riverbed

Test Plan

WAN Optimization Test Plan 8.0

Testing the Riverbed Steelhead Appliance

Riverbed Technical Marketing

© 2012 Riverbed Technology. All rights reserved.
Riverbed®, Cloud Steelhead®, Granite™, Interceptor®, RiOS®, Steelhead®, Think Fast®, Virtual Steelhead®, Whitewater®, Mazu®, Cascade®, Cascade Pilot™, Shark®, AirPcap®, SkipWare®, TurboCap®, WinPcap®, Wireshark®, and Stingray™ are trademarks or registered trademarks of Riverbed Technology, Inc. in the United States and other countries. Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed Technology. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein cannot be used without the prior written consent of Riverbed Technology or their respective owners.
Akamai® and the Akamai wave logo are registered trademarks of Akamai Technologies, Inc. SureRoute is a service mark of Akamai. Apple and Mac are registered trademarks of Apple, Incorporated in the United States and in other countries. Cisco is a registered trademark of Cisco Systems, Inc. and its affiliates in the United States and in other countries. EMC, Symmetrix, and SRDF are registered trademarks of EMC Corporation and its affiliates in the United States and in other countries. IBM, iSeries, and AS/400 are registered trademarks of IBM Corporation and its affiliates in the United States and in other countries. Linux is a trademark of Linus Torvalds in the United States and in other countries. Microsoft, Windows, Vista, Outlook, and Internet Explorer are trademarks or registered trademarks of Microsoft Corporation in the United States and in other countries. Oracle and JInitiator are trademarks or registered trademarks of Oracle Corporation in the United States and in other countries. UNIX is a registered trademark in the United States and in other countries, exclusively licensed through X/Open Company, Ltd. VMware, ESX, ESXi are trademarks or registered trademarks of VMware, Incorporated in the United States and in other countries.

Contents

© 2012 Riverbed Technology. All rights reserved.

© 2012 Riverbed Technology. All rights reserved.

© 2012 Riverbed Technology. All rights reserved.

WAN OPTIMIZATION TEST PLAN

Preface

This test plan is designed for the study of performance of (WAN) optimization and wide-area file service (WAFS) devices. Although there are sections of this test plan that refer to Riverbed Optimization System (RiOS™) devices and documentation, this test plan is designed so that you can test any WAN optimization or WAFS device and evaluate its functionality.

Audience

This document does not provide detailed “click-by-click” instructions; rather it assumes that you are familiar with common end user activities, such as file operations in Windows and MS Office applications. It is also assumed that you are familiar with basic setup and use of the Steelhead appliance, as well as any other WAN optimization or WAFS devices you intend to test. Finally, it is assumed that you are familiar with basic systems administration tasks such as installing operating systems and software and configuring network devices.

Questions to Ask Before Building Your Lab

The majority of this document describes tests you can set up and run in a lab environment, many of them focused on performance. However, the process of testing WAN Optimization technology gives an IT administrator much more than just the opportunity to test performance. It also gives an organization an opportunity to gain insight into product use, vendor support, and strategies for ultimately using the product within your production environment. While performing an evaluation, it is important to consider the following topics:

1. Will this product accelerate the broad range of applications used in my enterprise? Can the product do more than just accelerate CIFS traffic?
2. Will the product accelerate secure traffic as well?
3. How easy is this product to configure and deploy? Will I need support from the vendor for every device, or will my IT team be able to handle deployment? Will deployment scale as I add sites, or will it get harder as more locations are deployed?
4. What level of flexibility does the product have? Will I be able to easily change my deployment or configuration over time if I need to do so?
5. Can the product scale up and scale down to meet the needs of my largest and smallest sites?

In considering the feasibility of ultimately deploying a solution, the answers to these questions may provide even more guidance than the results of lab tests. It is advised that you seek the answers to questions like these before (or in parallel with) lab-focused testing.

Overview

This section provides an overview of the enterprise application acceleration solutions available in the industry today. It includes the following topics:

- The Enterprise Application Acceleration Challenge
- The Riverbed Approach
- Other Approaches

The Enterprise Application Acceleration Challenge

IT managers are stuck in the middle of a tug-of-war: on one hand, distributed employees need full, fast application access; but on the other, management wants to consolidate IT infrastructure in order to reduce expenses and meet regulatory compliance. How can an organization balance these needs? This challenge spans both commonly accessed end-user applications, such as MS Office, file sharing, and document management; and less visible administrative applications, such as backup and replication, software distribution, and database synchronization.

© 2012 Riverbed Technology. All rights reserved.

The Riverbed Approach

The Steelhead appliance takes a widely applicable, more scalable architectural approach to enterprise application acceleration. Based on the Riverbed Optimization System (RiOS), the Steelhead appliance is designed from the ground up to provide enterprises with one, easy to manage device that can accelerate all key applications that run over the wide area network.

Competitive systems typically started out as file caches, memory-based compression devices, or bandwidth-allocation systems. Although they may have subsequently adopted some elements of the Riverbed approach, their architectures have notable weaknesses in areas where Riverbed is strong.

Riverbed created RiOS to solve application acceleration challenges in a different way from other solutions. RiOS accelerates applications on three levels simultaneously:

1. Data streamlining: Data reduction for all TCP applications
2. Transport streamlining: TCP optimizations for all applications
3. Application streamlining: application-specific optimizations

In RiOS, each of these approaches happens independently, enabling all enterprise applications to benefit from data reduction and transport layer acceleration. Application layer acceleration is treated as one piece of the puzzle in the Riverbed architecture. The application-independent optimizations in RiOS enable email, file sharing, document management, ERP applications, CAD applications, network-based backup, software distribution, web-based applications, and even custom-built applications to see benefits.

The result of this approach enables massive acceleration for *all* applications that run over TCP. Users see up to 100 times faster application speed and up to 95% less bandwidth utilization at the same time. The system is designed to intelligently accelerate applications, while not creating the performance, consistency, or management problems that other approaches can create in today's networks.

Other Approaches and Their Problems

This section describes some of the alternatives to the RiOS approach to enterprise application acceleration. It includes the following topics:

- WAFS and Caching
- Bandwidth Allocation and QoS
- Per-peer data stores
- Tunnels
- Non-standard transports
- File Replication

WAFS and Caching

Caching is a narrowly-architected, partial approach to enterprise application acceleration that may introduce more problems than benefits into a network. Caching devices *seemingly* improve performance by short-circuiting the communication processes of application clients and servers. But, for the following reasons, this approach is limited in its efficacy, is error-prone, and is hard to manage:

- **Inability to Leverage Data Commonality Across Different Files**. File caching in itself does not leverage data commonalities that may exist among multiple files. Accessing files that have different filenames, but are identical in other respects, will result in the same raw data being sent across the WAN multiple times. The Riverbed approach, on the other hand, will identify repetitive byte patterns that may exist across different files, or different versions of files based on the same data. It will even recognize and leverage data commonalities from files transported by different applications and protocols (e.g., e-mail attachments, CIFS file shares, FTP, etc.). Some vendors combine a Riverbed-like approach with a file cache, but such a design is not the "best of both worlds" – instead, that combination forces the device to represent each file twice, once in a "cache" style and once in a "Riverbed-like" style. The Riverbed approach uses memory and disk more efficiently to have a larger useful optimization vocabulary and thus optimize more traffic.
- **Data Integrity and Coherency Issues**. The file caching approach inherently involves fragmenting storage used for files into multiple locations, and then relying on the file cache's automated synchronization mechanisms to ensure copies of the files

© 2012 Riverbed Technology. All rights reserved.

remain consistent with each other. However, such an approach is vulnerable to a “split-brain” phenomenon, and there are multiple scenarios where these synchronization mechanisms can break down, leading to inconsistencies in how a given file is stored at each physical location. Network outages, administrator intervention, file cache failure, and other unforeseen events can easily result in inconsistent or corrupted file collections. In contrast, the Riverbed Steelhead appliance operates transparently and does not interfere with client-server interactions. This includes data access and file locking, which is always end-to-end between client and server when using the Steelhead appliance.

Bandwidth Allocation and QoS

Various kinds of bandwidth allocation, bandwidth limiting, prioritization, and QoS tend to address issues of bandwidth more than performance. As a result, these technologies can sometimes improve end user performance, but the overall effect is typically much less dramatic than with other approaches Generally, any performance improvement comes from picking winners and losers when network resources are insufficient. These kinds of systems cannot improve performance unless there is congestion. In contrast, Steelhead appliances often produce dramatic improvements in performance even when there is no congestion, and without requiring tradeoffs of one traffic type against another.

QoS and bandwidth management techniques can be useful in many networks, but they are properly used for managing situations in which bandwidth is still scarce even after traffic has been accelerated, or where some traffic types must be protected against possible congestion. Actually accelerating traffic, as Riverbed does, requires multiple streamlining techniques -- and should not be confused with QoS. In environments that require QoS functionality, there are QoS enforcement features available in RiOS. For additional information on QoS you may wish to read the Riverbed white paper *Riverbed and QoS FAQ*.

Per-peer data stores

The data streamlining part of RiOS maintains a data store – a kind of dictionary of previously-seen byte patterns and their encoding on the WAN. An unusual aspect of RiOS is that an appliance has one single unified data store that is used for communicating with any number of peer appliances across the WAN. In competitive systems that imitate this Riverbed feature, the typical implementation is a per-peer data store: one data store devoted to each peer device. Why does this matter? Consider a simple example with a data center and two branches, and consider the common scenario in which the same file being fetched to both branches. In the Riverbed approach, there is only one representation of that file in the data-center appliance. In competitive per-peer approaches, there are two copies of that file in the data-center appliance – one representation per branch. Naturally, this difference becomes even more striking (and important) as the number of branches grows from two to tens, hundreds, or thousands.

Why don’t other products also use a unified data store? It turns out to be harder to build. It’s easier to build a small-scale system using per-peer data stores. Although Riverbed designed RiOS for its current market, most competitive systems have been adapted from other kinds of systems. The temptation must have been strong to just build something simple rather than doing the more sophisticated engineering required for high performance at high scale.

Tunnels

Steelhead appliances communicate via ordinary TCP connections that are usually set up in an auto-discovery process, occurring as a side-effect of a conventional TCP “three-way handshake” between client and server. Some competing approaches instead depend on tunnels connecting appliances, functioning as pipes for optimized traffic to reach the appropriate peer. When configuring tunnels, the amount of work grows significantly faster than the number of tunneling devices in the network. Each and every routed subnet must be identified and manually configured in each device, and large networks may have hundreds or thousands of IP subnets. Next, each device must be manually configured with the IP addresses of all other devices with which it will communicate. Tunnel bandwidth limits must be configured for each and every tunnel, and verification is necessary to ensure the tunnel bandwidth does not exceed the physical bandwidth available on the WAN link. For a large network with many sites, the amount of work and potential for error will be unmanageable, particularly for full-mesh MPLS-type networks, where n*(n-1) tunnel configurations must be manually performed. Although some of this process can be automated with a “tunnel manager” or “automatic tunnel configuration” (which is different from and should not be confused with Riverbed’s auto-discovery mechanism), there still remains a significant amount of work required to configure each and every subnet and tunnel bandwidth limit, especially in a large network with many subnets and tunnels.

Tunnels also frequently give rise to problems with network integration and/or visibility of optimized traffic. Multiple distinct flows are multiplexed together into a single flow, which can change the way that the network treats the traffic. Because tunnel-based

© 2012 Riverbed Technology. All rights reserved.

systems make it hard to distinguish the original flows, QoS mechanisms and reporting must be handled by the optimizing device – it "seizes control" of QoS and reporting. In contrast, the Riverbed approach offers flexibility for integration with a wide variety of reporting and QoS schemes.

Non-standard transports

Steelhead appliances by default send optimized traffic via conventional TCP. Alternative transports like HS-TCP and MX-TCP are available for the particular environments where they are needed, but they are not mandatory and they are not enabled by default. In contrast, some competing approaches are built around unconventional transports that are "unfriendly" to un-optimized traffic, stealing bandwidth improperly so that they appear to have better performance. With many competitive approaches, such an unconventional transport is mandatory – there is no way to avoid its use other than not optimizing traffic. In some other approaches, a safe and well-behaved transport is available – but must be explicitly configured, since it is not the default.

File Replication

Another approach to providing faster access to data is to simply make a copy (or "replica") of it available in every location. This may be facilitated through multiple approaches such as software on servers or features built into NAS systems. When this approach is applied to file data, the ultimate solution suffers from some of the same drawbacks as WAFS/caching, including inability to leverage data commonality among files and data integrity/coherency issues. File replication carries additional drawbacks as well:

- **Expense of Equipment**. Maintaining individual servers or NAS systems in multiple locations just for the purpose of maintaining local replicas is typically much more expensive than deploying appliance systems. If a server approach is used, there is additional administrative cost to ensuring patches and security configurations are current on each system.
- **Replication is Inefficient**. Replication typically requires allocating enough storage in every location to hold a full copy of data, even if much of that data is never accessed from some locations. This also leads to wasted bandwidth as distributed updates are sent to locations where data may never be accessed.

This use of replication for branches is different from data-center replication. Replication technologies are often utilized instead to create off-site replicas of data for purposes of disaster recovery. This is a separate use case in which Steelhead appliances typically provide significant optimization of the replication traffic.

Test Terminology

Following are some terms used throughout this document:

Term	Definition and Usage
DUT	Device Under Test. This abbreviation refers to any WAN OPTIMIZATION, WAFS and/or WAN optimization system you are testing.
Data Store	A disk or memory system in a DUT that learns about data that has passed through the network in order to eliminate the need to pass the full data sequence through the network in the future.
Un-optimized	A baseline operation, without the DUT, to provide a measured basis for comparison.
Optimized	Operations run with the DUT.
Cold	The first time the data passes through the DUT. A cold test can typically be ensured by "cleaning" the data store of the DUT in advance.
Warm	The second time the data passes through the DUT. A data store is typically leveraged in this case.

In most parts of this document, these terms apply to any WAN OPTIMIZATION, or WAFS device you are testing. The text explicitly states when something applies only to RiOS or the Steelhead appliance.

Setting Up Your Test Environment

This section describes how to set up a controlled test environment in a lab to thoroughly test the configuration and benefits of the DUT. The controlled environment ensures that tests are reliable and repeatable. This section includes the following topics:

- Simulating WAN Constraints

© 2012 Riverbed Technology. All rights reserved.

- Setting up Server Systems
- Setting up Client Systems
- Configuring a Network Topology
- Setting up the DUT

Simulating WAN Constraints

All WAN OPTIMIZATION and WAFS systems seek to improve performance in WAN environments. Therefore, it's crucial to simulate the effects of a WAN in your test environment, including constrained bandwidth and network latency.

The following table lists some WAN simulators and their respective strengths and weaknesses:

WAN Simulator	Description
Network Nightmare / GigEnn	An inexpensive and adequate solution is Network Nightmare (http://www.networknightmare.com/riverbed), which has a simple menu-driven interface and can reliably simulate WAN networks of up to 25 mbit/sec. The newer model, gigEnn-T3, can reliably simulate WAN networks of up to 45Mbps. And the gigEnn-OC12 can simulate up to 622 Mbps. Network Nightmare is based on open-source Dummynet technology that is available as part of FreeBSD.
Shunra	More deluxe WAN simulators are available from Shunra (http://www.shunra.com). While they offer many advanced features, they may be considered too expensive for a brief evaluation.
Apposite	Another off-the-shelf option is available from Apposite Technologies (http://www.apposite-tech.com). Their Linktropy Mini2 is an option of roughly similar cost, form-factor, and capability to the Network Nightmare.
Dummynet	For a technical "do-it-yourself" option, FreeBSD/Dummynet technology can be installed on a PC to create a bare-bones WAN simulator (http://info.iet.unipi.it/%7Eluigi/ip_dummynet). Use of Dummynet requires fairly deep UNIX® experience including kernel configuration and building firewall rules.
NISTnet	A similar "do-it-yourself" option for Linux is available called NISTnet (http://snad.ncsl.nist.gov/itg/nistnet). However, NISTnet is not recommended because it has shown itself to be fairly buggy.

It is recommended that you perform many of your tests with fairly constrained bandwidth, such as 256 Kbps. Testing with such links simulates small office environments and the share of the link that actual users see on larger links, such as T1. When performing tests that involve multiple simultaneous users, test faster links as well.

Set the network latency to one that is typical of a remote office (20 ms for regional, 100 ms for national, and 200 ms for international).

Attempting to add simulated packet loss to your test environment is not generally recommended. If you do have sustained random packet loss in your production environment (i.e. "dirty circuits") and absolutely wish to include this in your tests, ensure to configure any packet loss mitigation features on the DUT to properly mitigate the effects of the packet loss. *In the case of the Steelhead appliance, this means enabling use of MX-TCP.*

Different products excel in different WAN environments. Products that focus mainly on compression may not do as well in high-latency environments. Some products may lack a data store and therefore perform poorly in low-bandwidth environments. In the quest for the one product that helps under the most conditions, we advise testing under several different WAN environments. Consider at a minimum the following 2x2 matrix.

Low bandwidth, low latency	**High bandwidth, low latency**
Low bandwidth, high latency	**High bandwidth, high latency**

In terms of specific numbers, you may want to make selections based on what you think your stakeholders want to see. Riverbed recommends basing the 2x2 matrix on the four combinations of 20 & 100 ms of round-trip latency and 256k & 1536k of bandwidth.

© 2012 Riverbed Technology. All rights reserved.

