A Riverbed Technology White Paper

Scaling Riverbed Deployments

riverbed™

SCALING RIVERBED DEPLOYMENTS

Introduction

Since their introduction in May of 2004, thousands of Riverbed's Steelhead appliances have been deployed in production IP networks of every size, from one-person home offices to data center environments that rank among the largest in the world. The low-cost Steelhead 100 and 200 appliance models are ideal for deployments in remote offices with fewer than 15 persons, while larger Steelhead 6020 model are appropriate for large data centers used to support thousands of remote users.

This document explores various issues and methodologies used by Riverbed customers to scale Steelhead deployments for large IP networks. Unlike competitive products, the Riverbed solution has been successfully deployed and scaled in networks having hundreds of remote sites and tens of thousands of remote users. Furthermore, there is strong interest and planning by many customers to extend their deployments to their entire enterprise network infrastructure, which some cases involve thousands of remote sites.

Issues to Consider when Scaling Riverbed Deployments

When scaling deployment of large numbers of Riverbed Steelhead appliances, there are various issues that must be considered. The most appropriate deployment approach involves weighing the importance of each of these issues for the specific network environment, as each deployment methodology involves tradeoffs. The key issues are:

- Appliance Sizing
- Non-disruptive deployment
- Incremental scaling
- Device Configuration
- Asymmetric routing
- High availability
- End-to-end transparency
- Deployment timeline and process
- Centralized management

Each is described further below.

Appliance Sizing

The appropriate Steelhead appliance model should be used at each remote and data center site. Identifying the appropriate model involves evaluating several site-specific factors, including the number of users, the WAN bandwidth available, and the amount of data being accessed over the WAN at that site. In a majority of cases, the most important of these parameters is the number of users at each remote site. Riverbed's experience indicates that most customers should allocate an average of five TCP connections per end-user. For a given remote site, the appropriate Steelhead model will support at least the total number of TCP connections generated by remote users at that site; in a data center site, the one or more deployed Steelhead appliances should be able support the total number of TCP connections generated by remote users at all sites accessing data from the data center servers.

The largest Steelhead appliance, the 6020, is able to support 40,000 TCP connections. Using the above guideline, a single Steelhead 6020 appliance is able to support as many as 8000 remote users. Through clustering approaches described in this document, the Steelhead 6020 can be used to scale the Riverbed solution to support hundreds of thousands of remote users.

Non-Disruptive Deployment

For short-term deployments such as for a product evaluation or pilot deployment, the option for a non-disruptive deployment can be desirable. In these situations, a deployment that requires minimal re-cabling and network down-time would be the most favorable option, particularly if the data center network supports mission-critical traffic for tens or hundreds of thousands of users. Riverbed offers deployment options that require no changes to pre-existing cabling in the data center, which minimizes risk to the production traffic in the network.

Incremental Scaling

As new remote sites are supported by the WAN optimization solution, additional Steelheads may be needed in the data center to support new incremental traffic from the additional remote sites. In these cases, an ability to easily add WAN optimization, without disrupting applications and remote users that are already receiving optimized performance, can be desirable.

Fortunately, many of the deployment options discussed later in this document are easy to scale incrementally, without disrupting traffic for existing applications and sites that previously were receiving acceleration.

Device Configuration

When deploying large numbers of devices to hundreds or thousands of remote sites, individual device-specific configuration can be a huge problem. Fortunately, when deployed in an in-path configuration, where the Steelhead appliance is physically in the network path between the LAN and the WAN, Riverbed offers autodetection and discovery capabilities that allow each Steelhead to automatically optimize traffic without identification of host subnets and peer Steelhead appliances. As a result, Riverbed appliances can automatically discover each other and support full-mesh connectivity with no additional configuration, regardless of the number of sites in your network.

"Riverbed has been successfully deployed and scaled in IP networks having hundreds of remote locations and tens of thousands of remote users."

Beware of tunnel-based solutions—they struggle with this issue, because a tunneling device requires explicit identification of optimized host subnets and peer devices (and there may be thousands of subnets in a medium-size enterprise network). The configuration problem for tunneling-based products becomes worse when establishing full-mesh connectivity among a large number of remote sites, in which case the amount of configuration grows faster than the number of deployed tunneling devices. Similarly, file caching-based products require explicit identification of peer caching appliances, and also require either configuration changes on workstations and/or WCCP configuration on each site router or switch. This becomes a huge problem when scaling for large networks where the number of sites totals in the hundreds or thousands.

