WHITE PAPER

Optimizing Thin-client Traffic over the WAN

OPTIMIZING THIN-CLIENT TRAFFIC OVER THE WAN

Introduction

Thin-client traffic is a large and growing component of many enterprise network environments. Examples of thin-client traffic are Remote Desktop Protocol (RDP), which is used by Windows Terminal Server, and Citrix Independent Computing Architecture (ICA), which is the protocol used by Citrix's XenApp computing platform. There are multiple wide area network (WAN) optimization techniques that can be applied to thin-client traffic delivered over the WAN. However, techniques designed to optimize delivery of thin-client traffic should not be confused with the broader set of WAN optimization technologies used to produce dramatic performance gains for file-based traffic.

Where improved performance for thin-client computing over the WAN is the primary objective, Riverbed Steelhead product capabilities match or beat competitive WAN optimization solutions.

This document presents recent lab results showing that Riverbed Technology can improve the performance of thin-client traffic over the WAN. At the same time, the document explores various issues and methodologies around the challenge of optimizing thin-client traffic over the WAN, and shows how Riverbed Steelhead products are able to deliver not only compression and data deduplication of thin-client traffic, but also improved responsiveness for the remote end-user accessing the application over the WAN.

Can Riverbed optimize thin-client traffic?

Some competitive vendors have questioned the effectiveness of Riverbed's WAN optimization technology when applied to thinclient application traffic. The reality is that where improved performance for thin-client computing over the WAN is the primary objective, Riverbed Steelhead product capabilities match or beat competitive WAN optimization offerings. Riverbed Steelhead products are capable of providing excellent compression and data reduction results above and beyond that which is provided by the default compression capability embedded within a thin-client platform. This can easily be shown in tests such as those illustrated below in Figure 1. The tests script a fixed series of operations that open and close several files using the Citrix ICA thin client platform. Our test procedure measures the amount of KB sent over the WAN as a result of these operations under various test conditions, including with and without Riverbed WAN optimization. The results show that while the Citrix default compression capability was able to eliminate about 50% of the original raw uncompressed data, Riverbed data reduction technology was able to achieve better results. Specifically, in this test Riverbed delivered 21% better data reduction results on the 1st (cold) transfer, and 45% better data reduction results for the 2nd (warm) transfer compared to the default compression capability embedded in the Citrix ICA solution.

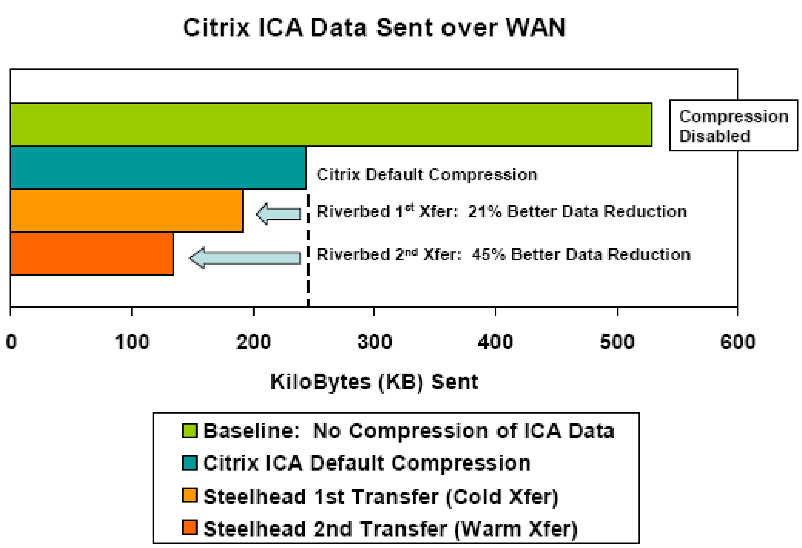

Figure 1: Riverbed achieves superior data reduction results compared to the default Citrix compression

Note that the above tests were achieved for applying Riverbed data reduction algorithms on a single individual Citrix ICA session. However, we have also observed additional cross-session data reduction benefits when the Steelhead stores redundant byte patterns that are leveraged across different Citrix ICA and RDP thin-client sessions. Generally, even better results than those shown above will be achieved in a network environment with large numbers of simultaneous thin-client sessions.