Setting up Server Systems

Your test environment likely should likely include a "headquarters" or "data center" location, including one or more server systems. Configure these servers in a way similar to your production environment. A typical environment might include some of the following servers:

- Windows environment
  - Microsoft Windows Sever operating system (2000, 2003, 2008, or 2008 R2)
  - Windows file shares
  - Microsoft Exchange Server (2000, 2003, 2007, or 2010). Note: you will need a 64-bit system for Microsoft Exchange Server 2007 and 2010
  - Web server software (such as IIS, which ships with Windows Server)
  - FTP server software (such as FileZilla)
  - Lotus Domino Server
  - Microsoft SQL Server
  - Backup/Replication server software
  - Other collaborative servers, such as Microsoft SharePoint
- UNIX® or Linux Server operating system
  - NFS server/exports
  - Samba fileserver with configured file shares
  - Web server software (such as Apache)
  - FTP server software (such as FileZilla)

Applications such as Exchange, IIS, or Apache should be configured and tested before you install the DUT.

The following table describes how to provision servers for the test environment:

Application	Procedure
File Sharing (CIFS)	Create a share directly on the CIFS server with permissions suitable for the client machine to access it. Ensure that the DUT allows existing operations to continue without administrative involvement. For example: • Ensure that clients access Windows File shares through the same UNC pathname and/or mapped drives as before the DUT were introduced. • Ensure that no client plug-ins, Active X controls, applets, or other special client-side software is required. • Confirm that the DUT do not need to be configured with knowledge of server hostname or share names. **Note on SMB signing**: The new CIFS Optimization for SMB Signed Traffic in RiOS version 5.5 enables Steelhead appliances to deliver the same high levels of CIFS performance even when the CIFS traffic is signed by the server. With a Windows 2003 server, RiOS 5.5 will provide CIFS Application streamlining even with SMB Signing enabled on the Windows 2003 server. Please refer to the 'SMB Signed traffic' test case in this document. For further information on SMB signing see the Riverbed Deployment Guide. Other WAN optimization products can only achieve their best performance when SMB signing is disabled. SMB signing is turned on by default on domain controllers and turned off on member servers. We recommend you configure SMB signing as you would expect it to be configured in your production environment. **Note on SMB v2:** SMB v2 is the native file transfer protocol for file sharing between Windows 7 and Windows Server 2008 R2. Starting with RiOS 6.5, Steelheads natively accelerates this protocol.
File Sharing (NFS)	Export a file system on the NFS server with permissions suitable for the client machine to mount it. Configure the exported file system with permissions that allow read and write permission for a non-root user as whom you'll perform tests. Ensure that the DUT allow existing operation to continue without administrative involvement. For example: • Ensure that clients mount NFS file systems using the same pathnames as before the DUT were introduced.

© 2012 Riverbed Technology. All rights reserved.

Application	Procedure
	• Ensure that no client plug-ins or other special client-side software is required. • Confirm that the DUT do not need to be configured with knowledge of server hostname or export names
E-Mail (Exchange/Domino)	Create users **test** and **test2** on the server, each with Exchange/Domino mailboxes. Ensure that the DUT allows existing operation to continue without administrative involvement. For example: • Ensure that clients access Exchange/Domino servers in the same way they did before the DUT were introduced. • Ensure that no client plug-ins, Active X controls, applets, or other special client-side software is required. • Confirm that the DUT do not need to be configured with knowledge of e-mail server hostname.
Web (HTTP/HTTPS)	Prepare web applications that you typically use in your work environment, such as a Perforce server or a web-based CRM system. If this is not applicable, stage large files on a new web server (for example, common applications, such as Windows Update or antivirus updates). Ensure that the DUT allows existing operation to continue without administrative involvement. For example: • Ensure that client browsers and applications do not need to “point” at the DUT. • Ensure that there is no risk of the DUT serving “stale” web objects • Ensure that the DUT supports and accelerates operations for all HTTP methods used in your environment (e.g. WebDAV)
File Transfer (FTP)	Configure your FTP server to default to binary downloads.

Setting up Client Systems

Your test environment should likely include a “branch” or “remote office” location, including one or more client (end user) systems. Configure these clients in a way similar to your production environment. A typical environment might include some of the following client systems:

- Windows environment
  - Microsoft Windows (XP, 2000, Vista, or 7)
  - Microsoft Office Professional (XP, 2000, 2003, 2007, or 2010)
  - Lotus Notes Client
  - Backup/Replication client agent software
  - Other software typically used in your environment (AutoCAD, Microsoft Project, etc.)
- UNIX® or Linux Server operating system
  - NFS client
- Mac environment
  - Mac OS 10.5 (Leopard) or later
  - Microsoft Office:Mac 2008
  - Other software typically used in your environment

Note: To simplify testing of Microsoft Office operations, turn off fast saves and background saves. For example, in Microsoft Office, choose *Tools*→*Options*→*Save*; under Save Options, uncheck Allow Fast Saves and Allow Background Saves.

Configuring a Network Topology

Many DUT have deployment options that allow them to integrate into networks in various ways. If one or more of these deployment options are of interest to you, you should focus on integration with your network topology as part of your testing. However, if a primary goal of testing is to focus on performance-enhancing benefits of the DUT, you might prefer a simple test environment.

© 2012 Riverbed Technology. All rights reserved.

The topology of a simple test network for Steelhead appliances is shown in the following illustration:

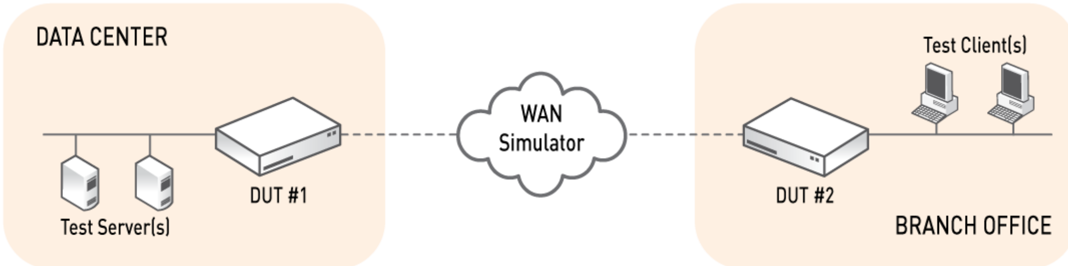

For network topology information on alternative DUT, contact the appropriate vendor.

Note: Ensure that network ports "match" in their speed and duplex settings. Mismatched speed and duplex can often cause inconsistent or poor results. We recommend you "hard set" ports to a common setting (e.g. 100 mbit/sec, full-duplex) before you begin your tests.

Setting up the DUT

To set up the RiOS system, follow the instructions in the Getting Started Guide or Steelhead Appliance Installation and Configuration Guide. You can download the RiOS system documentation at *http://www.riverbed.com/support*. Contact your Sales Engineer for a login and password. Riverbed recommends you obtain the following documents:

- Steelhead Appliance Installation and Configuration Guide
- Riverbed Deployment Guide
- Steelhead Management Console User's Guide
- Riverbed Command-Line Interface Reference Manual

In addition, Riverbed recommends you upgrade the software on the Steelhead appliance to the latest version.
For information on setting up alternative DUT, contact the appropriate vendor.

Measuring Test Performance

This section describes the recommended method for measuring test performance. It includes the following topics:

- Choosing Test Files
- Test Methodology
- Timing Operations
- Obtaining Performance Metrics

Choosing Test Files

The best test files are ones that your users actually work with on a daily basis. These files should be as typical as possible—some large, some small, different types, and so forth.

Ideally, for some files you will have active copies, such as a file copied twice—once on one day and once on another, having been edited in between. If you don't have "real world" examples of these slightly-edited files, create copies of some of your test files and perform your own edits on them.

Test Methodology

In each case, you want to perform:

1. **Un-optimized operations** to record a "baseline" performance without the DUT. In the interest of efficiency, Riverbed recommends you perform the baseline test first for all the different scenarios. You can perform a baseline test by using a network without the DUT or use a network that does not use the same network path as the DUT.
2. **Optimized operations** using the DUT:

© 2012 Riverbed Technology. All rights reserved.

a. **Cold Test.** Because most DUT include a data store and some type of compression, the first optimized test will often show some performance improvement and a reduction in bandwidth use.
b. **Warm Test.** The second time an operation is performed will often show more dramatic improvements.
c. **Partial Warm Test.** Performing the operation yet again, this time accessing the slightly-edited copy of the test file.

Note: with some applications, e.g. those that access back-end databases, the network may not be the bottleneck. It may be useful to record a "LAN baseline" to determine actual theoretical application performance with no network constraints. For details, please see the "HTTP Features in RiOS 5.0" whitepaper, available from your Riverbed representative.

Timing Operations

This section describes how to time test operations. It includes the following topics:

- Manual Timing
- Riverbench

Manual Timing

While a wall clock or the system clock has sufficient accuracy for many of the un-optimized tests, they are not sufficiently accurate for the optimized tests. Riverbed recommends that you use a stop watch for the tests.

Riverbench

Riverbench is a fully featured graphical benchmarking tool. It can automatically time many standard performance tests such as CIFS, HTTP, and FTP. In addition, if a Steelhead is present and optimizing traffic it will calculate the WAN bytes used. Having both presented in a graphical format makes the both the time and WAN bandwidth improvements easy to visualize.

Obtaining Performance Metrics

Start by recording baseline numbers, to understand the user experience across a high-latency, low-bandwidth WAN. Simple operations like a drag-and-drop of a file or downloading an e-mail attachment can take many seconds or even minutes.

The easiest way to measure performance improvement is to simply time it with a stopwatch: perform an operation with and without the benefit of DUT and record the times. While doing stopwatch testing, it is also important to measure the WAN utilization. For many DUTs, part of the performance gain comes as a result of removing repetitive data from the WAN. Therefore, it is important to log the bandwidth utilization for each test you run. Depending on which WAN simulator you choose, there are different ways to log this. For instance, the Apposite Linktropy Mini2 WAN simulator includes a "Monitor" web page that provides a count of the total number of bytes that have traversed an interface during the last 10 minutes. Sample output (showing approximately 11 MB of network traffic in the A->B direction) is shown in the following illustration:

© 2012 Riverbed Technology. All rights reserved.

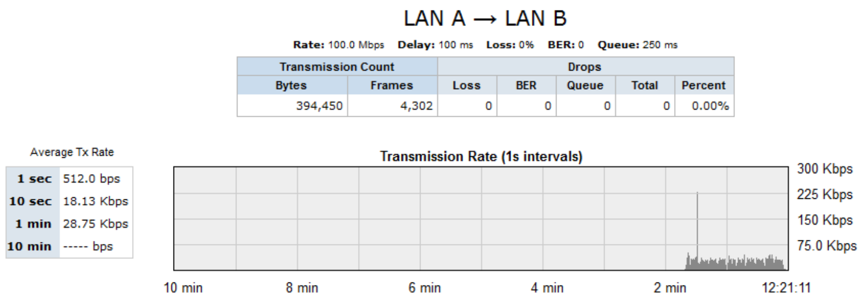

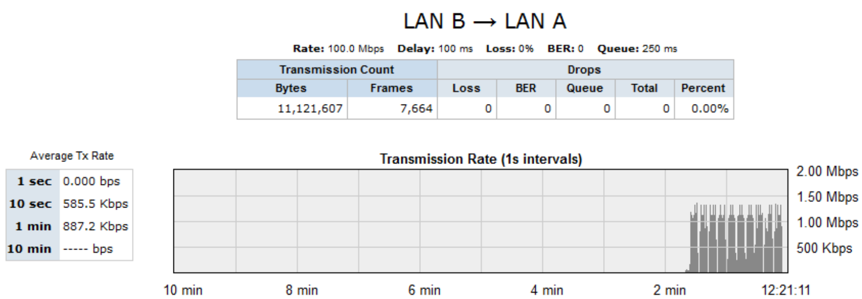

It is important to measure performance improvement and bandwidth savings separately because an improvement in one does not always imply an equal improvement in the other. For instance, many WAN Optimization devices minimize bandwidth utilization but do little or nothing address the chattiness of TCP or application-layer protocols and therefore do not improve performance much (if at all). Similarly, some products address protocol chattiness, but lack advanced data reduction methods, so they may improve performance somewhat but flood WAN pipes in the process. An ideal solution should manage to excel in both areas, so both should be measured.

Performance Statistics

The Steelhead appliance includes a statistical facility that collects and calculates performance statistics in a variety of areas. The key statistics available include:

- Bandwidth optimization
- Data reduction
- Throughput
- Traffic summary

Statistics can be viewed in the management console or in tabular format through the Steelhead appliance command-line interface (CLI). For detailed information about viewing statistics, see the *Riverbed Steelhead Management Console User's Guide* and the *Riverbed Steelhead Appliance Command-Line Interface Reference Manual*.

Testing Performance: Basic Tests

The benefits of a DUT can be showcased by performing fairly simple operations from a client system. Not only are these operations easy to test, they represent commonly-performed, heavyweight operations performed in typical production environments.

Each basic test follows the same format:

- START your timer and WAN bandwidth monitor
- Perform the operation from your client system
- STOP your timer and WAN bandwidth monitor
- Record your results

Test data should be put in place before you begin testing. Each test should be performed in un-optimized, cold, warm, and partial warm modes with each DUT you intend to evaluate. Some operations to perform in basic tests:

© 2012 Riverbed Technology. All rights reserved.

- Drag & drop from a remote mapped drive to the local C: drive. Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)
  - A folder containing many small files

- Copy (such as with “cp” in UNIX®/Linux) from a mounted NFS file system to local disk. Note that many NFS clients include a local cache, so you may need to “touch” files on the server between runs to invalidate the cache. Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)
  - A directory containing many small files

- Drag & drop from a remote mapped drive to the local Macintosh hard drive. Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)
  - A folder containing many small files

- Opening (or saving to local disk) an attachment from an email message (non-cached mode Outlook/Exchange and/or Notes/Domino). Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)

- Opening (or saving to local disk) an attachment from a public folder (non-cached mode Outlook/Exchange) and/or Notes database (Notes/Domino). Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)

- Use an FTP client (such as Filezilla) to download a file. Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)

- Use an HTTP client (such as wget) to download a file. Run tests for each of the following items:
  - A medium-sized file (2 MB)
  - A large-sized file (7 MB)

- Performing a full backup/restore of a large data repository.

Tip #1: Run through each test once or twice before timing it to familiarize yourself with the flow.

Tip #2: Run through all the tests and record baseline measurements before connecting the DUT and performing tests with optimization enabled.

Record your results in the following table:

© 2012 Riverbed Technology. All rights reserved.

Operation	Data	Baseline		Cold				Warm				Partial Warm			
				Steelhead Appliance		Alternative DUT		Steelhead Appliance		Alternative DUT		Steelhead Appliance		Alternative DUT	
		Bandwidth	Seconds	Bandwidth	Seconds	Bandwidth	Seconds	Bandwidth	Seconds	Bandwidth	Seconds	Bandwidth	Seconds	Bandwidth	Seconds
File-sharing (Windows CIFS)	2 MB														
	7 MB														
	Folder														
File-sharing (Mac CIFS)	2 MB														
	7 MB														
	Folder														
File-sharing (NFS)	2 MB														
	7 MB														
	Folder														
E-mail messages	2 MB														
	7 MB														
E-mail folders	2 MB														
	7 MB														
FTP	2 MB														
	7 MB														
HTTP	2 MB														
	7 MB														
Backup	Repository														
Restore	Repository														

Testing Performance: Advanced Tests

The differences between different DUT will often be shown by testing more advanced operations. These tests may take some effort to setup, measure, and understand, but many of them are more representative of what actual users really do. Some of the advanced operations covered by these tests include:

- **Advanced Protocol Operations.** Rather than simply helping to copy large items, a DUT should be able to accelerate operations such as browsing folders.
- **Cross-Application Functionality.** Data patterns "learned" via one protocol should be leveraged when they appear via another protocol.
- **Bi-directional Acceleration.** Data patterns "learned" when data traverses the network in one direction should be immediately usable to accelerate data patterns traversing the network in the opposite direction.
- **Working Within Applications.** Accessing files from within applications (such as through "File" drop-down menus) often generates network traffic that is different from that at the operating system level (such as through drag & drop). A DUT should be able to excel in both environments.
- **Incremental Operations.** A DUT should provide "warm" performance levels when slightly-modified versions of documents traverse the network.
- **Long-Term Data Reuse.** A properly-sized DUT should have the capacity to re-use data patterns learned weeks in the past. Its knowledge should also survive periodic power outages or scheduled shutdowns for maintenance.

© 2012 Riverbed Technology. All rights reserved.

Note: In order to achieve representative test results, some tests will require you to clear the DUT's Data Store and some will not. Please read the individual test for more information.

Testing File Operations

SMB Signed Traffic

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for SMB signed traffic

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">In order to optimize a CIFS signed connection the server-side Steelhead needs to authenticate the user making the connection. To authenticate the user the server-side Steelhead must be joined to a domain which is either the same as the domain of the user or has a trust relationship with the domain of the user. The domain to be joined should be a Windows 2003 Domain running in Native Mode.
A CIFS signed connection is optimized only if SMB signing is enabled on the server-side Steelhead.