Asymmetric Routing

Large data center networks have multiple Ethernet connections for redundancy and load balancing of traffic. As a result, the presence of asymmetric routing is almost assured, as some server response traffic will inevitably be routed along a different physical path than that used by the original client response. This is a consideration because Riverbed's Steelheads operate at the TCP layer, and each Steelhead must be able to observe and process traffic in both directions for a given TCP connection.

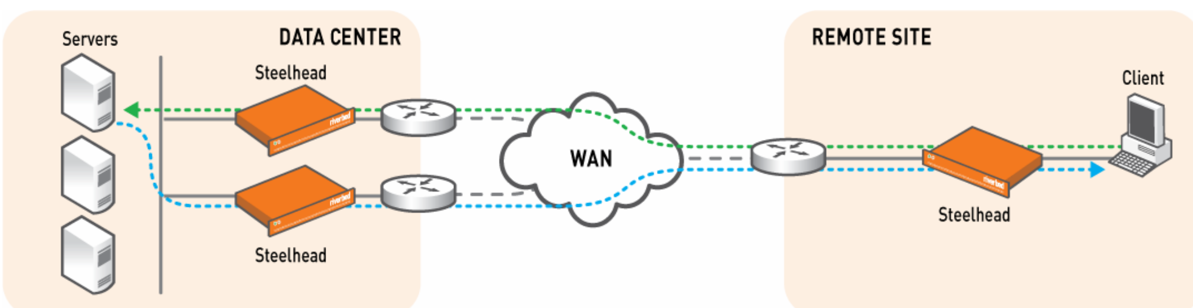

Figure 1*: Asymmetric routing common in large networks causes the server's response to take a different physical path from the original client's request.*

Fortunately, each of the deployment methodologies described later in this document have mechanisms that address asymmetric routing. Some deployment methodologies utilize Riverbed's patent-pending capability known as TCP connection forwarding, which allows Riverbed appliances to forward TCP flows to the "correct" Steelhead for optimization processing.

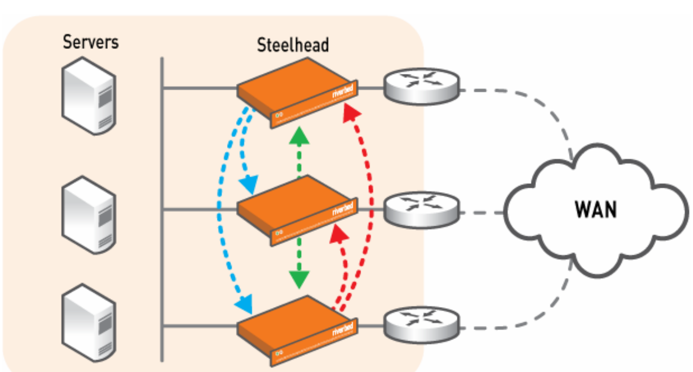

Figure 2: *The Connection Forwarding Solution: each server-side Steelhead appliance exchanges connection information with all the other server-side Steelhead appliances.*

High Availability

When supporting large numbers of remote users, it is important to ensure they continue receiving accelerated application performance in the event that a Steelhead appliance fails. Transparently passing-through traffic without optimization is often not an adequate solution.

To address this area, Riverbed offers a variety of high availability solutions. In each deployment methodology described below, multiple Steelhead appliances in the data center can be deployed not only for load balancing and scaling, but also for high availability through redundant Steelhead appliances. Some approaches deliver N+1 redundancy, while others provide 1:1 redundancy.

End-to-End Transparency

Some applications and network security environments require end-to-end transparency between the client and server. Specifically, select applications cannot communicate through proxies and NAT—for these applications the client must use the server's real IP address, and the server must respond to the real client's IP address in order for the application to function properly. This is the case with applications such as passive FTP, where the server will initiate a new data connection to the client.

Furthermore, in many cases security-related infrastructure may require that traffic be transparent end-to-end. For example with Intrusion Prevention Systems (IPS) that requires the traffic to be transparent to the real client and server.

Riverbed's in-path deployment options discussed below deliver end-to-end transparency for the client and server. Most Riverbed deployments use this approach, which delivers optimized performance for the broadest collection of TCP/IP-based applications.