Compression doesn't always yield better thin-client performance

While the above test results are certainly interesting, it is important to note that compression and data deduplication by themselves do not always result in "accelerated" thin-client performance. Application responsiveness is affected not only by limited bandwidth, but also by high latency in the WAN. This is especially important for applications accessed through a thin-client platform, because the nature of thin-client computing requires that interactions between the thin-client platform and the application delivery server be delivered with as little latency as possible.

Competitors often discuss acceleration benefits for thin-client traffic in terms of compression gains. However, a compression gain does not necessarily improve the responsiveness of the thinclient application.

For example, provided adequate bandwidth exists in the network, then it generally requires about the same amount of time for a 1500-byte packet to travel across the WAN as it does for a 60-byte compressed representation of the data in that original packet. Furthermore, the time required to transform data into and out of its compressed or deduplicated format may add additional latency and jitter.

This observation is particularly important for thin-client applications. Even if a WAN optimization product is able to deliver significant compression and deduplication for the thin-client traffic, this achievement does not necessarily indicate that the enduser's application performance will improve as a result. A pre-existing problem with an application's responsiveness and performance may in fact be due to the distance and latency that exists over the WAN between the application client and server; should this be the case, then compression and deduplication alone may deliver limited relief to performance and responsiveness issues experienced by the actual end-user of the thin-client-based application.

For our next test, we used WAN optimization techniques that go beyond just compression and data deduplication. Our objective is to improve the responsiveness of the thin-client application as seen by the remote end-user, in a network environment experiencing traffic congestion. In this test we measured access to a virtual Windows desktop accessed through Remote Desktop Protocol (RDP). The simulated network was a T1 WAN with 100ms of latency. Additional background traffic was added to the simulated WAN to create the congested network environment. The results of this test are illustrated in figure 2 below, which show a time-based responsiveness improvement of between 12% to 38% when using the Riverbed WAN optimization compared to only the native RDP compression without any Riverbed optimization.

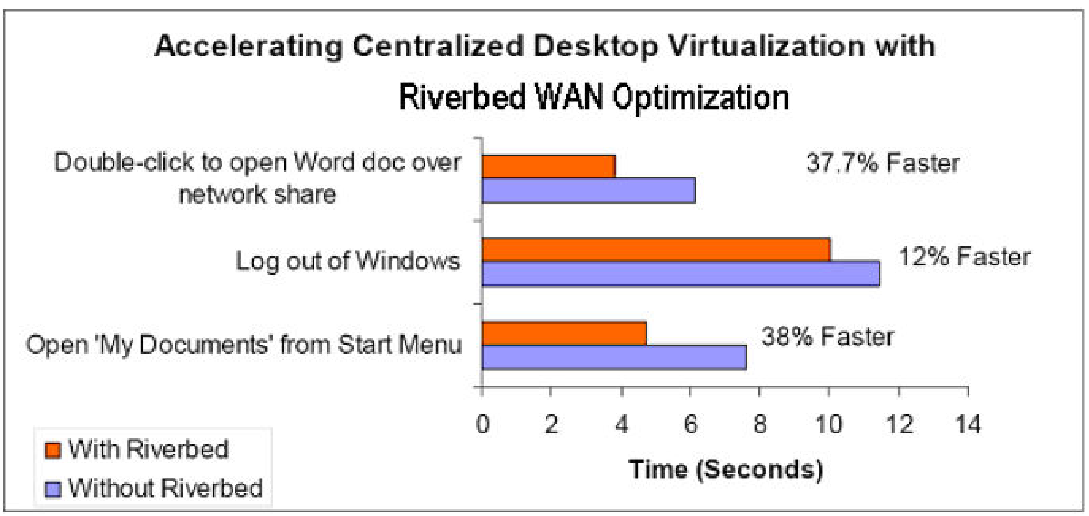

Figure 2: Riverbed optimization delivers faster application responsiveness compared to using only native RDP compression

Techniques for optimizing thin-client traffic

There are an assortment of features and capabilities used by WAN optimization solutions that are effective and beneficial for thinclient traffic. The following is an explanation of each of these features and capabilities:

- Compression and data deduplication of thin-client traffic
- Automatic suppression of Citrix ICA encryption and compression
- Compression and data deduplication of competing file-based traffic
- Enhanced TCP transport
- Quality of service: priority queuing
- Quality of service: bandwidth reservation
- Quality of service: hierarchical or inbound policy configuration
- Quality of service: unwrapping of Citrix ICA virtual channels