1. SMB Signing should be set to "required" on both the client and server machines.
2. Create a shared folder "shareme" on the server
3. Add some content to the shared folder
4. Mount the shared folder on the client
5. Navigate to the shared folder on the client
6. Start the timer.
7. Open a document in the shared folder on the client
8. Make a change and save it.
9. Stop the timer.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">If SMB Signing is disabled, signed connections will be passed through.
If SMB Signing is enabled, signed connections will be optimized.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>File Modify</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>File Modify</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>File Modify</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

SMB v2 Traffic

SMB v2 is the default protocol used when sharing files between Windows 7 and Windows Server 2008 R2 systems. This section requires a share to be set up between a Windows 7 client and a Windows Server 2008 R2.

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for both signed and unsigned SMB v2 traffic

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Enable SMB v2 optimization. This can be done on the Steelhead Management console by navigating to Configure -> Optimization -> SMB2 and selecting Enable SMB2 Latency Optimization.
2. Create a shared folder "shareme" on the server
3. Add some content to the shared folder
4. Mount the shared folder on the client
5. Navigate to the shared folder on the client
6. Start the timer.
7. Open a document in the shared folder on the client
8. Make a change and save it.
9. Stop the timer.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">The connection should be shown as a fully optimized connection</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>File Modify</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>File Modify</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>File Modify</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Testing MAPI (2003/2007 Cached Mode, Encrypted MAPI, and Outlook Anywhere)

The basic tests covered cases of MAPI operating in non-cached mode, which would be typical of Outlook 2000, Outlook XP, or Outlook 2003/2007/2010 when configured not to use cached mode. Because the cached mode configuration of Outlook 2003/2007/2010 creates different use cases, this section provides tests for those scenarios.

This section provides the following MAPI tests:

1. Downloading Email with Attachments

© 2012 Riverbed Technology. All rights reserved.

"Boomerang" Email
2. Encrypted MAPI
3. Outlook Anywhere

1. Downloading Email with Attachments

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for downloading email that contains attachments.

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Log in to the remote Exchange server from a local LAN client and send an email with attachments to the users test and test2. Attach files that are typically used in your workplace, such as Excel spreadsheets, PowerPoint slides, AutoCAD drawings, and so forth.
2. Log into the client machine as user test. Open Microsoft Outlook on the local client and wait until you see the message in the lower-right of the window that says "Updating Inbox".
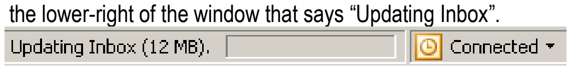

3. START your timer and WAN bandwidth monitor.
4. Observe the thermometer filling in the lower-right of the window. Wait for the new message to show up in the Inbox.
5. STOP your timer and close Outlook.
6. Record your results.
7. Log into the client machine as user test2. Open Outlook and wait until you see the "Updating Inbox" message.
8. START your timer.
9. Observe the thermometer and wait for the new message to appear in the Inbox.
10. STOP your timer and WAN bandwidth monitor.
11. Record your results.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">When the user test downloads the message, the Steelhead appliance has never seen this data before and therefore you will experience cold performance.

When the user test2 downloads the attachment, the Steelhead appliance already has the data and therefore can deliver LAN-like performance. If there were several more users in the branch (i.e. a bulk email scenario) each additional user would have experienced Warm performance.

Pay special attention to the amount of bandwidth used by the user test2 when receiving the attachment.</td></tr>
<tr><td rowspan="7">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td rowspan="2">Baseline</td><td>Email download (#6)</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td>Email download (#11)</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>Email download (#6)</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>Email download (#11)</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

2. "Boomerang" Email

Use Case

This procedure tests acceleration of the round-trip of an email between two users at the same location (both users access a centralized Exchange server across the WAN).

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? No

Procedure

1. Log into the client machine as user **test**. Open Microsoft Outlook and compose an email message with attachments to user **test2**.
2. START your timer and WAN bandwidth monitor.
3. Click **Send**. Observe the queued outgoing message "**Outbox [1]**" and the message "**Sending message 1 of 1**"

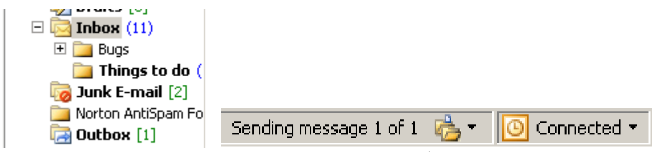

4. When the message disappears from the Outbox[1], STOP your timer and close Outlook.
5. Record your results.
6. Log into the client machine as user **test2** and wait until you see the "Updating Inbox" message.
7. START your timer.
8. Observe the thermometer and wait for the new message to appear in the Inbox.
9. STOP your timer and WAN bandwidth monitor.
10. Record your results.
11. Log into the client machine as user **test** and hit **Send/Receive**. The previously sent message created an instance of the e-mail in the user's "Sent" folder that will eventually be sent to the Exchange server. Because this traffic could interfere with repeated Outlook tests, the synchronization is forced here.

Steelhead Appliance expected results

When the user **test** sends the message, the Steelhead appliance has never seen this data before and therefore you will experience cold performance.

When the user **test2** downloads the message, the Steelhead appliance already has the data and therefore can deliver LAN-like performance. If there were several more users in the branch (i.e. a bulk email scenario) each additional user would have experienced Warm performance.

While the traffic in step #11 is not figured into the overall test, this does represent additional redundant traffic created by Outlook 2003 Cache Mode.

Results

		Bandwidth		Seconds	
Baseline	Email send (#5)				
	Email receive (#10)				
		Steelhead Appliance		**Alternative DUT**	
		Bandwidth	Seconds	Bandwidth	Seconds
Cold	Email send (#5)				
Warm	Email receive (#10)				

[1] There have been reports of some environments where the "Outbox [1]" does not seem to update on its own and instead the user must perform some activity (highlighting messages in the inbox, etc.) in order to trigger a screen update. This is a behavior that has been observed in baseline testing and is in no way specific to Steelhead appliances. The root cause is as yet unknown, but be prepared for the possibility it may occur in your testing. If you encounter this issue, another way to measure this in an isolated lab environment is by observing a spike in WAN utilization as the message is being sent.

© 2012 Riverbed Technology. All rights reserved.

3. Encrypted MAPI

Outlook 2007 and 2010 both encrypt MAPI traffic by default. In order for the Steelhead to optimize this traffic, the server-side Steelhead will need to join the domain that the exchange server is a part of.

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for downloading email that contains attachments.

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Add a user to the Windows domain.
2. Add the server-side Steelhead to the domain.
3. From the GUI on both Steelheads, enable 'Encrypted MAPI'
4. Send an email, preferably with an attachment, in Outlook. Start timer.
5. STOP your timer and close Outlook.
6. Record your results. This is the 'cold' performance.
7. Open Outlook and send another email with the same attachment.
8. STOP your timer and WAN bandwidth monitor.
9. Record your results. This is the warm performance.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">As in the previous test, the user will experience LAN-like performance for MAPI traffic on the 'warm' pass.
SH RiOS version 5.5 has the ability to run all Microsoft Exchange application streamlining optimizations on encrypted Microsoft Exchange traffic.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>MAPI Encryption – send mail</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>MAPI Encryption - Email send</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>MAPI Encryption - Email send</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

4. Outlook Anywhere

Microsoft Outlook Anywhere lets clients that use Microsoft Office Outlook 2010, Outlook 2007, or Outlook 2003 connect to their Exchange servers from outside the corporate network or over the Internet using the RPC over HTTP Windows networking component.

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for downloading email that contains attachments.

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Enable Outlook Anywhere in Microsoft Office Outlook 2003/2007/2010, Microsoft Exchange Server, and the DUTs. A how-to video is available that details the steps involved.
2. Send an email, preferably with an attachment, in Outlook. Start timer.
3. STOP your timer and close Outlook.
4. Record your results. This is the 'cold' performance.
5. Open Outlook and send another email with the same attachment.
6. STOP your timer and WAN bandwidth monitor.
7. Record your results. This is the warm performance.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">As in the previous test, the user will experience LAN-like performance for Outlook Anywhere traffic on the 'warm' pass.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>Outlook Anywhere Email send</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>Outlook Anywhere - Email send</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>Outlook Anywhere - Email send</td><td></td><td></td><td></td><td></td></tr>
</table>

Testing Native Kerberos

With RiOS 7.0 encrypted MAPI, signed SMB, signed SMB2 and HTTP (IIS, Sharepoint, etc) traffic authenticated using Kerberos will be optimized natively by the Steelhead.

See the *Steelhead Management Console User's Guide* for detailed information on configuration of Native Kerberos. For other DUTs, consult the appropriate vendor for information on configuring Native Kerberos. The basic steps are:

1. In Windows Server:
    a. Select a user to give delegate control to.
    b. Assign that user *Replicate Directory Changes* and *Replicate Directory Changes All* privileges
2. Add the Server side Steelhead appliance to your domain
3. Navigate to *Configure -> Optimization -> Windows Domain Auth* and enter the *Username* and *Password* of the user in step 1

Native Kerberos is now configured and any of the tests in the Encrypted MAPI, SMB v2 Traffic, SMB Signed Traffic, and etc. sections can now be run.

© 2012 Riverbed Technology. All rights reserved.

Testing Lotus Notes

Sending/receiving memos with attachments

Sending a memo with an attachment is optimized by the Lotus Notes blade. This test is to verify that write optimization is performing as expected. Attachments of varying sizes can be used.

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for downloading email that contains attachments.

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Open Lotus Notes and create a new memo.
2. Address the memo to another user.
3. Provide the following subject: "Lotus Notes Write Optimization Test"
4. Provide the following body: "The Lotus Notes Optimization blade will perform write optimization on attached files in memos."
5. Attach a medium sized file (10MB).
6. Send the memo and time how long it takes.
7. On another Notes client verify that the memo has been received by the addressed user and is intact.
8. Save the attachment and time how long it takes to save.
9. Delete the memo on both clients.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">With the introduction of RiOS 5.5, Lotus Notes traffic can now benefit from optimization, including sending of attachments sent and replication of data.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>LotusNotes- send mail</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>LotusNotes – send/receive email</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>Lotus Notes – send/receive email</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

1. Sending/receiving calendar meetings with attachments

Creating a calendar entry with an attachment will trigger a call which is optimized by the Lotus Notes blade. This test is to verify that write optimization is performing as expected.

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for LotusNotes calendar meetings with attachments

Clear the DUT’s data store before the start of the test? Yes
Clear the DUT’s data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Open Lotus Notes and open the calendar
2. Create a new appointment.
3. Provide the following subject: "Calendar Attachment Test"
4. Provide the following location: "Here"
5. Provide the following description: "The Lotus Notes Optimization blade will optimize sending an appointment with an attachment."
6. Attach a small sized file (< 32KB) and uncheck the "Compress" option in the Create Attachment dialog.
7. START timer and WAN Bandwidth monitor.
8. Save and close the item.
9. STOP timer and WAN Bandwidth monitor. This is the cold performance.
10. Open up the item and modify the subject to be: "Calendar Attachment Test Altered"
11. Delete the Attachment.
12. START timer and WAN Bandwidth monitor.
13. Save and close the item.
14. STOP timer and WAN Bandwidth monitor. This is the warm performance.
15. Delete the item.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">Calendar entries are added, modified and removed successfully. Performance improvement is observed.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>LotusNotes- calendar attachment</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>LotusNotes – calendar attachment</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>Lotus Notes – calendar attachment</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

2. Server-Server Replication

This test is to verify that the Lotus Notes Domino Server to Server replication traffic benefits from optimization.
The allocation should be set up with 2 servers and 2 clients, in the server replication in-path configuration with 2 clients and 1 server on one side and 1 server on the other side.
The 2nd server needs to be configured as a peer to the first server.

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests Lotus Notes Domino server replication

Clear the DUT’s data store before the start of the test? No
Clear the DUT’s data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">On the 1st server add two replication events:

1. Open the Admin Client
2. Go to the Replication Tab and open Replication Events
3. Add a new Replication Event that is initiated by Server 1.
4. Add a new Replication Event that is initiated by Server 2.
5. Perform (Lotus notes) tests 1 and 2 as outlined in above section
6. Trigger the first Replication Event.
7. Once the replication has completed, repeat step 1.
8. Trigger the second Replication Event.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">Replication completes successfully, with enhanced performance compared to baseline</td></tr>
<tr><td rowspan="5">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>Lotus Notes Replication</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Warm</td><td>Lotus Notes Replication</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Testing HTTP

This section provides the following HTTP tests:

1. Testing Web-File Access (Cross-Application)
2. Testing Web Applications

Test #1 assumes that the data store is still warm from the "boomerang" Email test.

1. Testing Web-File Access (Cross-Application)

<table>
<tr><td>Use Case</td><td colspan="6">This procedure tests application acceleration for web file access.

Clear the DUT's data store before the start of the test? No
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Create a folder called web on the web server. Place one of the attachments from the "boomerang" email test in the web folder.
2. From a client desktop, open a web browser and navigate to a page where there is a link to the file
3. START your timer and WAN bandwidth monitor.
4. Right-click the file and save it to the desktop.
5. STOP your timer and WAN bandwidth monitor.
6. Record your results.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">The Steelhead appliances data stores were cleared at the start of the "boomerang" email test. The attachment data had only passed through the Steelhead appliances via the MAPI 2003 protocol.

When the client accesses the file via the HTTP protocol, the user gets warm performance even though it's the first time the data pattern has crossed the network via HTTP. The test demonstrated the protocol-agnostic nature of the Steelhead appliance algorithm.

(Optional) You can repeat this test with the FTP protocol and you will see Warm performance.</td></tr>
<tr><td rowspan="5">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>Web File Access (#6)</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Warm</td><td>Web File Access (#6)</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

2. Testing Web Applications

Use Case

This procedure tests application acceleration for web applications.

To prepare for this test, you should prepare a written sequence of actions that represent a typical action (or set of actions) for the application. For example, if the application itself takes a long time to load, the test could be simply logging in. If it is a CRM system, perhaps opening a customer record, opening the frequently asked questions (Fads) within the system, and answering the query for several cases. The script should contain the tasks for which users report slowness and/or tasks that are suspected to use large amounts of bandwidth.

Real browsers should be used to run this test and the performance should be measured by using an application such as HTTP Watch. A free version is available. Test tools such as IXIA/Spirent/Loadrunner should not be used in this test as they do not behave in the same way as real browsers.

RiOS 6.0 adds a new object pre-fetch table to achieve an additional 20x -50x performance gain when dealing with web applications.

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes

Procedure

1. Prepare your written sequence of actions.
2. Configure the DUT for specific HTTP optimization
3. START HTTPWatch and WAN bandwidth monitor.
4. Run through the sequence of actions.
5. STOP HTTPWatch and the WAN bandwidth monitor.
6. Record your results.

Steelhead Appliance expected results

The first time the sequence of actions is performed; the Steelhead appliances improve the efficiency of TCP and perform some compression, providing cold performance. In addition, the HTTP application streamlining module learns about the structure of the web page so that it can pre-fetch certain objects in the future.

The second time the sequence of actions is performed, the Steelhead appliances leverage the data patterns learned during the first pass and provide Warm performance. Furthermore, since the HTTP application streamlining module already knows about the structure of the web page, it pre-fetches the majority of the objects before the client requests them. When the client requests the data, the Steelhead appliance can then serve them out locally. In the Steelhead appliance administrative GUI, the HTTP Statistics report includes prefetch statistics that show the effectiveness of the HTTP application streamlining module.

Some applications may require fine tuning of the HTTP application streamlining module to perform at its best. However, if your organization already uses the following applications, configure the HTTP optimization per the table below:

HTTP Application	URL Learning	Parse and Prefetch	Metadata Response
SAP/Netweaver	No	No	**Yes**
Microsoft CRM	**Yes**	No	**Yes**
Agile	No	No	**Yes**
Pivotal CRM	No	**Yes**	**Yes**
Sharepoint	**Yes**	No	**Yes**

Note that this warm performance benefit occurs regardless of whether the web application is dynamic or static HTML: The Steelhead appliances provide acceleration while preserving full end-to-end communication between client and server.

Results

		Bandwidth	Seconds
Baseline	Web Application Performance (#5)		
	Web Application Performance (#10)		

© 2012 Riverbed Technology. All rights reserved.

			Steelhead Appliance		Alternative DUT	
			Bandwidth	Seconds	Bandwidth	Seconds
	Cold	Web Application Performance (#5)				
	Warm	Web Application Performance (#10)				

Testing HTTPS

Because of the extra configuration involved, creating a test environment that uses HTTPS is more complex than for HTTP. Similarly, configuring DUTs to accelerate operations in an HTTPS environment is also more complex than for HTTP. However, once set up, the methodology for testing HTTPS is the same as HTTP. Once you've set up your environment correctly, you can refer back to the HTTP section and use similar tests in your HTTPS environment.

See the *Steelhead Management Console User's Guide* for detailed information on configuration of HTTPS (SSL). For other DUT, consult the appropriate vendor for information on configuring HTTPS.