> *"Since their introduction in May of 2004, thousands of Riverbed Steelhead appliances have been deployed in production IP networks of every size and scale..."*

Deployment Timeline and Process

Large Riverbed deployments involving hundreds or thousands of new Steelhead appliances are not enabled overnight. Rather, the recommended approach for scaling out a deployment is to implement it in phases. A flexible schedule will minimize risk and provide the means to adjust to unexpected contingencies that inevitably occur in large and complex deployments.

Typically, the first phase of a deployment may involve remote users at a small number of representative remote sites for a limited number of key applications. After validating the initial installation, phase two begins by adding more applications to the Steelhead optimization rules. The process is repeated until all TCP/IP applications are eventually optimized and accelerated through the Riverbed Steelhead appliances. At this point, the next deployment phase can begin, which involves a wider rollout in a larger number of remote sites. Lessons learned from the first two phases on configuration and tuning required for optimized applications are applied. The final deployment phase is the enterprise-wide rollout, which continues until all remote users at each remote location have every application optimized and accelerated.

Centralized Management

When large numbers of Riverbed appliances are deployed, monitoring and management of the many disparate devices can become an issue. An administrator without a central management facility will struggle to monitor, manage, and update each and every Steelhead appliance. The Riverbed Central Management Console (CMC) provides a single, central management point that the administrator can use to not only collect information on all Steelheads, but also configure and update Steelhead appliance configurations. Considered essential for networks with larger number of Steelheads, the CMC simplifies the deployment and management of numerous Steelhead appliances deployed at remote sites.

Customers with large deployments of Steelhead appliances have found some of the following CMC capabilities to be essential when managing hundreds of Steelhead devices:

- **Touchless Steelhead configuration –** New Steelhead appliances will automatically contact the CMC for configuration and operation information, and start accelerating applications immediately.
- **Flexible group management –** The powerful grouping feature provides administrators with the flexibility of creating groups of Steelhead appliances based on geographical location, business function, or personal preference. Once groups have been created, an administrator can broadcast a configuration to a specific group or to multiple groups as desired.

- **Intuitive configuration profiles –** Configuration profiles are a set of configuration values that may be created, edited, re-used from existing configurations, and applied to newly-deployed Steelhead appliances.
- **Configuration backup/restore –** Includes a facility to backup and restore binary configuration files to/from distributed Steelhead appliances.

Also useful are a host of reporting and management capabilities that facilitate the ongoing operation and monitoring of large numbers of remote Steelhead appliances.

Deployment Approaches

There are a number of alternative methodologies used for Riverbed deployments. Each approach has various advantages and disadvantages to considered, and some may be more scalable than others. Not all deployment approaches may be appropriate for a given network environment. Customers can be confident that at least one of the methodologies described in this document will be suitable for their network environment.

"Most Riverbed customers choose the physical in-path deployment approach, which leverages the Steelhead appliance's autodiscovery capabilities. This capability allows Riverbed Steelhead appliances to automatically connect and peer with thousands of remote Steelhead appliances without any additional configuration."

The following approaches have been used at different times and for various purposes by existing Riverbed customers:

- Out-of-path deployment
- Physical in-path deployment
- Virtual in-path deployment using WCCP
- Virtual in-path deployment using policy-based routing (PBR)
- Clustered deployment using Riverbed Interceptor 9200 appliances

Each approach, along with its strengths and weaknesses, is discussed below.

Out-of-Path Deployment

The primary advantage of a physical out-of-path deployment is that it is completely non-disruptive. Particularly in the data center, cables do not have to be pulled and network traffic will not be disrupted when installing and operating the Steelhead. This allows for quick and easy installation for short-term product evaluations and pilot testing.

On the other hand, this approach's chief disadvantage is that optimized traffic is not transparent to the server. The client still sees the server's real IP address, and apparently communicates directly with the server. However, the server sees requests as being initiated from the server-side Steelhead appliance, which is acting as a visible proxy on behalf of remote client workstations.

Because the server-side Steelhead appliance is a visible proxy for the real client, and is initiating client requests using its own IP address, asymmetric routing is not an issue for this deployment methodology. The server simply responds directly to the server-side Steelhead appliance in the data center and traffic is routed directly to the appropriate Steelhead appliance.

