/s↓768Kb/s↑	1.2	4.88	7.6	1.4	1.9	0.3
1.544 Mb/s	0.3	1.21	*31	4.9	7.4	1.2
256Kb/s	0.04	0.20	*228	*36	*46	7.2

* Test aborted after 1 hour to determine trend. Estimate Time for completion based upon trend.

The second set of tests examines the performance throughput gains with the Riverbed WDS product enabled at both sites. Typically, a 6.3x increase in bandwidth performance is achieved. For example, in comparison with the first set of test results, at 10Mb/s remote copy of 32GB of data completed in 7.8 hours versus 1.4 hours with the Riverbed solution – approximately 5.5 times faster. Note the performance increases from the COLD versus WARM runs. The 6.3x increase in bandwidth performance and 84% of data reduction over the wire was measured from the Riverbed appliance over the complete testing duration of all tests.

Compressed and non-compressed data testing results were similar in characteristics. The rate of data transferred being consistent in non-accelerated WAN profiles. Larger 32GB data sets typically benefited by a 10% performance increase in throughput over the 5GB data set. This can be rationalized by the inherent properties of the Riverbed bandwidth optimizations; a larger dataset will be better table lookup optimized by over 6 times the amount of data. Riverbed should increase in performance to a plateau characterized by the data sample. Data in itself can be performance enhanced by the pattern repetition found in the data.

The LeftHand Networks' Corporate Exchange data set volumes compared to the first data set. The noted exception is Log volumes performed a 7x increase in bandwidth performance over Information Store volumes.

CONCLUSION

The Riverbed Steelhead WDS solution complements any LeftHand Remote Disaster Recovery scenario. With a 6.3x improvement in bandwidth utilization, and replication windows that are 80% shorter, users can see a backup frequency increase or an increase in the amount of data replicated. An 84% data reduction over the wire reduces WAN use and possibly postpones planned bandwidth upgrades saving IT budget. Configuration of the LeftHand and Riverbed equipment can be performed in just minutes. Performance increased over time to a plateau in this configuration. As the LeftHand clustered network storage can be configured to scale in performance to the theoretical connection maximums, Riverbed also offers additional features that could trunk additional connections or units for additional performance. LeftHand Networks and Riverbed solutions are documented as interoperable.

LeftHand Networks, Inc.
1688 Conestoga
Boulder, CO 80301
Phone: 1.866.4.IPSANS (1.866.447.7267)

LeftHand Networks pioneered the Open iSCSI SAN market and is the only storage vendor to provide enterprise-class systems on industry standard platforms. LeftHand's patented SAN/iQ clustering storage software turns standard servers into scalable network storage ideal for disaster recovery and campus SANs. Each additional storage module added to the pool of storage scales both capacity and performance. SAN/iQ Network RAID capability ensures the highest continuous data availability in the event of network, disk, controller or an entire storage module failure. Advanced features including remote copy, snapshot, thin provisioning, auto grow, and automated load balancing simplify day-to-day management data.

sales@lefthandnetworks.com

http://www.lefthandnetworks.com

199 Fremont Street
San Francisco, CA 94105
Tel: 1-87-RIVERBED (1 877 483 7233)

Riverbed Technology is the performance leader in wide-area data services (WDS) solutions for companies worldwide. By enabling application performance over the wide area network (WAN) this is orders of magnitude faster than what users experience today, Riverbed is changing the way people work, and enabling a distributed workforce that can collaborate as if they were local.

info@riverbed.com

http://www.riverbed.com

SMC Networks North America
20 Mason
Irvine, CA 92618
Phone: 1.800.SMC.4YOU (1.800.762.4968)

A leading provider of networking solutions for the SMB/e and the home, SMC puts more than 30 years' experience to work in every product and solution the company manufactures. Leveraging strong industry relationships in combination with experienced and talented in-house design and engineering, SMC continues to bring to market a host of affordable, high-performance, easy to use products.

sales@smc.com

http://www.smc.com

PacketStorm Communications, Inc.
20 Meridian Road
Eatontown, NJ 07724
Phone: 1.732.544.2434

PacketStorm develops, manufactures, and supports high end testing solutions for the Internet Protocol (IP) communications market. PacketStorm is the market leader for advanced IP Network Emulators with dynamic and traffic conditioning emulation.

sales@packetstorm.com

http://www.packetstorm.com