Some basic steps to configure HTTPS acceleration on a Steelhead appliance for use with an existing HTTPS server:

1. SSL acceleration on the Steelhead appliance requires a separate license. If you do not have that license, please contact Riverbed Technical Support or your Riverbed Sales Engineer.
2. On the "data center" Steelhead appliance, add the server into the Protocol: SSL→Servers section. The server side Steelhead can automatically pull down intermediate certs from backend server if needed, and update on failure of SSL handshake with client
3. Establish the peering trust relationship between the Steelhead appliances. This must be done on both Steelhead appliances in order for the Steelhead appliances to establish a secure connection between them.
4. Create an in-path rule on the "branch" Steelhead appliance to optimize the traffic destined for this new server. Specific parameters for this in-path rule are:
   - Port: 443
   - Pre-optimization Policy: SSL
   - Latency Optimization Policy: HTTP

   By default, the Steelhead appliance will pass through port 443 due to the fact that it's part of the "Secure" port-label. Therefore, you must ensure this rule appears at the top the in-path rules list.
5. You're now ready to test HTTPS using the same steps as HTTP. Note that if you've used a self-signed certificate (as described in step 1 above), your browser will present a warning about an "unknown CA". This is expected behavior as the browser will not have the Steelhead appliance's root CA certificate listed in its known list of root CA.

Testing Client Authentication

New in RiOS 6.5 is client authentication. For deployments that require client's to authenticate using a smart card, such as a Common Access Card (CAC), enabling the client authentication features allows the Steelhead appliance to optimize as usual while the client authenticates to the SSL server.

1. Follow steps in Testing HTTPS section above to enable SSL optimization.
2. On the server side Steelhead appliance, navigate to *Configure* -> *Optimization* -> *Advanced Settings* and select *Enable Client Certificate Support*.
3. On the server side Steelhead appliance, navigate to *Configure* -> *Optimization* -> *SSL Main Settings* and import the private key and certificate use by the SSL server.
4. Attempt to authenticate the client and perform using the smart card. Perform a test such as those in the Testing HTTP section above. Verify that the Steelhead is optimizing this traffic.

© 2012 Riverbed Technology. All rights reserved.

Testing Long-Term Data Reuse

This section includes tests to confirm the degree of long-term data reuse enabled by the DUT. This includes the following tests:

1. Selective Restore
2. Surviving a Power Outage

1. Selective Restore

<table>
<tr><td>Use Case</td><td colspan="6">A DUT's resources may be stressed by branches that perform "data heavy" operations over the WAN such as backup/restore, remote software installation, or active use from a large and varied user population. The amount of "cold" data that may be generated by this type of usage could quickly overrun the limited capacity of some DUT systems, causing many data patterns to be deleted in favor of new data before learned patterns can be effectively leveraged. To ensure that "warm" benefits are seen over the long term, it's important that a DUT's memory capacity be taxed so that disk capacity is also leveraged. This long-term reuse is tested here by transferring a large number of files to "warm" the DUTs, then selectively transferring some of the files a second time.

To prepare for the test, you should create a set of unique files, each one containing uncompressible data. The total size of the data set should be at least twice the size of the memory capacity of your DUT, which will ensure that the memory capacity of the DUT is overrun when the full set is transferred. If you have access to a Linux system, a simple script could be used to create the data set. For example, the following sample script creates 8 GB of test data (800 files of ten megabytes each) which would be certain to overrun a DUT with 4 GB of memory:

<pre>#!/bin/sh
NUMFILES=800 ; FILE=1
while [${FILE} -le ${NUMFILES}] ; do
  dd count=20000 if=/dev/urandom of=file${FILE}.dat
  FILE=`expr ${FILE} + 1`
done</pre>
The transfers described in this test are similar to operations that would be performed with a backup/restore application. If you don't have access to such software, the test could also be run by transferring the data using FTP or any other generic TCP-based data transfer protocol.

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? No</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Remove any bandwidth or latency settings on your WAN simulator (to save time during the initial warming phase)
2. Using your backup or file transfer application, transfer the entire data set from the local client to the server
3. Select a few files at random from the repository to restore (transfer back) from the server to the local client
4. Restore the bandwidth and latency settings on your WAN simulator
5. START your timer and WAN bandwidth monitor
6. Using your backup or file transfer application, transfer the selected files from the server back to the local client
7. When the files have been transferred back, STOP your timer.
8. Record your results.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">The Steelhead appliance has a disk-based data reduction mechanism to ensure long-term reusability of data. The most popular data patterns may often be available in memory for extra-quick access, but all data patterns are ultimately stored on disk as well. Even though the data set in this case represents far more than the Steelhead appliance could store in memory, "warm" access is still provided during the restore phase because the data patterns are recalled from disk.</td></tr>
<tr><td rowspan="4">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Warm</td><td>Restore (#8)</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

2. Surviving a Power Outage

<table>
<tr><td>Use Case</td><td colspan="6">In production, DUTs may be rebooted periodically due to power outages in remote branch locations or scheduled maintenance. Should this happen, a branch DUT should ideally not lose its knowledge of learned data patterns, otherwise the entire remote location would return to “cold” performance. Even more drastically, if a data center DUT was to lose this knowledge due to a reboot, all remote branch locations that depend on it would also return to “cold” performance.

This test ensures the knowledge of data patterns survives a power outage by once again transferring a selection of files from the server to the local client, this time after power cycling one of the systems.

Clear the DUT’s data store before the start of the test? No
Clear the DUT’s data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Power cycle one of the DUTs.
2. Once power is restored, ensure the DUT is fully booted and ready to provide optimization services
3. Select a few files at random from the repository to restore (transfer back) from the server to the local client
4. START your timer and WAN bandwidth monitor
5. Using your backup or file transfer application, transfer the selected files from the server back to the local client
6. When the files have been transferred back, STOP your timer.
7. Record your results.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">The Steelhead appliance has a disk-based data reduction mechanism to ensure long-term reusability of data. Even in the event of power loss, previously learned data patterns can be recalled from disk once power has been restored to the Steelhead appliance.</td></tr>
<tr><td rowspan="4">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Warm</td><td>Restore (#7)</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Testing Centralized Printing

<table>
<tr><td>Use Case</td><td colspan="6">When print servers are located in the data center, print traffic from branch offices will need to be sent out over the WAN. The following procedure will test the performance of these print operations.

Clear the DUT’s data store before the start of the test? Yes
Clear the DUT’s data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. Enable print optimization on both DUTs.
2. START timer and WAN bandwidth monitor.
3. Print a document.
4. Wait for the document to start printing.
5. STOP timer and WAN bandwidth monitor.
6. Record your results. This is the ‘cold’ performance.
7. START timer and WAN bandwidth monitor.
8. Print the same document.
9. Wait for the document to start printing.
10. STOP timer and WAN bandwidth monitor.
11. Record your results. This is the warm performance.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">With RiOS 6.0, Windows print operations are accelerated as any other CIFS operations. The first print will receive cold performance while the second print operation will receive warm LAN-like performance.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>Print document</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Cold</td><td>Print document</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Warm</td><td>Print document</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Testing Citrix ICA Optimization

To consolidate operations, some organizations install thin clients in their branch offices and install a Citrix XenApp Server in the data center to front-end the applications. The proprietary protocol that Citrix uses to move updates between the client and the server is called ICA (Independent Computing Architecture).

This section provides the following tests:

1. Testing Basic Citrix ICA Optimization
2. Testing Citrix CDM Optimization

1. Testing Basic Citrix ICA Optimization

RiOS 6.0 and later will seamlessly optimize traffic within an ICA session. Please consult the Riverbed Deployment Guide for instructions on configuring Steelhead to work with Citrix ICA.

<table>
<tr><td>Use Case</td><td colspan="4">This test involves streaming a video file

Prerequisites:

1. Obtain a reasonable size WMV format movie file (preferably with a bit rate that's just under the bandwidth of the WAN link) and the length of the movie should run for several minutes. Place this file in appropriate directory so that the client can stream this file.
2. Set the WAN simulator to the appropriate link speed and latency

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="4">1. STOP the optimization service on the DUT.
2. Open up the desktops of a user.
3. START the WAN Bandwidth Monitor.
4. Begin streaming the video.
5. After the video finishes, STOP the WAN Bandwidth Monitor.
6. Record your results, this is the Baseline performance.
7. START the optimization service on the DUT
8. Enable Citrix ICA optimization.
9. START the WAN Bandwidth Monitor.
10. Begin streaming the video.
11. After the video finishes, STOP the WAN Bandwidth Monitor.
12. Record your results, this is the Cold performance.
13. START the WAN Bandwidth Monitor.
14. Begin streaming the video.
15. After the video finishes, STOP the WAN Bandwidth Monitor.
16. Record your results, this is the Warm performance.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="4">The Steelhead optimizes the ICA session. There should be an improvement in between Baseline and Cold and a larger improvement between Cold and Warm.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td></tr>
<tr><td>Baseline</td><td>Streaming</td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Bandwidth</td></tr>
<tr><td>Cold</td><td>Streaming</td><td></td><td></td></tr>
<tr><td>Warm</td><td>Streaming</td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

2. Testing Citrix CDM Optimization

Citrix Client Drive Mapping (CDM) gives the user the ability to mount a local client drive in the session to the remote server. RiOS 7.0 improves the latency of copying files to and from the remote server, in addition to performing SDR.

Use Case	This test involves copying a file from the remote server to a local drive. Prerequisites: 1. Stage a large file to download from the ICA session's local drive to the client's local drive mapped inside the ICA session 2. Set the WAN simulator to the appropriate link speed and latency **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**
Procedure	1. STOP the optimization service on the DUT. 2. Open up the desktops a separate user. 3. START Timer and WAN Bandwidth Monitor. 4. Start the large file transfer to the local drive mapped inside the ICA session of one of the users. 5. Wait for the file transfer to complete 6. STOP timer and WAN bandwidth monitor 7. Record your results. This is the *Baseline* performance. 8. START the optimization service on the DUT 9. Enable Citrix CDM optimization on both Steelheads. This can be done in the Steelhead Management Console by navigating tp *Configure -> Optimization -> Citrix* and selecting *Enable Citrix CDM Optimization* 10. START timer and WAN bandwidth monitor. 11. Repeat the large file transfer. 12. STOP timer and WAN bandwidth monitor 13. Record your results. This is the *CDM Cold* performance. 14. START timer and WAN bandwidth monitor. 15. Once again repeat the large file transfer 16. STOP timer and WAN bandwidth monitor 17. Record your results. This is the *CDM Warm* performance.
Steelhead Appliance expected results	Enabling Citrix CDM optimization should lead to an impressive performance improvement and bandwidth reduction.

Results

		Bandwidth		Seconds	
Baseline	File Copy				
		Steelhead Appliance		**Alternative DUT**	
		Bandwidth	Seconds	Bandwidth	Seconds
CDM Cold	File Copy				
CDM Warm	File Copy				

© 2012 Riverbed Technology. All rights reserved.

Testing SDRF Optimization

EMC Symmetrix Remote Data Facility (SRDF) family of software is one of the industry's most powerful suites of remote storage replication solutions for disaster recovery and business continuity. The Steelhead provides native optimization of SRDF protocol, allowing your replication operations to both complete faster and take up less WAN bandwidth. Selective SRDF optimization, introduced in RiOS 6.5, further improves performance with the ability to fine tune performance on a per RDF-group basis. Some common RDF groups can be selectively optimized as such:

- Oracle Logs: *LZ Compression only* will be used as it the data is highly compressible but not redundant
- VM Images: *SDR-Default* will be used as the data is both compressible and redundant
- Encrypted/Compressed Images: *No Optimization* will be used as the data is not compressible or redundant

The below test will assume the above three RDF groups.

Use Case	Testing SRDF optimization **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**
Procedure	1. STOP the optimization service on the DUT 2. START timer and WAN bandwidth monitor 3. Initiate an SRDF replication 4. STOP timer and WAN bandwidth monitor 5. Record your results. This is the *Baseline* performance. 6. Enable SRDF optimization on both DUTs. This can be done on the Steelhead management console by navigating to *Configure -> Optimization -> SRDF* and selecting *Enable SRDF*. 7. START timer and WAN bandwidth monitor. 8. Repeat the SRDF replication 9. STOP timer and WAN bandwidth monitor. 10. Record your results. This is the *Base SRDF* performance. 11. Configure selective SRDF optimization. This can be done on the Steelhead CLI by entering: • protocol srdf symm id <*symm id*> rdf group <group> optimization lz-only description OracleLogs • protocol srdf symm id <*symm id*> rdf group <group> optimization sdr-default description VMImages • protocol srdf symm id <*symm id*> rdf group <group> optimization none description EncryptedImages 12. START timer and WAN bandwidth monitor. Clear the DUT's data store 13. Repeat the SRDF replication. 14. STOP timer and WAN bandwidth monitor. 15. Record your results. This is the *Selective SRDF* performance.
Steelhead Appliance expected results	Enabling SRDF optimization will greatly increase performance over the Baseline case. Configuring Selective SRDF optimization will increase performance even further. Clearing the data store on step 12 ensures that the additional performance increase is not because the DUT is warm but because of the built in optimizations.
Results	

		Bandwidth		Seconds	
Baseline	SRDF replication				
		Steelhead Appliance		**Alternative DUT**	
		Bandwidth	Seconds	Bandwidth	Seconds
Base SRDF	SRDF replication				
Selective SRDF	SRDF replication				

© 2012 Riverbed Technology. All rights reserved.

Testing Satellite Optimization

Satellite links are used extensively for WAN connectivity for people in remote locations or in mobile environments such as ships or an oil rigs. Due to its nature, Satellite links have several unique characteristics that can lead to transmission errors, packet loss, and intermittent connectivity, all of which adversely affect performance. TCP satellite optimization, native in RiOS 6.5, attempts to alleviate these issues and improve overall throughput for critical resources over satellite links, including, MAPI, CIFS, HTTP, and FTP traffic.

In addition, to address satellite transmission errors and intermittent connectivity, satellite optimization in RiOS 6.5 includes the lossy link error recovery feature. This feature shields TCP from any packet loss that has not been corrected by the satellite modem and helps to improve overall throughput.

This section contains the following tests:

1. Testing Basic Satellite
2. Testing Basic Satellite Optimization with Error Recovery

1. Testing Basic Satellite Optimization

This section tests basic satellite optimization.

Use Case	Testing Basic Satellite optimization **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**					
Procedure	1. STOP the optimization service on the DUT 2. START timer and WAN bandwidth monitor 3. Initiate an FTP, CIFS, or some other form of data transfer 4. STOP timer and WAN bandwidth monitor 5. Record your results. This is the *Baseline* performance. 6. Enable Satellite optimization on the client side DUT. This can be done on the Steelhead CLI by entering: *tcp sat-opt bw-est mode always*. 7. START timer and WAN bandwidth monitor. 8. Repeat the data transfer 9. STOP timer and WAN bandwidth monitor. 10. Record your results. This is the *Basic Satellite Optimization cold* performance. 11. START timer and WAN bandwidth monitor. 12. Repeat the data transfer. 13. STOP timer and WAN bandwidth monitor. 14. Record your results. This is the *Basic Satellite Optimization warm* performance.					
Steelhead Appliance expected results	Enabling Satellite optimization will greatly increase performance over the Baseline case.					
Results			Bandwidth		Seconds	
	Baseline	Data Transfer				
			Steelhead Appliance		**Alternative DUT**	
			Bandwidth	Seconds	Bandwidth	Seconds
	Basic Satellite Optimization cold	Data Transfer				
	Basic Satellite Optimization warm	Data Transfer				

© 2012 Riverbed Technology. All rights reserved.

2. Testing Basic Satellite Optimization with Error Recovery

This section tests basic satellite optimizations with error recovery. This is useful for lossy satellite links.

Use Case	Testing Basic Satellite optimization with error recovery **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**
Procedure	1. STOP the optimization service on the DUT 2. START timer and WAN bandwidth monitor 3. Initiate an FTP, CIFS, or some other form of data transfer 4. STOP timer and WAN bandwidth monitor 5. Record your results. This is the *Baseline* performance. 6. Enable Satellite optimization on the client side DUT. This can be done on the Steelhead CLI by entering: *tcp sat-opt bw-est mode always*. 7. Configure Error Recovery on the client side DUT. This can be on the Steelhead CLI by entering: *tcp err-recovery loss-recovery mode always* 8. START timer and WAN bandwidth monitor. 9. Repeat the data transfer 10. STOP timer and WAN bandwidth monitor. 11. Record your results. This is the *Satellite Optimization with Error Recovery cold* performance. 12. START timer and WAN bandwidth monitor. 13. Repeat the data transfer. 14. STOP timer and WAN bandwidth monitor. 15. Record your results. This is the *Satellite Optimization with Error Recovery warm* performance.
Steelhead Appliance expected results	Configuring Error Recovery will further improve performance on a lossy satellite link.
Results	

		Bandwidth		Seconds	
Baseline	Data Transfer				
		Steelhead Appliance		**Alternative DUT**	
		Bandwidth	Seconds	Bandwidth	Seconds
Satellite Optimization with Error Recovery cold	Data Transfer				
Satellite Optimization with Error Recovery warm	Data Transfer				

© 2012 Riverbed Technology. All rights reserved.

Testing Satellite Optimization with the SCPS Protocol Suite

RiOS 7.0 now integrates the SCPS protocol suite. SCPS is designed to allow communication over challenging environments. Originally, it was developed jointly by NASA and DOD's USSPACECOM to meet their various needs and requirements. Through a collaborative, multiyear R&D effort, the partnership created the Space Communications Protocol Standards-Transport Protocol (SCPS-TP)

Using SCPS provides additional performance benefits on top of the optimization discussed in the

© 2012 Riverbed Technology. All rights reserved.

Testing Satellite Optimization section above. Usage of SCPS functionality with RiOS requires an additional license. Previously this functionality was only available through the addition of an RSP package.

This section contains the following tests:

1. Testing Per-Connection SCPS
2. Testing Error-Tolerant SCPS

1. Testing Per-Connection SCPS

This section tests the performance of the SCPS protocol suite integrated with RiOS 7.0.