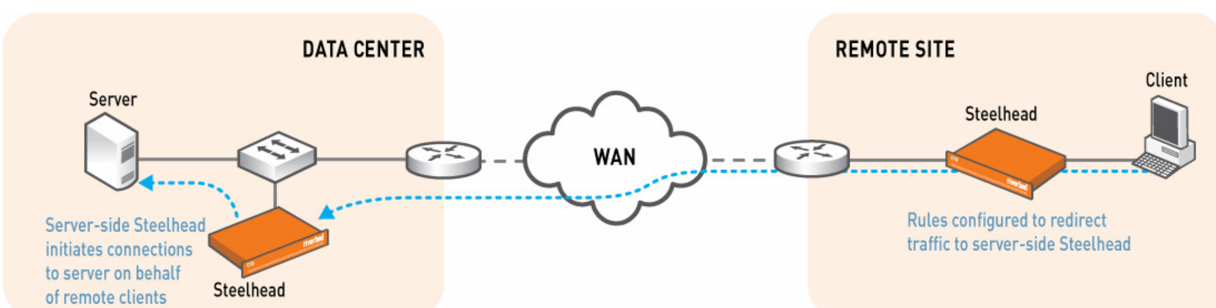

***Figure 3** – Out-of-Path deployments are easy and non-disruptive to deploy. However, rules must be configured on the client-side Steelhead appliance and the optimized traffic is not transparent end-to-end between client and server.*

From the client side, routing of optimized connections is achieved by installing rules in the client-side Steelhead appliance at the remote site. Each client-side Steelhead appliance in the remote site must specify the traffic to be optimized and the IP address of the server-side Steelhead. In the simplest case, this can be accomplished by a single rule.

Incremental scaling of this solution is rather straightforward—adding server-side Steelhead appliances proportionately increases optimization capacity in the data center. There is no limit to the number of Steelhead appliances that can be added for scaling; however, each new Steelhead added in the data center must have corresponding rules configured in the designated client-side Steelheads in the remote site. The Riverbed Central Management Console (CMC) can aid in the configuration of rules in large numbers of client-side Steelhead appliances.

High availability is similarly achieved by configuring rules in the client-side Steelhead appliance; each rule is accompanied by the IP address of the backup Steelhead in the data center. If an active data center-side Steelhead experiences a failure, then traffic continues to be optimized by the designated backup Steelhead. If the client-side Steelhead in the remote side experiences a failure, then traffic will continue to be optimized if there is a redundant Steelhead at the remote site; otherwise it will be passed-through without optimization.

Physical In-Path Deployment

Most Riverbed customers choose the physical in-path deployment approach, which leverages the Steelhead appliance's autodiscovery capabilities. This capability allows Riverbed Steelhead appliances to automatically connect and peer with remote Steelhead appliances without any additional configuration. Through this methodology, full-mesh connectivity among hundreds or thousands of Steelhead appliances at different remote sites can be formed without explicitly identifying the name or IP address of peer devices or the subnets reachable through them.

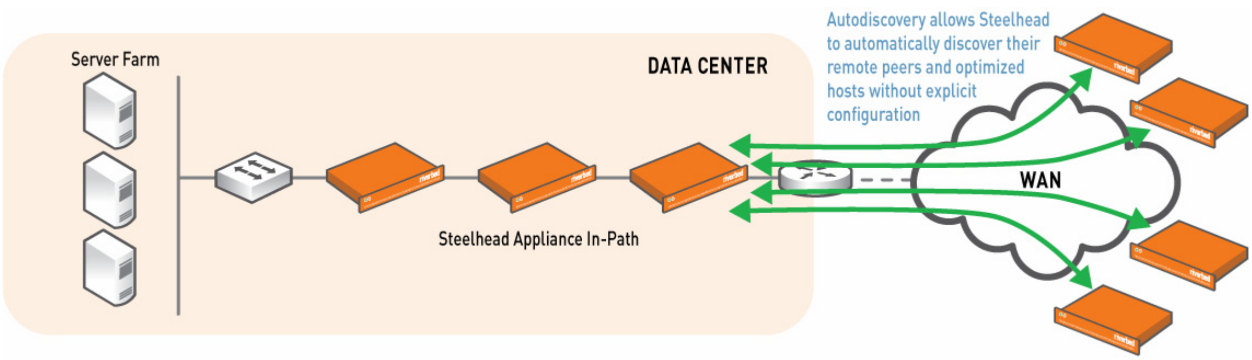

Figure 4: *Physical in-path configurations place the Steelhead appliances in the physical path of the WAN traffic flow. This approach is much easier to maintain and configure due to the autodiscovery capabilities of the Steelhead appliance.*

One concern of this approach is that it does require re-cabling and insertion of the Steelhead appliances into the physical network bottleneck, typically between the WAN router and LAN switch. Some customers may pause at the thought of introducing a new device in such a sensitive location of their production network.