Compression and data deduplication of thin-client traffic

Data compression and deduplication mechanisms can be applied directly to thin-client traffic. To improve the effectiveness of the data reduction algorithms, the native compression and/or encryption in the thin-client application platform should be disabled. Compression yield on the thin-client traffic varies depending on the type of application data being sent across the WAN and the effectiveness of the data reduction algorithms that are applied. Because real-time thin-client traffic is particularly sensitive to delay, best performance results are achieved by using a data reduction mechanism that only stores data in memory. For some solutions, this means applying conventional LZ or GZIP-based compression algorithms that are implemented completely in memory. In the case of Riverbed, the Steelhead product solution offers a memory-only data reduction mode known as SDR-M. SDR-M is based on Riverbed's advanced Scalable Data Reduction technology that identifies byte-level data redundancies as very small granularities averaging 100-bytes, which is far more granular than competitive solutions.

Riverbed's SDR-M typically delivers 10% to 40% data reduction above and beyond the embedded compression capabilities in thin-client platforms such as Citrix ICA and RDP. For some types of application data delivered over the thin-client platform, the data reduction results can be even better, as illustrated earlier in figure 1 of this document.

In the Riverbed Steelhead solution, SDR-M is selectable on a per-flow basis. This allows the same Steelhead to optimize a mix of thin-client and traditional network traffic, applying SDR-M appropriately to thin-client traffic flows while also applying disk-based SDR algorithms to CIFS, Exchange, FTP, and other common enterprise traffic types. Note that for some competitive products, memory-only data reduction mode is a global setting for all flows optimized by the WAN optimization device. Using such products may require separate sets of WAN optimization devices—one to apply memory-based algorithms to thin client traffic, and the other to apply disk-based algorithms to traditional file-based network traffic.

Automatic configuration of Citrix ICA encryption and compression

Starting with RiOS 6.0, Riverbed's capabilities include the ability to dynamically disable the default encryption and compression in the XenApp presentation server so that the Steelhead can optimize the raw Citrix ICA data. Manual reconfiguration of the XenApp server is not necessary—disabling of the encryption and compression is handled automatically as long as Citrix ICA optimization is enabled in the Steelhead management interface.

Citrix WANScaler/Branch Repeater offers a similar capability to suppress the default encryption and compression performed by the XenApp server, but Citrix requires additional configuration changes in the XenApp server itself to enable this feature. In contrast, Riverbed's feature to configure ICA encryption and compression works transparently, without any administrator reconfiguration or patching of the XenApp server.

Compression and data deduplication of competing file-based traffic

As previously noted, applying SDR-M to thin client traffic typically achieves 10% to 40% data reduction above and beyond the default compression achieved by the thin client platform. On the other hand for bulky file-based traffic such as CIFS, Exchange and FTP, SDR typically achieves 70% to 95% data reduction. In many or most enterprise networks, file-based traffic consumes far more overall bandwidth than thin-client traffic.

When this is the case, a significant benefit delivered by Riverbed Steelhead products is their ability to eliminate most of the traffic that previously would caused network congestion in the WAN. As a result, the remaining thin-client traffic can be delivered across the network with a minimal amount of jitter and latency, because router buffer queues have been freed from vast amounts of file-based traffic that otherwise would have caused network congestion, packet delay, and packet loss.

Enhanced TCP Transport

Most thin-client applications use TCP as a transport. However, in a congested or poorly-performing network environment, packet loss and retransmission can lead to significant delays in delivery of the real-time data.

To optimize thin-client applications, use of an alternative transport to TCP may be appropriate. Specifically, the alternative transport should not exhibit the slow start and delayed packet retransmission behavior that TCP does. A drawback of using an alternative transport is that its use must be carefully designed and configured in order to avoid creating serious network congestion problems.

Riverbed offers MX-TCP as an alternative transport to normal TCP. The Steelhead product's MX-TCP feature avoids delays associated with slow packet retransmit times often encountered in normal TCP.

Quality of Service (QoS)

QoS policy enforcement is one of the most important mechanisms available for optimizing thin-client traffic over WANs. Although it is not possible to cause thin-client traffic to travel faster than the speed of light, nevertheless through QoS mechanisms it is possible to ensure that thin-client traffic is not unnecessarily delayed as it travels through the network. QoS allows packets containing thin-client data to avoid network congestion, thus allowing them to be delivered with minimal latency and jitter.