Use Case	Testing Per-Connection SCPS integration **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**					
Procedure	1. STOP the optimization service on the DUT 2. START timer and WAN bandwidth monitor 3. Initiate an FTP, CIFS, or some other form of data transfer 4. STOP timer and WAN bandwidth monitor 5. Record your results. This is the *Baseline* performance 6. Enable SCPS Satellite optimization on the client side DUT. This can be done on the Steelhead Management Console by navigating to *Configure -> Optimization -> Transport Settings* and enabling *Skipware Per-Connection TCP Optimization* 7. START timer and WAN bandwidth monitor 8. Repeat the data transfer 9. STOP timer and WAN bandwidth monitor 10. Record your results. This is the *SCPS cold* performance 11. START timer and WAN bandwidth monitor. 12. Repeat the data transfer 13. STOP timer and WAN bandwidth monitor 14. Record your results. This is the *SCPS warm* performance					
Steelhead Appliance expected results	Enabling SCPS optimization will greatly increase performance over the Baseline case.					
Results			Bandwidth		Seconds	
	Baseline	Data Transfer				
			Steelhead Appliance		**Alternative DUT**	
			Bandwidth	Seconds	Bandwidth	Seconds
	SCPS cold	Data Transfer				
	SCPS warm	Data Transfer				

© 2012 Riverbed Technology. All rights reserved.

2. Testing Error-Tolerant SCPS

This section tests Error Tolerant SCPS integrated with RiOS 7.0. This is useful for lossy satellite links

<table>
<tr><td>Use Case</td><td colspan="6">Testing Per-Connection SCPS integration

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="6">1. STOP the optimization service on the DUT
2. START timer and WAN bandwidth monitor
3. Initiate an FTP, CIFS, or some other form of data transfer
4. STOP timer and WAN bandwidth monitor
5. Record your results. This is the Baseline performance
6. Enable SCPS Satellite optimization on the client side DUT. This can be done on the Steelhead Management Console by navigating to Configure -> Optimization -> Transport Settings and enabling Skipware Error-Tolerant TCP Optimization
7. START timer and WAN bandwidth monitor
8. Repeat the data transfer
9. STOP timer and WAN bandwidth monitor
10. Record your results. This is the SCPS cold performance
11. START timer and WAN bandwidth monitor.
12. Repeat the data transfer
13. STOP timer and WAN bandwidth monitor
14. Record your results. This is the SCPS warm performance</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="6">Configuring Error Tolerant SCPS will further improve performance on a lossy satellite link.</td></tr>
<tr><td rowspan="6">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td><td colspan="2">Seconds</td></tr>
<tr><td>Baseline</td><td>Data Transfer</td><td colspan="2"></td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Seconds</td><td>Bandwidth</td><td>Seconds</td></tr>
<tr><td>Error Tolerant SCPS cold</td><td>Data Transfer</td><td></td><td></td><td></td><td></td></tr>
<tr><td>Error Tolerant SCPS warm</td><td>Data Transfer</td><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Testing Live Microsoft Silverlight Video

Microsoft Silverlight Smooth Streaming is an adaptive streaming protocol that runs entirely over HTTP. Microsoft Silverlight is adaptive in that it delivers the best quality stream possible to each viewer. Viewers with higher bandwidth will receive a higher quality stream, while lower bandwidth viewers will receive a lower quality stream. In addition it is responsive to changing conditions, so if the high bandwidth user initiates a large download, the quality of the stream will be lowered to accommodate.

It does this by splitting the video up into small chunks, approximately 2 seconds in length each. Each 2 second chunk is simultaneously encoded at multiple quality levels. If the viewer is not able to handle the video properly at any quality level, the next 2 second chunk will be served up at the lower quality level. And vice versa if the viewer is able to handle a higher quality stream.

New in RiOS 7.0 is the ability to optimize live Microsoft Silverlight video. When a live video is being watched by multiple users, they will all be requesting the same 2 second chunk of video at roughly the same time. In RiOS 7.0, only the first of these requests will be sent out across the WAN, the rest will be held back by the Steelhead. Once the first request is satisfied the Steelhead will duplicate it and send it out to all the other viewers.

<table>
<tr><td>Use Case</td><td colspan="4">Test Live Microsoft Silverlight Video

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="4">1. Configure a Microsoft Silverlight Smooth Streaming server and set up a live stream. An easy to follow walkthrough and full deployment guide are available.
2. START WAN bandwidth monitor
3. Start the stream on two client machines.
4. STOP WAN bandwidth monitor
5. Record your results. This is the Baseline performance.
6. Enable HTTP Stream Splitting. This can be done on the Steelhead GUI by navigating to Configure -> Optimization -> HTTP.
7. START WAN bandwidth monitor.
8. Start the stream on two client machines
9. STOP WAN bandwidth monitor.
10. Record your results. This is the Steelhead Optimized performance.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="4">With two clients and HTTP Stream Splitting enabled the WAN bandwidth consumption should be halved. Any number of clients can be added without increasing WAN bandwidth consumption.</td></tr>
<tr><td rowspan="5">Results</td><td colspan="2"></td><td colspan="2">Bandwidth</td></tr>
<tr><td>Baseline</td><td>Live Streaming</td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Bandwidth</td><td>Bandwidth</td></tr>
<tr><td>Steelhead Optimized</td><td>Live Streaming</td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Testing Live Adobe Flash HTTP Dynamic Streaming Video

Adobe Flash HTTP Dynamic Streaming is an adaptive streaming protocol that functions in a similar manner to Microsoft Silverlight described in the Testing Live Microsoft Silverlight Video section above.

New in RiOS 7.0 is the ability to optimize live Adobe Flash HTTP Dynamic Streaming video. When a live video is being watched by multiple users, they will all be requesting the same 2 second chunk of video at roughly the same time. In RiOS 7.0, only the first of these requests will be sent out across the WAN, the rest will be held back by the Steelhead. Once the first request is satisfied the Steelhead will duplicate it and send it out to all the other viewers.

Use Case	Test Live Adobe HTTP Dynamic Streaming Video **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**			
Procedure	1. Configure a Microsoft Silverlight Smooth Streaming server and set up a live stream. An easy to follow Tutorial is available 2. START WAN bandwidth monitor 3. Start the stream on two client machines. 4. STOP WAN bandwidth monitor 5. Record your results. This is the *Baseline* performance. 6. Enable *HTTP Stream Splitting*. This can be done on the Steelhead GUI by navigating to *Configure -> Optimization -> HTTP*. 7. START WAN bandwidth monitor. 8. Start the stream on two client machines 9. STOP WAN bandwidth monitor. 10. Record your results. This is the *Steelhead Optimized* performance.			
Steelhead Appliance expected results	With two clients and *HTTP Stream Splitting* enabled the WAN bandwidth consumption should be halved. Any number of clients can be added without increasing WAN bandwidth consumption.			
Results			Bandwidth	
	Baseline	Live Streaming		
			Steelhead Appliance	**Alternative DUT**
			Bandwidth	Bandwidth
	Steelhead Optimized	Live Streaming		

© 2012 Riverbed Technology. All rights reserved.

Testing IPv6 and UDP optimization

RiOS 7.0 performs SDR packet optimizations on TCP traffic over IPv6 and UDP traffic over IPv4. UDP based applications such as Aspera, TFTP, and older versions of NFS are now accelerated by the Steelhead Appliance.

Use Case	This procedure tests application acceleration for UDP and IPv6 traffic **Clear the DUT's data store before the start of the test? Yes** **Clear the DUT's data store at the end of the test? Yes**
Procedure	1. STOP the optimization service on the DUT 2. START timer and WAN bandwidth monitor 3. Initiate a TFTP, NFS, or some other form of UDP data transfer 4. STOP timer and WAN bandwidth monitor 5. Record your results. This is the *Baseline* performance. 6. On both Steelheads, navigate to *Configure -> Optimization -> General Service Settings* and *Enable Packet Mode Optimization* 7. On both Steelheads, navigate to *Configure -> Optimization -> In-Path Rules* and select *Add a New In-Path Rule*. Fill in the fields as follows: a. **Type**: Fixed-Target (Packet Mode Optimization) b. **Protocol**: UDP c. **Channel Mode**: Per-Flow d. **Target Appliance IP Address**: Remote Steelhead IP Address 8. START timer and WAN bandwidth monitor. 9. Initiate the same UDP transfer 10. Wait for the transfer to complete. 11. STOP timer and WAN bandwidth monitor. 12. Record your results. This is the *Cold* performance. 13. START timer and WAN bandwidth monitor. 14. Initiate the same UDP transfer 15. Wait for the transfer to complete. 16. STOP timer and WAN bandwidth monitor. 17. Record your results. This is the *Warm* performance.
Steelhead Appliance expected results	The connection should be shown as a fully optimized connection

Results

		Bandwidth		Seconds	
Baseline	UDP Transfer				
		Steelhead Appliance		**Alternative DUT**	
		Bandwidth	Seconds	Bandwidth	Seconds
Cold	UDP Transfer				
Warm	UDP Transfer				

© 2012 Riverbed Technology. All rights reserved.

Testing Functionality

In addition to improving performance and reducing bandwidth utilization, a DUT must include many additional features in order to integrate into many production environments. While it may not be feasible to test every feature of a DUT, you should ensure that you test the functions that are most important for your needs.

While the Steelhead appliance includes features to address each one of these functional areas, many other DUT may not. Consult the appropriate vendor for information on how to address your needs in each area.

Ease of Deployment

As your production network grows, you will inevitably need to add appliances to additional locations. An additional DUT should be easy to integrate into a network that already contains other DUT units. If a DUT requires explicit configuration of tunnels or remote subnet information, you may find that each DUT becomes successively more difficult to integrate into the network and simply does not scale.

To test this, reconfigure your network to create a second "branch" location as shown in the diagram below. If you are using a WAN simulator, you will need to replace it with some sort of router that has at least three ports. Establish routing and confirm that each "branch" location can connect to the "data center" location, and that the "branches" can also connect to each other.

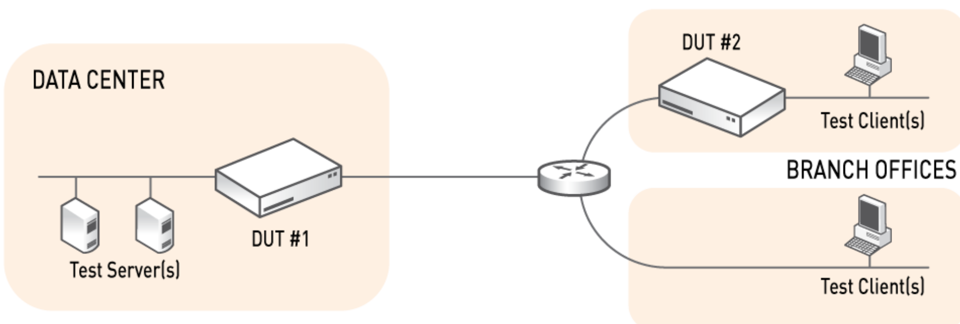

Once the network is configured and tested in this topology, add a fresh DUT to the network as shown in the diagram below. For adding a Steelhead appliance, follow the instructions in the *Steelhead Appliance Getting Started Guide*. You may wish to time how long it takes you to fully configure and test the system at the new location.

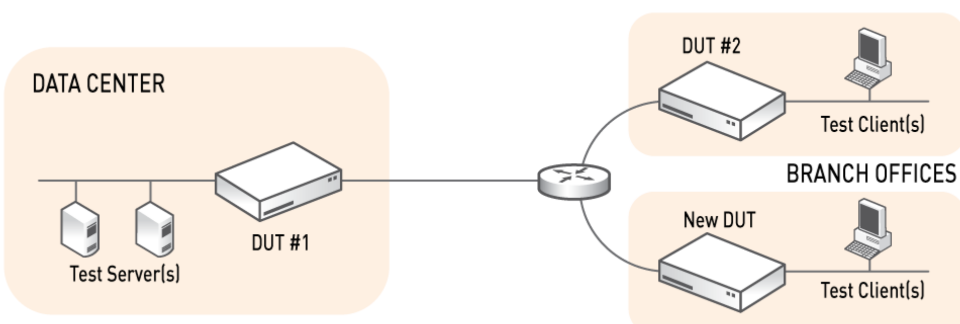

After completing the minimal configuration, establish connections between each "branch" location and the "data center" location, as well as a connection between the two "branch" locations. Confirm through the administrative interface of the DUT that all connections are optimized.

Use Case	Testing ease of deployment
Steelhead Appliance expected	You should find that very little configuration is required to bring up the additional location. You will also find that "any-to-any" connectivity between locations is established without the need to configure the Steelhead appliances with routing information other than a gateway.

© 2012 Riverbed Technology. All rights reserved.

results				
Results	**Steelhead Appliance**		**Alternative DUT**	
	Length of time to bring additional DUT online	Pass/Fail	Length of time to bring additional DUT online	Pass/Fail

WAN Visibility Modes/Network Transparency

In most situations, it is desirable to be able to distinguish the pre-optimized traffic (LAN) from the post-optimized traffic (WAN). The most common way of doing this is to use different IP addresses between the pre-optimized traffic and post-optimized traffic. In some very specific situations, however, it may be necessary to have the pre-optimized traffic appear the same way as post-optimized traffic or be able to keep the destination port intact.

There are three situations where it may be necessary to keep the TCP headers intact for pre and post-optimized traffic:

1. If the DUT doesn't support Netflow and needs to rely on the router to collect Netflow data
2. If the DUT doesn't support QoS and must rely on router for QoS functionality
3. The network administrator would like to continue to use the QoS functionality on the router even though the DUT supports QoS.

The final case is particularly true when the router classifies traffic based on the original IP address and port numbers and cannot be modified. For routers that enforce QoS based on destination port only, it is only necessary to keep the destination port intact for the post-optimized traffic while using the DUT's own IP address as the source and destination IP.

Refer to the whitepaper titled *Understanding the Hidden Caveats of Network Transparency*, for more information regarding the caveats associated with network transparency. Given the various scenarios, a DUT should have the flexibility to choose which addressing is the most appropriate for that particular network.

Use Case	Testing the different WAN visibility modes and how they work within your network.			
Procedure	1. Place a sniffer between the Steelheads 2. Login to the Steelhead appliance's GUI as "admin" 3. Click on Configure > Optimization > In-path Rules > Add a New In-Path Rule 4. Change the "WAN Visibility Mode" to "Full Transparency" 5. Click on "Add" 6. Initiate a connection from the client to the server and observe the traffic traversing the WAN. 7. Repeat the above changing the "WAN Visibility Mode" to "Port Transparency" and "Correct Addressing"			
Steelhead Appliance expected results	The default mode for the Steelhead appliance is to use something known as "correct addressing" (CA). CA uses the Steelhead appliance's own IP address for the post-optimized traffic. This makes the Steelhead appliance simple to deploy in a complex network without having to worry about any routing issues. The second option is "port transparency" (PT). PT keeps the original destination port of the connection but uses the Steelhead appliance's own IP address for the post-optimized traffic. Once again, this allows for simple deployment within a complex environment without having to worry about routing issues. Finally, there is "full transparency" (FT). FT provides the ability to make the post-optimized traffic appear the same as the pre-optimized traffic by keeping the original source/destination ip/port.			
Results	**Steelhead Appliance**		**Alternative DUT**	
	Ability to configure different WAN visibility modes	Pass/Fail	Ability to configure different WAN visibility modes	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Role-Based Management

There may be times when you want to grant particular users read/write access only to certain parts of the DUT. For example, you may want to grant read/write access to the Exchange administrator for items related to MAPI, but read-only access to the rest of the DUT. Or, for auditing reasons, you may want individual logins instead of sharing a common login amongst all the users. A DUT should be able to provide granular access control.

Some basic steps to test the different WAN visibility modes with a Steelhead appliance:

<table>
<tr><td>Use Case</td><td colspan="4">Testing Role-based management capabilities.</td></tr>
<tr><td>Procedure</td><td colspan="4">1. Login to the Steelhead appliance's GUI as "admin"
2. Click on Configure > Security > User Permissions > Add a New User
3. For "Account name", enter "jdoe"
4. Click on "Use a Password", and enter "12345". Reconfirm the password
5. Under "Roles and Permissions", select "Read-only" for "MAPI"
6. Click on "Add"
7. Logout of the Steelhead and login as "jdoe"
8. Click on Configure > Optimization > In-path rules and notice that an error message appears stating that you do not have permissions to view this page.
9. Click on Configure > Optimization > MAPI
10. Notice that you can view this page but notice that you cannot change any of the settings.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="4">Since the release of RiOS 5.0, it is possible to create individual users with different roles. For example, it's possible to create a role for the Exchange administrator to only allow him/her access to the MAPI specific items within the Steelhead appliance.</td></tr>
<tr><td rowspan="3">Results</td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Ability to configure different user roles</td><td>Pass/Fail</td><td>Ability to configure different user roles</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Caching DNS

A potential side effect of centralizing servers from a remote location back to the data center is the increase of traffic for services such as DNS queries. To minimize WAN congestion, a DUT should be able to reduce the number of DNS queries traversing the WAN.