However, most customers are comfortable with the in-path deployment approach because of the Steelhead appliance's fail-to-wire capability. The advantages of simplicity and easy configuration through this approach outweigh most issues involved with recabling, as this methodology avoids the need to configure or disturb other functioning network devices.

The physical in-path deployment methodology also has the additional advantage of being completely transparent to both the client and server. The end-to-end transparency provided through this approach ensures that all applications can seamlessly receive optimized performance.

Most large data centers support multiple redundant Ethernet connections. To address this for the physical in-path deployment approach, Steelhead appliances optionally support multi-link NIC interfaces that allow each appliance to be physically deployed in the path of up to six separate Ethernet connections. Use of multi-link NIC interfaces not only addresses scaling—no need to dedicate a separate appliance to each Ethernet link—but it also addresses asymmetric routing that might exist in the network. In addition, the use of multiple redundant links and multi-link interfaces means that Steelhead appliances can be inserted into the network without breaking connectivity: each appliance is inserted one link at a time, taking advantage of the network's existing failover mechanisms at the point where any one link is interrupted to rearrange cables.

Asymmetric routing is also addressed through the TCP connection forwarding capability discussed earlier. There may be scenarios where asymmetric paths may exist through geographically diverse locations, where it is impossible to deploy a single Steelhead appliance over all of the network connections. In this case the TCP connection forwarding capability allows each Steelhead to re-direct TCP flows to the appropriate Steelhead appliance.

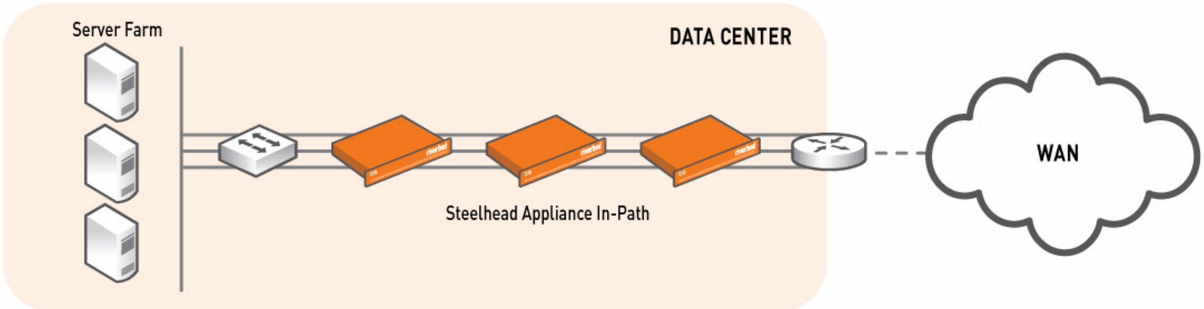

Figure 5: *Special multi-link NIC cards can be installed in Steelhead appliances, allowing each appliance to support multiple redundant Ethernet connections. Asymmetric routing issues can also be addressed in this manner.*

Load balancing and capacity expansion in the physical in-path deployment is also straightforward—just add more Steelhead appliances in series to achieve the desired scale. As each Steelhead appliance reaches its capacity in terms of optimized TCP connections, the next Steelhead in line starts intercepting and optimizing connections. This process repeats until a maximum of five Steelhead appliances are fully-utilized and optimizing TCP connections.

This approach also delivers high availability. If a Steelhead fails, then the TCP connections that it was formerly optimizing will spill over to the next Steelhead appliance in the series of appliances. N+1 availability is achieved by deploying one extra Steelhead in the chain of Steelhead appliances.

If in the amount of traffic exceeds that which five Steelhead appliances can handle, then a different deployment methodology should be used, such as the Interceptor device.

Virtual In-Path Deployment using WCCP

If the idea of placing a new hardware device physically in the path of the WAN traffic flow is of concern, then customers using certain Cisco routers or switches can use WCCP to deploy Steelhead appliances. WCCP can be a very effective methodology to scale out deployment of Riverbed Steelhead appliances; however, this methodology does depend heavily on the stability and performance of the WCCP implementation in the Cisco router or switch.