Without QoS, not only might the packets containing thin-client data be dropped, but they may also be delayed behind queues of packets from other traffic classes in network routing elements. Any compression or data reduction gains become meaningless if the latency-sensitive thin-client data is delayed in its delivery through the network. As a result, one should consider overweighting the importance of QoS enforcement capabilities in any WAN optimization solution under evaluation.

Delivering effective QoS for thin-client traffic involves two primary capabilities: (1) priority queuing and (2) bandwidth reservation (also known as traffic shaping). Both techniques are important and necessary to ensure optimal delivery of real-time traffic over the network, and WAN optimization devices must be capable of delivering both techniques effectively.

QoS: Priority Queuing

When multiple packets arrive at a network device providing QoS priority queuing services, the higher-priority packets are placed in the high-priority queue, while packets of lower priority are placed in lower-priority queues. This ensures that high-priority traffic, such as thin-client traffic, will always be given forwarding priority through the network regardless of the amount of lower-priority traffic that exists. Because thin-client traffic will have its own dedicated forwarding queues through the QoS device, it will not be affected by the amount of low-priority traffic that is in the network.

Priority queuing is an essential component of any QoS capability where a class of traffic, such as thin-client traffic, must always be given absolute priority for delivery through the network, regardless of the amount of bandwidth consumed by that class of traffic. A device that lacks priority queuing capability will not be able to guarantee that a given class of traffic will always receive priority forwarding service through the network.

Riverbed's priority queuing capability allows Steelhead products to guarantee that thin-client traffic will receive priority over all other types of network traffic. This ensures that thin-client traffic will receive priority processing and forwarding through the Steelhead product with a minimal amount of added latency.

> Surprisingly, Citrix's own WANScaler product lacks a priority queuing capability. As a result, the WANScaler product is unable to guarantee absolute priority for any class of traffic, including for Citrix ICA thin-client traffic.

QoS: Traffic Shaping and Bandwidth Guarantee

Traffic shaping reserves a pre-determined minimum and/or maximum amount of bandwidth for each traffic class. Bandwidth reservation is important for ensuring that a given class of traffic cannot consume more bandwidth than it is allowed. It is also important to ensure that a given class of traffic has a minimum amount of bandwidth available for delivery of data through the network. For example, a device providing QoS services can be configured to guarantee that given class of traffic, such as FTP, will always receive 3Mbps of bandwidth, while at the same time will never be allowed to use more than a maximum of 6Mbps of bandwidth.

The Riverbed traffic shaping/bandwidth reservation capability allows the Steelhead products to reserve a specified amount of bandwidth on the WAN for thin-client traffic. The ability to set maximum bandwidth limits for any class of traffic also allows the Steelhead product to ensure that other lower-priority classes of traffic do not consume more than their share of WAN bandwidth.

However, it is important to note that bandwidth reservation does not prioritize the handling of one traffic class over another. A higher relative amount of bandwidth allocated to a given class of traffic does not necessarily guarantee that it will be prioritized over a different class of traffic. Only a priority queuing mechanism (as described earlier) can guarantee that urgent traffic is always handled before non-urgent traffic.

QoS: Hierarchical or Inbound Policy Configuration

Hierarchical or inbound QoS policy configuration is important whenever multiple remote offices access data from a common shared central site. This allows separate QoS policies to be configured on a per-site or endpoint basis, and it is important because it allows QoS enforcement policies to account for the downstream bandwidth availability at each remote site. Without either a hierarchical or inbound QoS policy configuration capability, QoS policies cannot distinguish the amount of bandwidth available at each remote site; any bandwidth reservation configured at the central site would be indiscriminately applied to all traffic being delivered to all remote sites, regardless of the actual amount of bandwidth at each individual remote site.

Riverbed hierarchical QoS policy configuration allows each Steelhead product to enforce bandwidth guarantees and constraints based on not only the local bandwidth available, but also bandwidth constraints that may exist at each of different remote sites. This capability allows the Riverbed QoS enforcement capabilities to be integrated into shared network infrastructures such as MPLS or ATM clouds. Hierarchical QoS addresses the situation where the amount of WAN bandwidth in the data center available to access the shared WAN cloud may not necessarily match the WAN bandwidth available at each remote site.

QoS: Unwrapping of Citrix ICA Virtual Channels

Citrix ICA is an implementation of thin client traffic that uses encoded virtual channels that each denote one of four different delivery priority levels. A single TCP-level stream of ICA traffic may thus contain multiple channels of traffic that need different handling, and a simple TCP-level QoS solution cannot effectively apply QoS techniques to these channels. Delivering effective QoS specific for Citrix ICA traffic requires that the solution be able to distinguish the virtual channels, and apply the above two QoS capabilities to allocate bandwidth and prioritize delivery for each of these four virtual channels. Steelhead products include patent-pending capabilities to support the application of QoS capabilities to these ICA virtual channels.