Some basic steps to test the caching DNS feature with a Steelhead appliance:

<table>
<tr><td>Use Case</td><td colspan="4">Testing the capability of the DUT to reduce the number of DNS queries</td></tr>
<tr><td>Procedure</td><td colspan="4">1. Login to the Steelhead appliance’s GUI as “admin”
2. Click on Configure > Branch Services > Caching DNS
3. Check the box for “Enabling Caching DNS” and “Primary Interface Responding to DNS Requests”
4. Click on “Apply”
5. Click on “Add a New DNS Name Server”
6. Enter the IP address of a DNS server and click “Add”
7. Change the IP address of the PC to point to the IP address of the primary interface of the Steelhead
8. Launch your web browser and navigate to your intranet.
9. Log back into the your browser and click on Reports | Branch Services | DNS Cache Hits
10. A cache hit means that the DNS request was served locally from the Steelhead and the request didn’t traverse the WAN.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="4">The Steelhead appliance supports the caching of DNS entries. While it doesn’t act as an authoritative server, it does cache DNS entries to reduce the amount of DNS queries across the WAN.</td></tr>
<tr><td rowspan="3">Results</td><td colspan="2">Steelhead Appliance</td><td colspan="2">Alternative DUT</td></tr>
<tr><td>Ability to offload DNS requests</td><td>Pass/Fail</td><td>Ability to offload DNS requests</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Redundant Power

DUT units deployed in a data centers are particularly vital to the network because they often support optimization functions for many remote locations. As a result, a data center DUT should be resilient in the face of issues such as a partial loss of power or a power supply failure. Steelhead 3000-class, 5000-class, and 6000-class systems are designed for this type of data center environment, as they contain redundant power supplies.

To test redundant power with a data center class DUT, follow these steps:

Use Case	Testing the resiliency of the DUT in the face of power related issues.	
Procedure	1. Cable the DUT to two separate power sources. 2. While the DUT is up and running, disable the power source going to one of the DUT power supplies. 3. Listen for an audible alarm and confirm that the DUT is still functioning normally. 4. Restore the power source you previously disabled. 5. Confirm that the DUT is still functioning normally. 6. While the DUT is up and running, pull one of the power supplies completely out of the DUT. 7. Confirm that the DUT is still functioning normally. 8. Restore the power source you disabled. 9. Confirm that the DUT is still functioning normally.	
Steelhead Appliance expected results	The Steelhead appliance should alarm when it loses power and continue to function normally. You should also find you are able to replace the power supply while the system is running.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Pre-population

To save end users from enduring a "cold" hit, you may wish to specify that the data stores of the DUT be "warmed" in advance with certain data. This operation may be performed on a scheduled basis so that it can be performed during times when WAN links are otherwise idle.

This type of operation can be performed on the Steelhead appliance using the Pre-population feature. See the *Steelhead Management Console User's Guide* for detailed information on configuration of Pre-population. For other DUTs, consult the appropriate vendor for information on configuring population.

This section contains the following tests:

1. CIFS Pre-population
2. HTTP Pre-population
3. Microsoft Silverlight Pre-population

1. CIFS Pre-population

Some basic steps to test CIFS Pre-population with a Steelhead appliance:

<table>
<tr><td>Use Case</td><td colspan="2">Testing the pre-population capabilities of the DUT</td></tr>
<tr><td>Procedure</td><td colspan="2">1. Create a folder on the data center server and enable it for Sharing (e.g. "\\server\prepop"). Populate the folder with several documents. Make the same documents available in another location, such as on another network share or through another protocol such as HTTP, FTP, or e-mail.
2. Clean the data stores on both your "branch" and "data center" Steelhead appliances.
3. Enable the Pre-population service on your "branch" Steelhead appliance.
4. Add a new Pre-population share on your "branch" Steelhead appliance, specifying the share on your data center server as the Remote Path. Set the "Sync Interval" to 1 minute. Save the configuration.
5. Configure the Pre-population share you just created for "Sync Enable". Confirm that the Status changes to "Get initial copy in progress…", indicating that the initial pre-population has begun. Confirm through the Current Connections report in the administrative GUI that you see an optimized CIFS connection from your "branch" Steelhead appliance's primary IP address to your data center server's IP address.
6. Once the Pre-population shares screen in the administrative GUI reflects that the share has a status of "Share idle", the initial synchronization has completed. Access the documents across the simulated WAN through the method you prepared in advance (file share, HTTP, FTP, etc.) and ensure that you experience "warm" performance.
7. Observe the status for your share in the administrative GUI for a minute and observe that it briefly changes to a "SYNC" status, ensuring that any additional documents that have been added to the share will also be pre-populated.
8. To ensure there's no synchronization traffic on the WAN during future tests, you may wish to remove the Pre-population configuration upon completion of the test.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">The Pre-population service ensures that the data store of the Steelhead appliances are always "warmed" by the data contained in the documents on the data center share. Because the pre-population is performed over an efficient, optimized connection, the same data sequences will never have to cross the WAN twice, even if there's similarity among the documents in the Pre-population share. Also, because the Steelhead appliance's SDR disk-based data reduction technology is protocol-independent, the "warm" data sequences learned through pre-population can be leveraged when the same data sequences appear in any other data protocol (HTTP, FTP, e-mail, etc.)

While we used a short 1-minute sync interval for testing purposes, note that this can easily be changed to suit the needs of your production environments, e.g. performing pre-population at night during idle periods on the network.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

2. HTTP Pre-population

HTTP Pre-population is new in RiOS 7.0. The section tests the HTTP Pre-population capabilities of the DUT.

<table>
<tr><td>Use Case</td><td colspan="2">Testing the HTTP pre-population capabilities of the DUT</td></tr>
<tr><td>Procedure</td><td colspan="2">1. Deploy a web server, such as Apache or Microsoft IIS, in the data center.
2. Create a folder on the data center web server and populate the folder with several documents. Make the same documents available in another location, such as as FTP, CIFS, or e-mail.
3. Clean the data stores on both your “branch” and “data center” Steelhead appliances.
4. Initiate a pre-population operation on the branch Steelhead appliance. This is done by navigating to the CLI and entering:
a. <i>conf t</i>
b. <i>protocol http prepop list</i> <<i>name</i>> <i>url</i> http://<server>/<file>
c. Repeat the previous step for every file to be pre-populated
d. <i>protocol http prepop list</i> <<i>name</i>> <i>start</i>
5. The CLI will be unresponsive as the files are pre-populated.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">The Pre-population service ensures that the data store of the Steelhead appliances are always “warmed” by the data contained in the documents on the data center share. Because the pre-population is performed over an efficient, optimized connection, the same data sequences will never have to cross the WAN twice, even if there’s similarity among the documents in the Pre-population share. Also, because the Steelhead appliance’s SDR disk-based data reduction technology is protocol-independent, the “warm” data sequences learned through pre-population can be leveraged when the same data sequences appear in any other data protocol (CIFS, FTP, e-mail, etc.)

While we initiated the pre-population immediately in this example, a job can be set to initiate the pre-population at a time you specify. Please refer to the Steelhead Appliance Deployment Guide for details on creating a job.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

3. Microsoft Silverlight Pre-population

Microsoft Silverlight Pre-population is new in RiOS 7.0. Please refer to the Testing Live Microsoft Silverlight section for an overview of Microsoft Silverlight. The section tests the Microsoft Silverlight Pre-population capabilities of the DUT.

<table>
<tr><td>Use Case</td><td colspan="2">Testing the HTTP pre-population capabilities of the DUT</td></tr>
<tr><td>Procedure</td><td colspan="2">1. Configure a Microsoft Silverlight Smooth Streaming server and set up an on-demand stream. An easy to follow walkthrough and full deployment guide are available.
2. Clean the data stores on both your "branch" and "data center" Steelhead appliances.
3. Initiate a pre-population operation on the branch Steelhead appliance. This is done by navigating to the CLI and entering:
a. <i>conf t</i>
b. <i>protocol http prepop list <name> silverlight-url <manifest></i>
i. Where <<i>manifest</i>> is the URL for your Microsoft Silverlight manifest file. It will be of the form <i>http://server_name/Silverlight/SilverlightStream.isml/Manifest</i>
c. Repeat the previous step for every on-demand video to be pre-populated
d. <i>protocol http prepop list <name> start</i>
4. The CLI will be unresponsive as the files are pre-populated.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">The Pre-population service ensures that the data store of the Steelhead appliances are always "warmed" with the on-demand video. Because the pre-population is performed over an efficient, optimized connection, the same data sequences will never have to cross the WAN twice.

While we initiated the pre-population immediately in this example, a job can be set to initiate the pre-population at a time you specify. Please refer to the Steelhead Appliance Deployment Guide for details on creating a job.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Quality-of-Service (QoS) Enforcement

For cost reasons, it is common to have oversubscribed links in a production environment. For non-interactive traffic such as file transfer, a delay of several seconds wouldn't be significant. However, when the network carries delay sensitive traffic such as voice or video, delay of even one second can be detrimental to the quality of the interactive experience. Therefore, implementing QoS is crucial to prioritize delay sensitive data over any other traffic.

QoS is typically enforced on the router but it may not be possible to do so in certain situations (e.g. provider managed router or the router simply does not support QoS). In this case, implementing QoS on the DUT may be the next best option.

The Steelhead appliance can classify and enforce QoS ensuring the proper handling of delay sensitive traffic across the WAN when the link is oversubscribed. The implementation of the QoS in RiOS is based on Hierarchical Fair Service Curve (HSFC) which can guarantee both bandwidth and latency.

This section contains the following tests:
1. Outbound QoS (Basic)
2. Outbound QoS (Advanced)
3. Inbound QoS
4. With Inbound QoS, introduced in RiOS 7.0.1, QoS policies can be applied to traffic on the inbound direction. This applies to traffic coming for another Steelhead appliance or for general internet traffic. Inbound QoS is useful for split tunnel deployments, where there is direct internet access at the branch office. In these deployments general internet traffic, such as personal email, web browsing, etc. that are commonplace at the office are not coming from a peer Steelhead appliance, and as a result have not been subject to an appropriate QoS policy. Inbound QoS is needed in order to prevent this traffic from overwhelming traffic coming from the peer Steelhead appliance in the data center.

Use Case	This procedure tests Inbound QoS Enforcement functionality with the Steelhead. This test involves streaming a YouTube video while doing an FTP transfer of a large file. Prerequisites: 1. Install FTP server on a server running Windows Server 2003/2008/2008R2. All these components come standard with Windows Server 2003/2008/2008R2 but may either need to be installed or enabled. 2. Select a reasonable length YouTube video. Here is a good one. Stream the video at a fixed setting, ie 720p, 480p, etc. 3. Select a large file to download from a public FTP. Here is a good one. 4. Set the WAN simulator to the appropriate bandwidth and latency. Set the bandwidth to be slightly more than the bandwidth required by the video. A 480p video requires approximately 1.5 Mbps of bandwidth. **Clear the DUT's Data Store before the start of the test? No** **Clear the DUT's Data Store at the end of the test? No**
Procedure	1. START the FTP download 2. START streaming the video file 3. Notice the jitter with the video file while the FTP transfer is in progress. 4. STOP the FTP download and the video streaming. 5. Create one new QoS classes: *Class-video* with a *Latency Priority* of Real-Time. Ensure sufficient bandwidth is allocated to this class but also cap the maximum bandwidth so that it doesn't monopolize the link. 6. Create the following QoS rules on the "branch" DUT - Place YouTube traffic into *Class-video*. Use the *Application Protocol* drop down menu and select YouTube and set the Class Name to *Class-video*. Everything else can stay the same. - Best effort delivery for all other traffic. 7. Repeat steps 1) – 2) above. Note that if you previously allowed the streaming video to play to completion, you'll need to clear the Client's local browser cache before streaming it again. 8. There should be no jitter with the video file. 9. In the Steelhead administrative GUI, click on Reports → QoS Statistics.
Steelhead Appliance	When there is delay sensitive traffic in the "real-time" queue, the Steelhead will send that traffic out first ahead of the non-delay sensitive traffic. This ensures a smooth delivery of video file without any jitter even during an FTP transfer.

© 2012 Riverbed Technology. All rights reserved.

expected results		
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

5. Citrix ICA Optimization with Enhanced QoS

1. Outbound QoS (Basic)

<table>
<tr><td>Use Case</td><td colspan="2">This procedure tests outbound QoS Enforcement functionality with the Steelhead. This test involves streaming a video file while doing an FTP transfer and in addition an interactive telnet session.

Prerequisites:
1. Install Streaming Media Server, FTP server and a Telnet server on a server running Windows Server 2003/2008/2008R2. All these components come standard with Windows Server 2003/2008/2008R2 but may either need to be installed or enabled. In the Windows Media Services admin tool, be sure to “Allow New Connections” in the Publishing Point configuration to enable streaming.
2. Obtain a reasonable size WMV format movie file (preferably with a bit rate that’s just under the bandwidth of the WAN link) and the length of the movie should run for several minutes. Place this file in appropriate directory so that the client can stream this file.
3. Stage a file in the FTP directory so that the client can download this file.
4. Set the WAN simulator to the appropriate link speed and latency

Clear the DUT’s Data Store before the start of the test? Yes
Clear the DUT’s Data Store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="2">1. STOP the optimization service on the DUT
2. START the FTP download
3. START streaming the video file (typical URL format: rtsp://<server>/filename.wmv)
4. START a telnet session
5. Notice the jitter with the video file and also the unresponsiveness of the telnet session while the FTP transfer is in progress.
6. STOP the FTP download, streaming and the telnet session.
7. Create two new QoS classes:
Class-video with a priority of Real-Time. Ensure sufficient bandwidth is allocated to this class but also cap the maximum bandwidth so that it doesn’t monopolize the link.
Class-interactive with a priority of Interactive. Ensure that it’s given a 5% guaranteed bandwidth.
8. Create the following QoS rules on the both the “branch” and “data center” DUT
- Place RTSP (server side TCP port 554) into Class-video.
- Place Telnet traffic (server side TCP port 23) into Class-interactive.
- Best effort delivery for all other traffic.
9. START the service on the DUT
10. Repeat steps 2) – 4) above. Note that if you previously allowed the streaming video to play to completion, you’ll need to clear the Client’s local IE browser cache before streaming it again.
11. There should be no jitter with the video file and telnet should be responsive.
12. In the Steelhead administrative GUI, click on Reports → QoS Statistics.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">When there is delay sensitive traffic in the “real-time” queue, the Steelhead will send that traffic out first ahead of the non-delay sensitive traffic. This ensures a smooth delivery of video file without any jitter even during an FTP transfer. In addition, interactive pass-through traffic (i.e. telnet) will also have priority over FTP traffic. When there is extra bandwidth available, the extra bandwidth will be given to the FTP process.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

2. Outbound QoS (Advanced)

With Advanced QoS, introduced in RiOS 6.5, traffic can now be classified at an application level. The configuration is as simple as selecting the application from a drop down list and assigning it to QoS class. Both minimum and maximum bandwidths along with latency sensitivity can now be set at the application level.

Use Case	This procedure tests advanced outbound QoS Enforcement functionality with the Steelhead. This test involves streaming a video

© 2012 Riverbed Technology. All rights reserved.

<table>
<tr><td></td><td colspan="2">file while doing an FTP transfer and in addition an interactive telnet session.

Prerequisites:
1. Install Streaming Media Server, FTP server and a Telnet server on a server running Windows Server 2003/2008/2008R2. All these components come standard with Windows Server 2003/2008/2008R2 but may either need to be installed or enabled. In the Windows Media Services admin tool, be sure to “Allow New Connections” in the Publishing Point configuration to enable streaming.
2. Obtain a reasonable size WMV format movie file (preferably with a bit rate that’s just under the bandwidth of the WAN link) and the length of the movie should run for several minutes. Place this file in appropriate directory so that the client can stream this file.
3. Stage a file in the FTP directory so that the client can download this file.
4. Set the WAN simulator to the appropriate link speed and latency

Clear the DUT’s Data Store before the start of the test? Yes
Clear the DUT’s Data Store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="2">1. STOP the optimization service on the DUT
2. START the FTP download
3. START streaming the video file (typical URL format: rtsp://<server>/filename.wmv)
4. START a telnet session
5. Notice the jitter with the video file and also the unresponsiveness of the telnet session while the FTP transfer is in progress.
6. STOP the FTP download, streaming and the telnet session.
7. Create the following QoS rules on the both the “branch” and “data center” DUT
- Application Protocol RTSP into Class Default-Site$$Realtime</i>.
- Place Telnet traffic (server side TCP port 23) into Default-Site$$Interactive.
8. START the service on the DUT
9. Repeat steps 2) – 4) above. Note that if you previously allowed the streaming video to play to completion, you’ll need to clear the Client’s local IE browser cache before streaming it again.
10. There should be no jitter with the video file and telnet should be responsive.
11. In the Steelhead administrative GUI, click on Reports → QoS Statistics.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">When there is delay sensitive traffic in the “real-time” queue, the Steelhead will send that traffic out first ahead of the non-delay sensitive traffic. This ensures a smooth delivery of video file without any jitter even during an FTP transfer. In addition, interactive pass-through traffic (i.e. telnet) will also have priority over FTP traffic. When there is extra bandwidth available, the extra bandwidth will be given to the FTP process.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

3. Inbound QoS

With Inbound QoS, introduced in RiOS 7.0.1, QoS policies can be applied to traffic on the inbound direction. This applies to traffic coming for another Steelhead appliance or for general internet traffic. Inbound QoS is useful for split tunnel deployments, where there is direct internet access at the branch office. In these deployments general internet traffic, such as personal email, web browsing, etc. that are commonplace at the office are not coming from a peer Steelhead appliance, and as a result have not been subject to an appropriate QoS policy. Inbound QoS is needed in order to prevent this traffic from overwhelming traffic coming from the peer Steelhead appliance in the data center.