WCCP allows a Cisco switch or router to re-direct traffic to a Steelhead that is configured and deployed so it is physically out-of-path. This is sometimes referred to as a "one-arm" configuration. As a result, no re-cabling or other physical changes to the pre-existing network infrastructure are necessary—Steelhead appliances can be easily added to the network in the same process as any other server host. Furthermore, end-to-end transparency is preserved, as the client and server both use their real IP addresses to communicate.

WCCP scales easily through use of WCCP groups. Up to eight WCCP groups can be configured on a Cisco router or switch, and each WCCP group can accommodate up to 31 Steelhead appliances each. As a result, hundreds of Steelhead appliances can theoretically be deployed through use of WCCP.

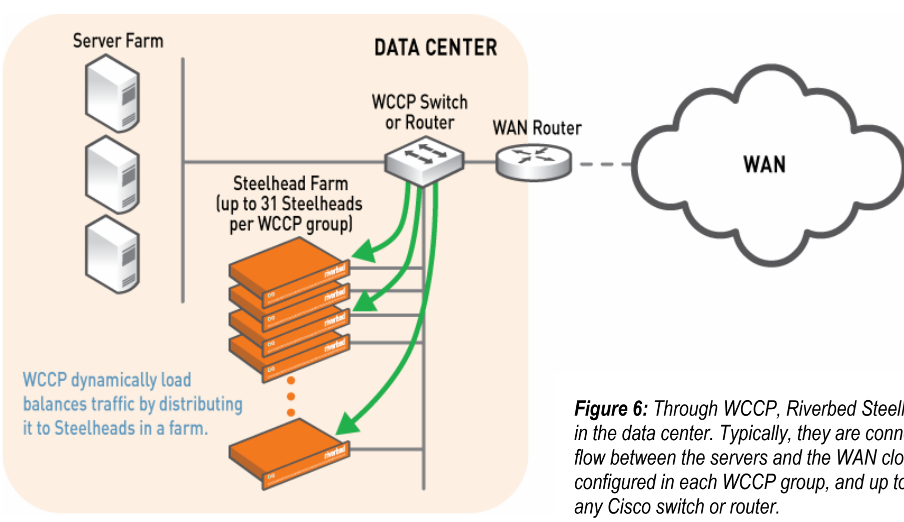

Figure 6: *Through WCCP, Riverbed Steelhead appliances can be deployed anywhere in the data center. Typically, they are connected to a switch in the path of the network flow between the servers and the WAN cloud. As many 31 Steelhead appliances can be configured in each WCCP group, and up to eight WCCP groups can be configured on any Cisco switch or router.*

WCCP has its drawbacks. It requires additional configuration on the Cisco router or switch. The addition of the new WCCP configuration and associated Access Control Lists (ACL's) introduces the potential for configuration error. Also the packet forwarding responsibilities for the WCCP switch or router are increased compared to the previous case prior to deployment of the WCCP and the Steelhead appliances. As a result, traffic can bottleneck and performance be affected due to the packet processing limitations of the WCCP switch or router.

Another drawback of WCCP is that some implementations are either not supported, or not supported very well in Cisco's high-end switching routers. Riverbed has documented a fair number of bugs and other implementation issues in specific Cisco IOS releases, and an appropriate IOS version that is known to be free of WCCP-related bugs should be used.

Finally, when performing load balancing of Steelheads through WCCP, there are additional issues to address. When a Steelhead appliance falls out of a WCCP group for whatever reason, it could cause the load balancing algorithm to recalculate for all TCP connections, causing TCP connections to reset when they are redistributed to different Steelhead appliances. The impacts of this behavior can be mitigated through use of the TCP connection forwarding capability in the Steelhead appliances (which allows each Steelhead to forward TCP connections to the original Steelhead processing that connection), and by limiting the number of Steelheads in each WCCP group.

Despite the above issues, WCCP remains a viable approach for scaling Riverbed deployments. Packet forwarding is performed natively by Cisco equipment, and no Riverbed equipment needs to be physically deployed in the path of the network traffic.