Other considerations for optimizing thin-client traffic

There are a number of other optimization approaches and mechanisms that are often discussed as technology components for optimizing thin client traffic. The following is a discussion on some of the techniques employed by some WAN optimization solutions.

Tunneling

Although Riverbed's Steelhead product does not use tunnels, there are a number of other WAN optimization devices that use tunnels to intercept and process the WAN traffic. A consequence of tunneling is that visibility to each individual flow is obstructed for the WAN router, MPLS network service, and any other downstream network element. In WAN optimization devices that use tunnels, each inter-product tunnel coalesces packets from multiple distinct optimized flows. Although the WAN optimization devices at each end of the WAN can distinguish and unpack those flows, devices in the WAN cannot understand it as anything except a single large flow. Because traffic bottlenecks can occur downstream from the WAN optimization device, it is important to that QoS policies be enforceable at these same downstream bottleneck points, such as routers and in the MPLS network cloud.

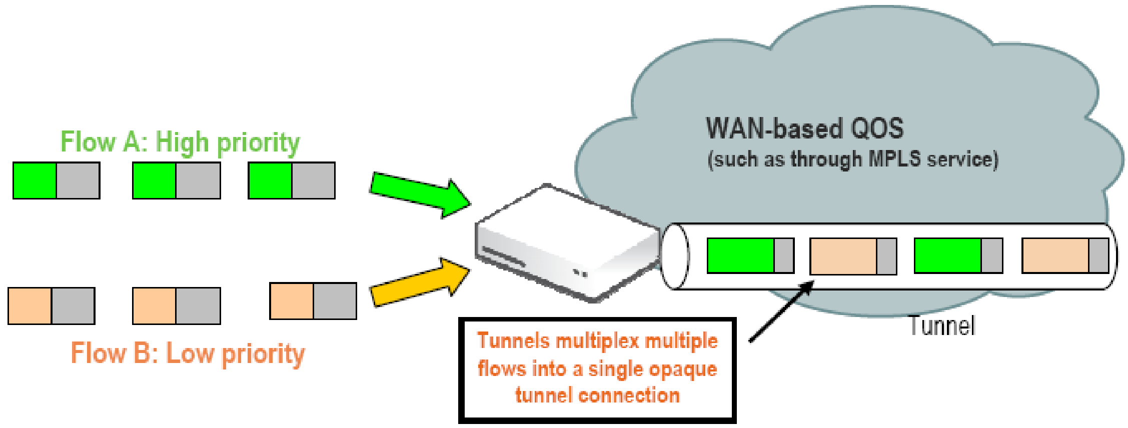

Figure 3: Tunneling solutions obscure visibility into different flows

One of the most important aspects of the Riverbed Steelhead products is that they do not use tunneling. As a result, the Steelhead products integrate well with router and MPLS-based QoS mechanisms. The TCP proxy architecture of the Steelhead product processes each TCP flow independently of other TCP flows and non-TCP datagrams. Each TCP connection remains visible and distinguishable over the WAN because there is a one-for-one mapping between proxy TCP connections on the WAN and the original TCP connections between the client and server. This allows downstream network elements to perform QoS enforcement mechanisms in order to expedite high-priority flows and datagrams while dropping or delaying lower-priority traffic. Because the Steelhead processes each TCP connection independently, downstream QoS devices can apply their QoS policies to the traffic.

VMware View and PC over IP (PCoIP)

PC over IP is an alternative transport option first available in VMware View 4. RDP is the other transport option for VMware View, and using RDP still remains an option for VMware View 4 users. Because PC over IP was originally designed for use in a LAN environment, it uses a UDP-based transport. PCoIP also consumes a substantial amount of bandwidth--between 250kbps to multiple Mbps of bandwidth per user--compared to the typical RDP-based thin client session which typically consumes about 20- 50kbps. Network congestion that may result from using PCoIP in a WAN environment also affects other hosts and applications that attempt to access the shared bandwidth resources available in the WAN infrastructure.