<table>
<tr><td>Use Case</td><td>This procedure tests Inbound QoS Enforcement functionality with the Steelhead. This test involves streaming a YouTube video while doing an FTP transfer of a large file.

Prerequisites:
5. Install FTP server on a server running Windows Server 2003/2008/2008R2. All these components come standard with Windows Server 2003/2008/2008R2 but may either need to be installed or enabled.
6. Select a reasonable length YouTube video. Here is a good one. Stream the video at a fixed setting, ie 720p, 480p, etc.
7. Select a large file to download from a public FTP. Here is a good one.</td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

	8. Set the WAN simulator to the appropriate bandwidth and latency. Set the bandwidth to be slightly more than the bandwidth required by the video. A 480p video requires approximately 1.5 Mbps of bandwidth. **Clear the DUT's Data Store before the start of the test? No** **Clear the DUT's Data Store at the end of the test? No**
Procedure	10. START the FTP download 11. START streaming the video file 12. Notice the jitter with the video file while the FTP transfer is in progress. 13. STOP the FTP download and the video streaming. 14. Create one new QoS classes: *Class-video* with a *Latency Priority* of Real-Time. Ensure sufficient bandwidth is allocated to this class but also cap the maximum bandwidth so that it doesn't monopolize the link. 15. Create the following QoS rules on the "branch" DUT - Place YouTube traffic into *Class-video.* Use the *Application Protocol* drop down menu and select YouTube and set the Class Name to *Class-video*. Everything else can stay the same. - Best effort delivery for all other traffic. 16. Repeat steps 1) – 2) above. Note that if you previously allowed the streaming video to play to completion, you'll need to clear the Client's local browser cache before streaming it again. 17. There should be no jitter with the video file. 18. In the Steelhead administrative GUI, click on Reports → QoS Statistics.
Steelhead Appliance expected results	When there is delay sensitive traffic in the "real-time" queue, the Steelhead will send that traffic out first ahead of the non-delay sensitive traffic. This ensures a smooth delivery of video file without any jitter even during an FTP transfer.

Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

4. Citrix ICA Optimization with Enhanced QoS

RiOS 6.0 adds enhanced QoS capabilities to accelerate Citrix ICA performance. This is done by ensuring one user does not starve other users when running bandwidth intensive operations.

<table>
<tr><td>Use Case</td><td colspan="2">This test involves streaming a video file with one user while transferring a file with another user.

Prerequisites:
1. Obtain a reasonable size WMV format movie file (preferably with a bit rate that's just under the bandwidth of the WAN link) and the length of the movie should run for several minutes. Place this file in appropriate directory so that the client can stream this file.
2. Stage a large file to download from the ICA session's local drive to the client's local drive mapped inside the ICA session
3. Set the WAN simulator to the appropriate link speed and latency

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="2">1. STOP the optimization service on the DUT.
2. Open up the desktops of two separate users.
3. START the WAN Bandwidth Monitor.
4. Start the large file transfer to the local drive mapped inside the ICA session of one of the users. Ensure that it is consuming the majority of the bandwidth.
5. With the second user begin streaming the video.
6. Notice the jitter with the video file and while the file transfer is in progress.
7. STOP the file transfer and the streaming video.
8. Enable Citrix ICA optimization.
9. Create one new QoS class:
Class-video with a priority of Real-Time. Ensure sufficient bandwidth is allocated to this class but also cap the maximum bandwidth so that it doesn't monopolize the link.
10. Create the following QoS rules with Citrix ICA as the Application Protocol. The QoS rules need to be created on the both the "branch" and "data center" DUT
- Place Class 1 into Class-video.
- Best effort delivery for all other traffic.
11. START the optimization service on the DUT
12. Repeat steps 2) – 4) above. Note that if you previously allowed the streaming video to play to completion, you'll need to clear the Client's local IE browser cache before streaming it again.
13. There should be no jitter with the video file.
14. In the Steelhead administrative GUI, click on Reports → QoS Statistics.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">When there is delay sensitive traffic in the "real-time" queue, the Steelhead will send that traffic out first ahead of the non-delay sensitive traffic. This ensures a smooth delivery of video file without any jitter even during a file transfer. When there is extra bandwidth available, the extra bandwidth will be given to the file transfer.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

5. Classification of PCoIP with Enhanced QoS

RiOS 8.0 adds enhanced QoS capabilities to classify PCoIP traffic. VMware View is a desktop virtualization solution that simplifies IT manageability and control while delivering the highest fidelity end-user experience across devices and networks. PCoIP is the proprietary protocol for VMware View, a remote workstation and desktop solution, developed by Teradici. PCoIP enables remote access to workstations and servers, normally hosted in a machine room or data center, from a thin client. The PCoIP protocol transmits only the changing pixels across a standard IP network to stateless PCoIP technology enabled devices. PCoIP both compresses and encrypts data over the wire.

Note: The Steelhead appliance can only classify PCoIP traffic at this time since it runs over UDP.

Please consult the Teradici website at http://www.teradici.com/pcoip-technology.php for more detailed information on the PCoIP protocol.

<table>
<tr><td>Use Case</td><td colspan="4">This test involves streaming a video file

Prerequisites:

1. Obtain a reasonable size WMV format movie file (preferably with a bit rate that's just under the bandwidth of the WAN link) and the length of the movie should run for several minutes. Place this file in appropriate directory so that the client can stream this file.
2. Set the WAN simulator to the appropriate link speed and latency

Clear the DUT's data store before the start of the test? Yes
Clear the DUT's data store at the end of the test? Yes</td></tr>
<tr><td>Procedure</td><td colspan="4">1. On the server-side Steelhead, navigate to the Outbound QoS (Advanced) settings page.
2. Enable QoS Shaping and Enforcement.
3. Create 4 new QoS classes and a new QoS rule for PCoIP traffic.
4. Navigate to the Outbound QoS Reports page on the server-side Steelhead.
5. Map the 4 different QoS classes using different colors (e.g. Realtime = blue, Interactive = brown, Business-Critical = green, Best-Effort = red).
6. Open up the virtual desktop of a user.
7. Begin streaming the video.
8. As the video plays, note the different types of PCoIP traffic being classified.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="4">The Steelhead classifies the PCoIP session. Note the different traffic types (e.g. realtime, interactive, business-critical, best effort).</td></tr>
<tr><td rowspan="5">Results</td><td colspan="2"></td><td colspan="2">Yes / No</td></tr>
<tr><td>Baseline</td><td>PCoIP classified?</td><td colspan="2"></td></tr>
<tr><td colspan="2" rowspan="2"></td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Yes / No</td><td>Yes / No</td></tr>
<tr><td>Cold</td><td>PCoIP classified?</td><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Disconnected Operations

While most DUT rely on the WAN to provide accelerated access to centralized data, you may have a desire to allow for access to data even when the WAN is down. Because of coherency and consistency concerns, it is not possible to guarantee global read/write access to collaborative data through a WAN outage (see the Riverbed white paper *No Free Lunch* for details on this), but a DUT should be able to at least provide limited, safe access to Windows file data during WAN outages.

This type of disconnected operation is provided on the Steelhead appliance using the Proxy File Service feature. See the *Steelhead Appliance Deployment Guide* for detailed information on configuration of PFS. For other DUT, consult the appropriate vendor for information on configuring disconnected operations.

Some basic steps to test disconnected operations with a Steelhead appliance:

Use Case	Testing the ability of the DUT to alleviate WAN outages
Procedure	1. Configure your "data center" server as a Windows Domain Controller 2. Create two folders on the data center server and enable them for sharing (e.g. "\\server\broadcast_source" and "\\server\local_backup"). Populate the "broadcast_source" folder with several documents. 3. Ensure that your "branch" Steelhead appliance has adequate network connectivity and IP/DNS configuration such that it can "ping" the data center server by hostname. Also ensure that your Steelhead appliances are time synchronized with the "data center" server (such as through configuring NTP). 4. Complete the PFS configuration on your branch Steelhead appliance to join the system to the Windows domain 5. Start the PFS service on the branch Steelhead appliance 6. Add a new proxy share called "broadcast" on the branch Steelhead appliance, specifying mode "Broadcast". Associate the share with the "broadcast_source" share on the data center server. Set the "Incremental Sync Interval" to one minute. Save the configuration. 7. Add a new proxy share called "local" on the branch Steelhead appliance, specifying mode "Local". Associate the share with the "local_backup" share on the data center server. Set the "Incremental Sync Interval" to one minute. Save the configuration. 8. Configure each share for "Syncing Enable" In the Shares screen in the Steelhead appliance administrative GUI, note that the status changes to reflect that synchronization is taking place. Confirm through the Current Connections report in the administrative GUI that you see an optimized CIFS connection from your "branch" Steelhead appliance's primary IP address to your data center server's IP address. 9. Once the Shares screen in the administrative GUI reflects that both shares have a status of "Share idle", the initial synchronization has completed. Configure each share for "Sharing Enable". 10. On the branch client system, map a drive to the "broadcast" share. Confirm you see the documents that you had staged. 11. On the branch client system, map a drive to the "local" share. Copy several documents to the location. 12. On the "Shares" screen in the Steelhead appliance administrative GUI, confirm that the status for each share updates within a minute to indicate they are being synchronized. Most of the data flowing will be the documents you added to the "local" share being backed up to the "local_backup" folder in the data center. 13. Once the statuses of the shares indicate that synchronization has completed, create a network outage by disconnecting the WAN simulator. 14. From the branch client, access the shares again via the map drives. Confirm that you can still read the documents in both shares. Copy additional documents to the local mode share to confirm you can still write to it. Confirm that you cannot write to the broadcast mode share. 15. From the data center side, populate the "broadcast_source" folder with several more documents. 16. Reconnect the WAN simulator. 17. On the Shares screen in the Steelhead appliance administrative GUI, confirm that the shares synchronize again within a minute. 18. After synchronization is complete, confirm from the branch client that the additional documents added to the broadcast mode share are now visible. Also confirm from the data center server that the additional documents added to the local mode share are now visible in the "local_backup" folder. 19. To ensure there's no synchronization traffic on the WAN during future tests, you may wish to remove the PFS configuration.

© 2012 Riverbed Technology. All rights reserved.

<table>
<tr><td rowspan="1">Steelhead Appliance expected results</td><td colspan="2">The broadcast mode share provides read-only access to centralized data even when the WAN is down. Any changes made to the central data are synchronized to the remote Steelhead appliance when the WAN returns.

The local mode share provides read/write access to local data even when the WAN is down. Any changes made to the local data are synchronized to the central backup location when the WAN returns.

All synchronization traffic is subject to Steelhead appliance optimization through an efficient TCP stream.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Pass/Fail</td><td>Pass/Fail</td></tr>
<tr><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Servicing Multiple Links

For purposes of redundancy, many networks contain multiple switch/router paths. Such a topology could be created in your test network as shown in the following diagram:

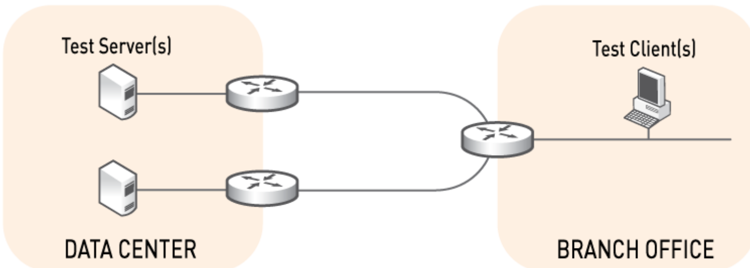

Using routing, create traffic flows through the test network such that both LAN/router paths are utilized. Add your DUT to the network such one DUT straddles both network paths. The Steelhead appliance offers a four-port card that can be used to permit this. See the *Bypass Installation Guide* for detailed information on installing a four-port card. See the *Steelhead Appliance Deployment Guide* for detailed information on configuring multiple in-path interfaces.

The final topology is shown in the following diagram:

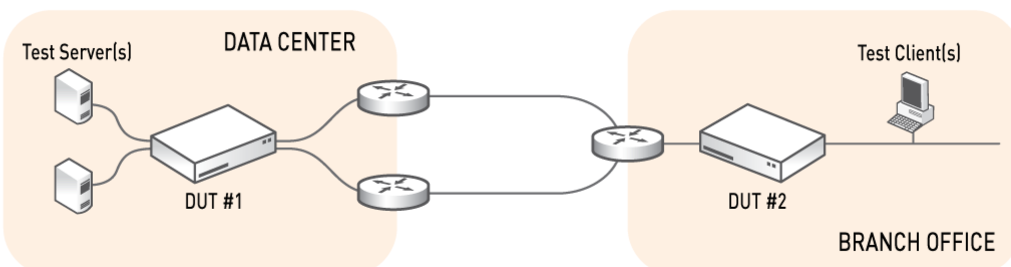

Perform operations through both network paths and use the DUT administrative interface to confirm that connections are optimized.

Steelhead Appliance expected results	Traffic is optimized through both network paths.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Asymmetric Routing Auto-Detection

To provide maximal benefits, RiOS requires that traffic associated with optimized TCP sessions pass through Steelhead appliances in both directions. Complex networks often contain asymmetric routing (a common by-product of multi-path redundant networks) which may result in traffic flows that bypass a Steelhead appliance in one direction. In an ideal situation, the network administrator should be aware of this and account for it by ensuring symmetric routing or configuring connection forwarding on the Steelhead appliances (covered in the next section). However, given a very complex network, it simply may be not possible determine the exact traffic flow for all traffic. The Steelhead appliance has the ability to detect asymmetric routing in a network and to pass through the traffic, alerting the administrator so the issue may be addressed.

To test asymmetric routing auto-detection, use routing to set up an asymmetric traffic flow in your network. This could be done on either the client-side or server-side of the network, or you could test each separately. Begin by configuring your network with a traffic flow such as what's shown in one of the following diagrams:

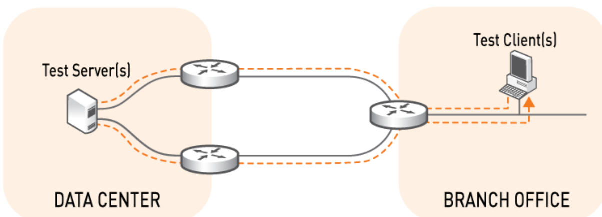

Server-Side Asymmetric Routing

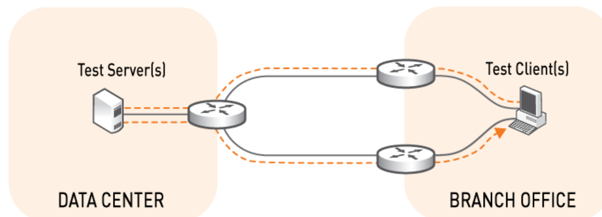

Client-Side Asymmetric Routing

Add your DUT to the network as shown in the following diagrams:

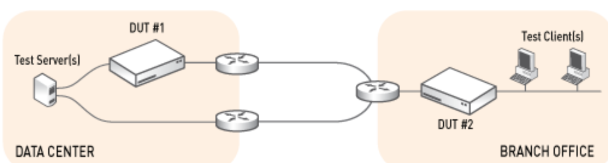

Server-Side Asymmetric Routing (with DUT)

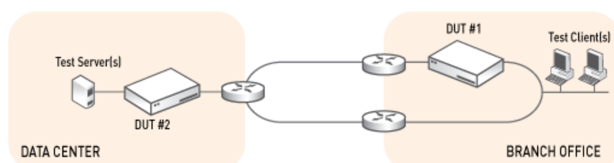

Client-Side Asymmetric Routing (with DUT)

Use Case	Testing asymmetric routing detection capabilities of the DUT	
Procedure	1. Open a connection between client and server that follows an asymmetric path. 2. Check logs of the Steelhead and look for the following entry: **asymmetric routing between <*source-IP*> and <*dest-IP*> detected** 3. Note the status of the Steelhead appliance is now in "Degraded" mode 4. The connection will be passed through	
Steelhead Appliance expected results	When the Steelhead appliance detects an asymmetric routed connection, it will start passing through the traffic for a set amount of time. This ensures that any asymmetric routed connections will continue to work albeit without optimization.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Connection Forwarding

In environments that are known to contain asymmetric traffic flows, configuration of the connection forwarding feature in the Steelhead appliances can ensure that all traffic is subject to optimization. See the *Steelhead Appliance Deployment Guide* for details on configuring connection forwarding.