Virtual In-Path Deployment using Policy Based Routing (PBR)

Policy Based Routing offers many capabilities that are similar to WCCP; customers that are hesitant about deploying Riverbed equipment in the physical path of the network traffic may consider PBR as an alternative to WCCP. However, unlike WCCP, PBR is widely available in non-Cisco networking switches and routers. Any layer-4 load balancer/switch can be used to forward traffic to the Steelhead appliances through PBR.

One of the potential drawbacks of the WCCP approach is the increased traffic load and processing requirements for the WCCP switch or router. Policy Based Routing (PBR) is an alternative approach that continues to use inline Cisco equipment to forward traffic to the Steelhead appliances. Its main advantage is that forwarding throughput is higher than the WCCP approach in many cases, as PBR in the Cisco equipment can be supported through Cisco Express Forwarding (CEF). As a result, forwarding throughput can be in gigabit-per-second range.

However, the main drawback between PBR is that it requires individual rules for forwarding of every traffic class. Each Steelhead appliance must be associated with at least one forwarding rule. For large numbers of Steelhead appliances and many different classes of traffic, this can add up to a significant amount of configuration work, making it harder to scale.

Similarly, load balancing through PBR is accomplished statically with fixed rules. Each load balanced Steelhead has at least one corresponding rule in the switch or router forwarding traffic to it through PBR.

High availability for PBR can be achieved through the "Multiple Tracking Options" capability in Cisco switches and routers; non-Cisco equipment has similar features that allow traffic to be handled differently when the Steelhead appliance fails to function.

"Riverbed's largest customers use the Interceptor 9200 to scale out their deployment of Steelheads in the data center. The Interceptor 9200 scales Riverbed deployments to support 100,000 or more remote users and thousands of remote sites".

Clustered Deployment using Riverbed Interceptor 9200 appliances

The Steelhead appliance's TCP-flow oriented optimization technique creates a unique set of challenges. Riverbed supports a broad range of TCP-based applications, including many that use long-lived TCP connections that can persist for days or months. On the other hand, traditional layer-4 load balancers are designed to support web applications that use short-lived TCP connections that persist for a few seconds or minutes. While appropriately-configured layer-4 load balancers that are commercially available are certainly an option for deployment with the Steelhead appliance, they may not handle the broader class of TCP connections ideally in some cases.

Furthermore, existing layer-4 load balancers perform their responsibilities in a uni-directional manner, without regard to the existence of asymmetric routing that may exist in the network. For Riverbed Steelhead appliances, this is non-optimal because each Steelhead must be able to process all packets in both directions of a given TCP connection, even in the presence of asymmetric routing.

For the above reasons, Riverbed has designed and created its own specialized traffic director device called the Interceptor 9200. The Interceptor 9200 was designed to gracefully distribute TCP sessions—both short-lived and long-lived connections—for clusters of Steelhead appliances. Unlike traditional layer-4 load balancers, the Interceptor 9200 will consistently ensure that the same Steelhead appliance will service and process a given TCP connection—for the entire life of that TCP connection.