The main problem in applying WAN optimization to PC over IP traffic is that it is encrypted and compressed. With VMware View 4, the default encryption and compression cannot be disabled, meaning that Riverbed and other WAN optimization devices will not be able to deliver compression and data deduplication benefits for PC over IP traffic. In order to apply the WAN optimization techniques to the thin client traffic, RDP should be selected as the transport option for VMware View 4.

VMware View has a feature that adaptively adjusts graphics resolution according to the network conditions. This feature should not be confused for a substitute for WAN optimization, nor does it make WAN optimization any less of a requirement. Rather, this adaptive bandwidth utilization feature basically allows VMware View to gracefully ratchet down the graphics resolution as the amount of available network bandwidth decreases. With half the bandwidth, you will obtain roughly half of the graphics resolution. However, an inadequate amount of available bandwidth will still result in a very unsatisfactory end-user experience.

Header Compression

Header compression is a technique that can be explicitly applied (or not applied) to thin client traffic by tunneling-based devices. But for a TCP proxy-based solution such as the Riverbed Steelhead product, header compression is the natural outcome of the traffic interception and optimization process itself. A TCP proxy already performs the equivalent of header compression on TCPbased traffic (including RDP and Citrix ICA) by coalescing small TCP segments into larger TCP segments.

When using header compression, tunneling devices coalesce small packets together for delivery using a single packet header. This technique leverages the observation that most packets containing thin client data are usually small, especially when compared to file-based TCP/IP traffic. By bundling up the payloads of many smaller packets and delivering them in the payload of a single larger packet, the headers of each individual small packet do not have to be sent over the WAN. Header compression as a technique can only be applied to small packets; it is not possible to coalesce larger packets due to maximum transfer unit (MTU) limitations found on most WAN links.

Blindly applying the header compression process may add jitter and latency to the real-time traffic. This occurs when header compression is applied to the extreme, by unnecessarily delaying earlier-arriving packets so that their data payloads can be buffered up and coalesced with those of other later-arriving packets. While some bandwidth can be saved, the added jitter and delay can affect the performance and responsiveness of the thin-client application.

Comparing Riverbed with competitive solutions

Riverbed is the market leader offering the premier WAN optimization solution. Some vendors, in their attempts to compete, have suggested that Riverbed WAN optimization technology is not effective when applied to thin-client traffic. Nothing could be further from the truth. From surveying the various capabilities and features that are effective for optimizing thin-client traffic, the analysis indicates that Riverbed provides comparable or superior capabilities to those offered by competitive vendors.

Figure 4 below surveys the various vendors that have made prominent claims of being able to optimize thin-client traffic. This comparison provides a high-level overview of what each vendor offers with regard to the capabilities discussed earlier in this document.

	Riverbed	Citrix Branch Repeater	Juniper WXC	Expand	Silver Peak	Cisco WAAS
DRAM-only data deduplication of thin-client data	✓	✓	✓	✓	✓	LZ-Based Only
Automatic suppression of ICA encryption and compression	✓	✓	⛔	⛔	⛔	⛔
Compression and deduplication of competing file-based traffic	✓	✓	✓	✓	✓	✓
Per-flow application of memory-based vs. disk-based algorithms	✓	✓	✓	⛔	⛔	⛔
Enhanced TCP transport	✓	✓	✓	✓	✓	✓
QoS: priority queuing	✓	⛔	✓	✓	✓	⛔
QoS: bandwidth reservation	✓	✓	✓	✓	✓	⛔
QoS: hierarchical/inbound policy configuration	✓	⛔	✓	✓	✓	⛔
QoS: unwrapping of Citrix ICA virtual channels	✓	✓	⛔	✓	⛔	⛔
No tunneling: QoS can be applied downstream in the WAN cloud	✓	✓	⛔	⛔	⛔	✓

Figure 4: Comparison of WAN optimization capabilities for thin-client traffic

One of the more notable vendors is Citrix, which is the vendor that offers not only the XenApp thin-client computing platform, but also the Citrix WANScaler WAN optimization solution. Since Citrix is the original inventor of the ICA thin client protocol, one might naturally reason that they may have technical advantages when optimizing ICA traffic over the WAN.

However, WAN optimization is not one of Citrix's core technologies, and this is particularly evident upon closer evaluation of the Citrix WANScaler/Branch Repeater product. Experience indicates that the data deduplication algorithms used by Citrix's WANScaler/Branch Repeater product are less effective in identifying relatively short repetitive byte patterns. In comparison after comparison, Riverbed's Steelhead solution has consistently delivered better data reduction results for all traffic types (including Citrix ICA traffic) than the Citrix WANScaler/Branch Repeater product.