To test connection forwarding, start with one of the asymmetric routed configurations from the previous exercise, this time adding a 3rd DUT as shown in one of the following diagrams:

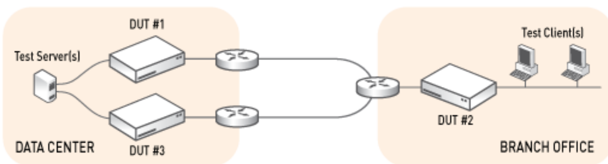

Server-Side Asymmetric Routing (with DUT)

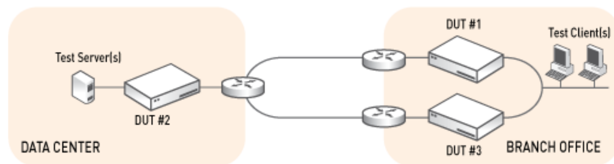

Client-Side Asymmetric Routing (with DUT)

Use Case	Testing the ability of the DUT to optimize in the face of asymmetric routes	
Procedure	1. Enable connection forwarding on DUT #1 and DUT #3. 2. Configure DUT #1 and DUT #3 as Peers by configuring each as a neighbor of the other. 3. Open a connection between client and server that follows an asymmetric path. Via the Steelhead appliance administrative interface, confirm that the connection is optimized. 4. Using tcpdump, study the source/destination IP addresses of packets for the connection as they enter/leave each Steelhead appliance LAN/WAN interface. Note the point where the destination address is changed to forward traffic between the two Steelhead appliances. 5. In the Steelhead appliance administrative GUI, observe the “Current Connections” report and observe that your connection is flagged as “Forwarded”	
Steelhead Appliance expected results	The Steelhead appliance that handles the initial TCP SYN forwards information about the connection to the peer Steelhead appliance before the connection is completed to the server. As traffic returns from the server via the asymmetric path, the peer Steelhead appliance forwards packets associated with the connection back to the other Steelhead appliance. This allows proxy of the TCP session so that full optimization can be performed for the connection.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Filling Large Data Pipes

While most DUT are designed for conditions where latency is high and bandwidth is constrained, sometimes bandwidth is not a concern. In many environments, links as large as OC-12 (622 mbps) are available, but due to high network latency, protocols used for tasks such as bulk file transfer may not be able to adequately use the available bandwidth. A DUT should be capable of performing latency optimization at these high speeds to take advantage of the available bandwidth. However, it should do so in a manner that shares the pipe with other, lower-speed or un-optimized flows.

To test this type of high-speed traffic, you first need to acquire a high-speed WAN simulator. Systems that should be capable of operating at these speeds include the products from Shunra and Apposite or the gigEnn-OC12. You will also need high-end client and server systems in order to send traffic through the network at this rate. One way to do this would be to install several GB of RAM in both the client and server systems and create RAM disks that can be used to host/receive data, such as through FTP and HTTP.

The Steelhead appliance addresses the filling of large pipes through a feature called high-speed TCP (HSTCP). In order to scale to speeds as high as OC-12, it is possible to disable first-pass compression and/or SDR disk-based data reduction and instead concentrate Steelhead appliance resources on removing the effects of latency in order to fill the pipe with data. See the *Steelhead Management Console User's Guide* for detailed information on configuring HSTCP.

Use Case	Testing the ability of the DUT to fill large data pipes using HSTCP	
Procedure	1. Create a simple network topology such as was used for the basic performance tests. 2. Configure a high-bandwidth, high-latency link on the WAN simulator. 3. Configure interface queues in the WAN simulator such that their size corresponds to the bandwidth-delay product of the simulated link. 4. Perform a bulk transfer between the client and server using FTP. Note the low performance due to the latency of the link. 5. Enable HSTCP on the Steelhead appliances. 6. Configure the LAN/WAN Send/Receive buffer size on the Steelhead appliances to correspond to double the bandwidth-delay product of the simulated link. 7. If you are testing with a very high-speed link (e.g. OC-12), add an in-path rule to both Steelhead appliances such that all traffic is subject to optimization policy “None”. This disables first-pass compression and SDR disk-based data reduction. 8. Perform the bulk FTP transfer again between the client and server. Note that the effective throughput is much higher and the WAN pipe is now kept full. 9. Add an additional in-path rule to both Steelhead appliances such that all HTTP traffic (port 80) is “Pass-Through”. 10. Perform the bulk FTP transfer again. During the FTP transfer, run an HTTP transfer. Note that the HSTCP-enabled FTP transfer “backs off” somewhat to make room for the un-optimized HTTP transfer. 11. After the HTTP transfer is complete, note that the FTP transfer fills the pipe once again.	
Steelhead Appliance expected results	Steelhead appliance HSTCP allows for TCP across the WAN to “ramp up” more quickly and “back off” less in the face of congestion or loss. This causes the FTP transfer to quickly fill the pipe and keep it full. However, HSTCP is fundamentally TCP, therefore when you introduce the HTTP “cross-traffic”, the high-speed FTP backs off enough to share the pipe with the cross-traffic. Once the cross-traffic is gone, the high-speed FTP once again “ramps up” to fill the pipe.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Netflow

Netflow is a valuable tool for network administrators as it provides very useful data for capacity planning, billing and may even assist in troubleshooting. Netflow is often enabled on routers, but may also be available on a DUT to provide statistics on connections traversing the system. These statistics are likely to include client and server source/destination IP addresses and TCP port numbers, as well as byte count statistics. See the *Steelhead Management Console User's Guide* for detailed information on configuring Netflow on the Steelhead appliance.

In order to perform this test, you will need to have a Netflow collector running in your network to collect the statistics from the DUT. If you do not have one already setup, there are free Netflow collectors available on the Internet such as the Plixer Scrutinizer Netflow Analyzer.

In RiOS, by default, Netflow traffic is collected on inbound traffic only. The test below will consist of a FTP client uploading a file to a FTP server.

Use Case	Testing the Netflow capabilities of the DUT	
Procedure	1. Ensure that the Netflow collector is setup correctly and listening for incoming Netflow connections 2. Ensure that the FTP server is setup and that it allows uploading of files 3. Login to the Steelhead and click on Setup and then Netflow 4. Enable Netflow Export and then add the information for the Netflow collector. For the “Capture Interface”, choose LAN and add the collector to the list 5. Initiate the FTP upload. Wait until the upload is done, then go back and check the Netflow collector.	
Steelhead Appliance expected results	With the Netflow feature, the Steelhead can export all the information that the network administrator needs. The Steelhead is capable of exporting traffic regardless of whether it is optimized or pass-through.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Virtualization Services

<table>
<tr><td>Use Case</td><td colspan="2">This procedure tests the capability to run additional services and applications on the DUT. The Steelhead Riverbed Services Platform (RSP) uses VMware Server 2.0 to provide dedicated resource instances for virtual machines to run on the Steelhead appliance. . It provides the capability to run up to five additional services and applications virtually on VMware's VMserver 2.0 in a protected partition on the Steelhead appliance. Having RSP on the Steelhead minimizes the hardware infrastructure footprint at the branch office enabling companies to further consolidate IT.

Prerequisites (Steelhead):
1. Install the appropriate RSP package
2. Install the RSP license

Please refer to the Riverbed Services Platform Installation and Configuration Guide for details on installing and configuring RSP.</td></tr>
<tr><td>Procedure</td><td colspan="2">1. Obtain the Squid pre-built RSP package. Squid is open source Web cache and Proxy.
2. Follow the documentation to install and enable the Squid RSP package.</td></tr>
<tr><td>Steelhead Appliance expected results</td><td colspan="2">After installing and enabling the Squid RSP package the Steelhead will be able to cache and proxy general internet traffic including YouTube videos.</td></tr>
<tr><td rowspan="3">Results</td><td>Steelhead Appliance</td><td>Alternative DUT</td></tr>
<tr><td>Virtualization Services</td><td>Virtualization Services</td></tr>
<tr><td></td><td></td></tr>
</table>

© 2012 Riverbed Technology. All rights reserved.

Packet Loss Mitigation

The TCP-based approach to Data Streamlining present in the Steelhead appliance addresses network congestion, which is the most common root cause of packet loss in production environments. As a result, for most environments, no special configuration beyond the optimizations already tested in this document should be required to address this common root cause of packet loss. However, some environments may include network impairments such as “dirty circuits” or under-buffered router interfaces. These impairments may require specific loss-mitigation features in order to maintain consistent WAN utilization.

The MX-TCP feature in RiOS provides the network administrator the ability to configure a target WAN utilization that should be maintained on a dedicated circuit for optimized traffic. Once configured, MX-TCP will maintain that level of WAN utilization despite any lost packets that occur.

Because MX-TCP is configured as part of the RiOS QoS Classification system, granular controls are available to apply MX-TCP to only particular subsets of traffic. For simplicity, the test steps described here will configure MX-TCP to affect the throughput of all optimized traffic between Steelhead appliances.

Use Case	The test involves creating a lossy network environment, confirming that operations are negatively affected in the environment, and then introducing MX-TCP on the Steelhead appliances to mitigate the effects of the packet loss. To ensure you experience the effects of the MX-TCP feature, you should select one of the longer Basic Tests from earlier in the test plan to re-run here. A recommended test would be the HTTP 7 MB file download test with 1536k bandwidth and 100 ms round-trip latency.	
Procedure	1. Review the prior results for the Un-optimized and Cold runs for your selected Basic Test (runs earlier without packet loss) 2. Configure a packet loss rate on your WAN simulator (e.g. 1% in each direction). 3. STOP the optimization service on the DUT 4. START your timer 5. Repeat the run of your selected Basic Test 6. Once the operation has finished, STOP your timer. 7. Record your result 8. START the optimization service on the DUT and clean the data store to ensure a Cold run 9. START your timer 10. Repeat the run of your selected Basic Test 11. Once the operation has finished, STOP your timer. 12. Record your result 13. Configure MX-TCP in the Steelhead administrative GUI on both the “branch” and “data center” DUTs: - Enable QoS Classification with a WAN Bandwidth equal to the bandwidth configured in your WAN simulator - Create a new QoS class called *Blast* with Normal Priority, a Guaranteed Bandwidth of 99%, and Queue “mxtcp”. - Create a QoS rule to place Optimized traffic for all Source/Dest Subnets and Ports into the newly created *Blast* class. - Save the configuration 14. RESTART the optimization service on both DUTs and clean the data stores to ensure a Cold run 15. Repeat the run of your selected Basic Test 16. Once the operation has finished, STOP your timer. 17. Record your result	
Steelhead Appliance expected results	On long-lived flows such as HTTP file download, packet loss may cause poor performance. Each lost packet is interpreted by TCP as congestion such that TCP “backs off” to make room for the perceived congestion. In higher latency environments (such as the 100 ms latency recommended for this test) it may take extra time for TCP to recover. In the presence of loss, while the Cold run still significantly outperforms the Un-optimized run, performance in the Cold run can be improved even further if the sustained loss did not cause TCP to back off at all. The MX-TCP feature allows the Steelheads to continue sending out traffic at the configured rate rather than backing off. Any packets that were lost are quickly retransmitted by MX-TCP, allowing the Cold run in the lossy environment to perform almost as well as the Cold run in the loss-free environment. The Warm run was not repeated here because it may complete too quickly for any loss events to occur during the duration of the test.	
Results		**Seconds**

© 2012 Riverbed Technology. All rights reserved.

			Steelhead Appliance	Alternative DUT
	Without packet loss	Un-optimized HTTP download (copied from earlier run)		
		Cold HTTP download (copied from earlier run)		
	With packet loss	Un-optimized HTTP download (#7)		
		Cold HTTP download – w/o MX-TCP (#12)		
		Cold HTTP download – w/ MX-TCP (#17)		

© 2012 Riverbed Technology. All rights reserved.

Adaptive MX-TCP

Packet Loss Mitigation was addressed with the MX-TCP feature in the previous section. New in RiOS 8.0, the Adaptive MX-TCP feature is designed to take advantage of Bandwidth beyond the guaranteed limit. With Adaptive MX-TCP, Minimum Bandwidth (Guaranteed Limit) and Maximum Bandwidth (Upper Limit) can now be specified in QoS classes.

The benefits for Adaptive MX-TCP are that available bandwidth beyond the Minimum Bandwidth (Guaranteed Limit) can now be used up to the Maximum Bandwidth (Upper Limit). In Disaster Recovery (DR) scenarios during off-peak hours where there is a lot of bandwidth available but can go unused; Adaptive MX-TCP will certainly help.

Use Case	The test involves creating a lossy network environment, confirming that operations are negatively affected in the environment, and then introducing MX-TCP on the Steelhead appliances to mitigate the effects of the packet loss. To ensure you experience the effects of the Adaptive MX-TCP feature, you should select one of the longer Basic Tests from earlier in the test plan to re-run here. A recommended test would be the HTTP 7 MB file download test with 1536k bandwidth and 100 ms round-trip latency. <u>Configure the Minimum Bandwidth at 30% and the Maximum Bandwidth at 100%.</u>
Procedure	1. Configure a packet loss rate on your WAN simulator (e.g. 1% in each direction). 2. Configure MX-TCP in the Steelhead administrative GUI on the "data center" DUT: - Enable QoS Classification with a WAN Bandwidth equal to the bandwidth configured in your WAN simulator - Create a new QoS class called *Blast* with Normal Priority, a Minimum Bandwidth of 30%, a Maximum Bandwidth of 100% and Queue "mxtcp". - Create a QoS rule to place Optimized traffic for all Source/Dest Subnets and Ports into the newly created *Blast* class. - Save the configuration 3. RESTART the optimization service on the DUT and clean the data store to ensure a Cold run 4. Navigate to the Outbound QoS Reports page on the DUT. 5. Map the "mxtcp" traffic to any particular color (e.g. brown). 6. Begin the HTTP file download of your selected Basic Test. 7. Note as the file download saturates the WAN, the "mxtcp" traffic class peaking steadily at Maximum Bandwidth (100%).
Steelhead Appliance expected results	The Adaptive MX-TCP feature allows the Steelheads to continue sending out traffic at the **<u>Maximum Bandwidth (Upper Limit)</u>** rate rather than backing off. Any packets that were lost are quickly retransmitted by MX-TCP, allowing the Cold run in the lossy environment to perform almost as well as the Cold run in the loss-free environment.

Results			Yes / No	
	Baseline	MX-TCP at Maximum Bandwidth (Upper Limit)?		
			Steelhead Appliance	**Alternative DUT**
			Yes / No	Yes / No
	Cold	MX-TCP at Maximum Bandwidth (Upper Limit)?		

© 2012 Riverbed Technology. All rights reserved.

Embedded Cascade Shark

New in RiOS 7.0 is the ability to integrate the Riverbed Steelhead appliance with Cascade Pilot. With this new feature packets captured on the Steelhead appliance, using the built-in tcpdump, will immediately be accessible within Cascade Pilot, turning the Steelhead appliance into a Mini Cascade Shark.

Use Case	This procedure tests ability of the DUT to integrate with Cascade Pilot.	
Procedure	1. In the Steelhead Management Console: a. Navigate to *Reports* | *Diagnostics* | *TCP Dumps* and *Enable Cascade Shark* b. Navigate to *Configure* | *Security* | *User Permissions* and create a password for user *shark* under *Role-Based Accounts* 2. In Cascade Pilot: a. Navigate to the *Remote* tab and select *Add Probe* b. Enter the IP address of the Primary or Auxiliary interfaces of the Steelhead appliance c. Enter *shark* for the *User Name* d. Enter the password created in the previous step 3. In the Steelhead Management Console initiate a tcpdump	
Steelhead Appliance expected results	The tcpdump capture should now be accessible from within Cascade Pilot.	
Results	**Steelhead Appliance**	**Alternative DUT**
	Pass/Fail	Pass/Fail

© 2012 Riverbed Technology. All rights reserved.

Diagnostics & Troubleshooting

Use Case	This procedure tests the troubleshooting and reporting features available on the appliance			
Procedure	1. Log into the DUT 2. Check the overall health and optimization status. 3. List the 50 "top-talkers' in the past hour 4. Perform a tcpdump capture of traffic going to a server			
Steelhead Appliance expected results	The Steelhead appliance has a comprehensive range of reports available, including showing WAN bandwidth usage for conversations, senders, receivers and applications for top Talkers in 5 minute intervals. The ability to perform a tcpdump capture of traffic from the GUI is very useful for troubleshooting purposes.			
Results	**Steelhead**		**Alternative DUT**	
	Diagnostics/Troubleshooting capability	Pass/Fail	Diagnostics/Troubleshooting capability	Pass/Fail

Conclusion

Riverbed Steelhead appliances can accelerate all applications that run over TCP, while at the same time streamlining deployment and ongoing management of branch office IT infrastructure. Steelhead appliances are designed to have the flexibility and the intelligence to operate as organizations would like to operate, and not force IT staff or users to make dramatic changes to their normal business processes. Because of this unique approach, Steelhead appliances have become the technology of choice for enterprise application acceleration.

Contacting Riverbed Technical Support

For technical support, please contact your Sales Engineer or Riverbed Technical Support at 1-888-RVBD-TAC (1-888-782-3822) in the United States and Canada or +1 (415) 247-7381 outside the United States.

About Riverbed

Riverbed delivers performance for the globally connected enterprise. With Riverbed, enterprises can successfully and intelligently implement strategic initiatives such as virtualization, consolidation, cloud computing, and disaster recovery without fear of compromising performance. By giving enterprises the platform they need to understand, optimize and consolidate their IT, Riverbed helps enterprises to build a fast, fluid and dynamic IT architecture that aligns with the business needs of the organization. Additional information about Riverbed (NASDAQ: RVBD) is available at http://www.riverbed.com/.

riverbed

Riverbed Technology, Inc.
199 Fremont Street
San Francisco, CA 94105
Tel: (415) 247-8800
www.riverbed.com

Riverbed Technology Ltd.
One Thames Valley
Wokingham Road, Level 2
Bracknell. RG42 1NG
United Kingdom
Tel: +44 1344 31 7100

Riverbed Technology Pte. Ltd.
391A Orchard Road #22-06/10
Ngee Ann City Tower A
Singapore 238873
Tel: +65 6508-7400

Riverbed Technology K.K.
Shiba-Koen Plaza Building 9F
3-6-9, Shiba, Minato-ku
Tokyo, Japan 105-0014
Tel: +81 3 5419 1990

© 2012 Riverbed Technology. All rights reserved.