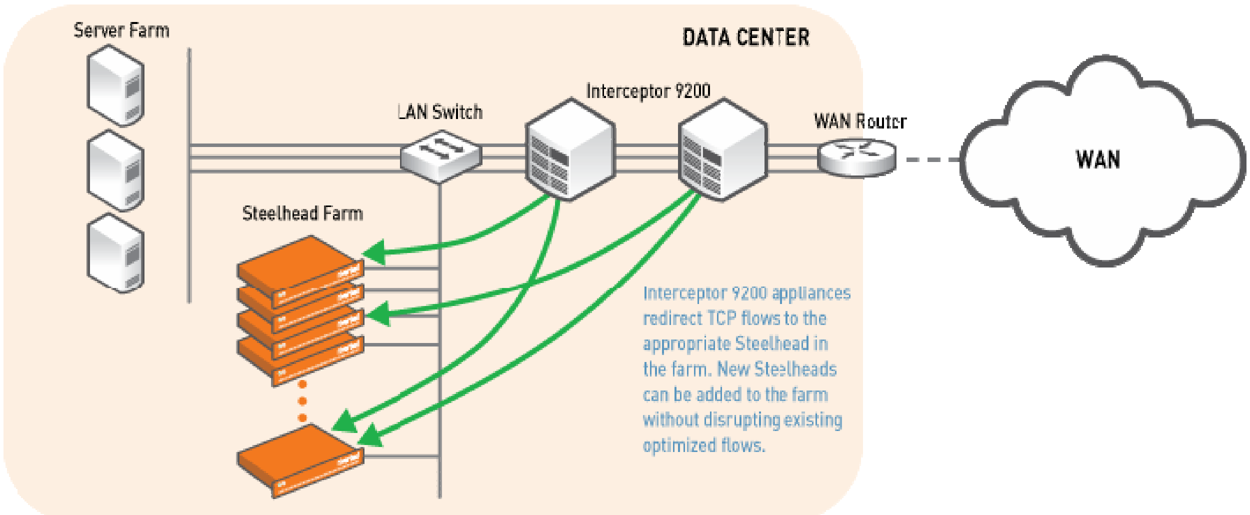

Figure 7: *Interceptor 9200 in a redundant configuration. Interceptor appliances forward traffic to the Steelhead farm deployed anywhere in the data center. Steelhead appliances can be added "on demand" without disrupting existing optimized traffic.*

Riverbed's largest customers use the Interceptor 9200 appliance to scale out their Steelhead deployment in the data center. The Interceptor 9200 appliance allows Steelhead appliances to support hundreds of thousands of remote users and thousands of remote sites.

The Interceptor 9200 appliance is deployed physically in-path. Other than re-cabling of Ethernet connections, it does not require changes to the existing network. Since every Interceptor appliance has multi-link network interfaces, deployment into a network with redundant links does not require breaking network connectivity. Once deployed and configured, the Interceptor device can support forwarding for optimization by the Steelhead cluster of up to 4Gbps of traffic throughput and one million TCP connections. Once the Interceptor 9200 is deployed and operational, Steelhead appliances can be added or removed on demand, with minimal risk and disruption to existing TCP connections.

Interceptor 9200 devices can be deployed in pairs in order to deliver high availability. A redundant pair of Interceptor 9200 devices will share TCP connection information with each other. Should one Interceptor device experience a failure, the remaining operational device will forward TCP flows to the correct Steelhead appliance Interceptor without disruption to the application data flow.

Up to 32 Steelheads appliances can form a cluster "farm" that can be deployed anywhere in the data center network. The Interceptor 9200 uses a connection forwarding mechanism that re-directs TCP connections to the Steelhead cluster across routed subnets.

The Interceptor 9200 is key to the plans of Riverbed's largest customers for scaling deployment of their Steelhead appliances. The design parameters of the Interceptor 9200 allow it to support 100,000 or more remote users and thousands of remote sites.

Summary

As with any enterprise-wide solution, full-scale deployment of the Steelhead appliances requires careful planning and detailed consideration of the various scaling and deployment issues for each specific network environment. Scaling a Riverbed deployment for hundreds of thousands of remote users is a challenge. However, it can be achieved through selection of the appropriate deployment methodology and a gradual phased deployment process.

Riverbed is the clear market leader in scaling WAN optimization solutions for large enterprise networks. Steelhead appliances have been deployed in greater numbers at more production network environments than any other competitive WAN optimization product. Customers deploying Riverbed can draw upon the experience and lessons learned from previous customers with similarly large and complex network environments.

About Riverbed

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

The Wall Street Journal named Riverbed the winner of its 2005 Technology Innovation Award in the Network/ Broadband/Internet category, and *InfoWorld* has named Riverbed's Steelhead appliance "Technology of the Year" in both 2005 and 2006, as the "Best WAN Accelerator". In addition, Network Computing named Riverbed the winner of its 2006 Well-Connected Award for Remote Office Network Infrastructure, and *eWeek* named Riverbed the 2006 Excellence Award winner for Networking Infrastructure. Riverbed's award-winning solutions are available worldwide from resellers who are members of the Riverbed Partner Network, from Riverbed OEM partners, or directly from Riverbed.

Riverbed Technology, Inc.
501 Second Street, Suite 410
San Francisco, CA 94107
Tel: (415) 247-8800
www.riverbed.com

Riverbed Technology Ltd. UK
200 Brook Drive, Green Park
Reading RG2 6UB
United Kingdom
Tel: +44 118 949 7002

Riverbed Technology Pte. Ltd.
350 Orchard Road #21-01/03
Shaw House
Singapore 238868
Tel: +65 68328082

Riverbed Technology K.K.
Shibuya Mark City W-22F 1-12-1
Dogenzaka, Shibuya-ku Tokyo
Japan 150-0043
Tel: +81 3 4360 5357