One of the most glaring weaknesses of the Citrix WANScaler product is its limited capabilities for delivering QoS policy enforcement. Because WANScaler lacks priority queuing mechanisms, it is unable to prioritize delivery of thin-client traffic over other traffic types that might be sharing the network, and incapable of guaranteeing high-priority service for latency-sensitive thinclient traffic. Furthermore, because WANScaler lacks any hierarchical or in-bound QoS policy configuration capability, its limited QoS enforcement functions are even further hindered when WANScaler is deployed in a shared WAN infrastructure such as an MPLS or ATM cloud.

When compared to tunneling-based WAN optimization products such as Juniper, Expand, and Silver Peak, Riverbed advantages go beyond delivering superior performance and data reduction results. Specifically, the tunneling approach used by these vendors create challenges when applying QoS, because it multiplexes thin client traffic with other traffic types into the same tunnel over the WAN. Since high-priority thin-client traffic has been mixed with other low-priority traffic flows, the tunneling approach limits flexibility on where QoS enforcement can be performed. This becomes a problem when a bandwidth bottleneck develops in the network downstream from the tunneling device. Certainly, there will be deployment scenarios such as WCCP-based deployments where the WAN optimization device itself is not the bandwidth bottleneck; with a tunneling-based solution, it will not be possible to apply QoS policy enforcement at the desired downstream bottleneck points in the network.

Of all of the optimization techniques discussed in this white paper, QoS is of particular importance for optimizing delivery of thinclient traffic over the WAN. When QoS is ineffective, absent, or improperly-applied, the negative impact on thin-client traffic can far outweigh any positive benefit that might be obtained through compression, data reduction, or any other WAN optimization technique used optimize delivery of the thin-client traffic. As a result, the QoS-related deficiencies discussed above should not be dismissed lightly when evaluating competitive products.

In the case of Cisco WAAS, it's important to note that WAAS does not have a memory-only DRE mode. Rather, any use of DRE will automatically store the byte-level data on disk, and when applying DRE to thin-client traffic, the amount of time required to read and write data on the spinning disk media will add jitter and latency, which adversely affects thin-client application performance. A practical attempt to use Cisco WAAS to compress thin-client traffic therefore requires that only LZ compression be used as the optimization policy. But there are two things to consider with the use of LZ compression when applied to thinclient traffic:

- It's not clear that LZ compression will yield better results than the embedded compression algorithms already available in thin-client platforms. LZ is a widely-implemented compression algorithm, and may already be used by many thinclient computing platforms.
- LZ compression does not provide additional leverage for redundant data patterns observed across multiple thin-client sessions. This is a distinct advantage for Riverbed's SDR-M capability, which is able to leverage byte-level redundancies observed across different thin-client sessions accessing common screen shots.

Conclusion

Thin-client computing is an important component of many enterprise networks. The Riverbed solution is capable of providing superior WAN optimization capabilities for thin-client traffic, just as it is capable of optimizing file-based network traffic. More specifically, the Riverbed solution is able to deliver superior compression and data deduplication results compared to that achieved by the default compression capability in thin-client computing platforms. The Riverbed solution is also capable of improving the responsiveness of the thin-client application as experienced by the remote end-user.

One of the distinguishing characteristics of the Riverbed solution for optimizing thin-client traffic is the flexibility with which QoS policy enforcement can be applied. Unlike tunneling offerings, the Riverbed transparent TCP proxy approach allows QoS to be enforced anywhere in the network. Furthermore, the Riverbed Steelhead product solution offers the full range of QoS policy enforcement capabilities that are essential for optimizing delivery of thin-client traffic over the WAN.

About Riverbed

Riverbed Technology is the IT infrastructure performance company. The Riverbed family of wide area network (WAN) optimization solutions liberates businesses from common IT constraints by increasing application performance, enabling consolidation, and providing enterprise-wide network and application visibility – all while eliminating the need to increase bandwidth, storage or servers. Thousands of companies with distributed operations use Riverbed to make their IT infrastructure faster, less expensive and more responsive. Additional information about Riverbed (NASDAQ: RVBD) is available at www.riverbed.com

riverbed
Think fast.™

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

Riverbed Technology Ltd.
Farley Hall, London Road, Level 2
Binfield, Bracknell
Berks
RG42 4EU
Tel: +44 1344 401900

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