

The 3 Barriers to IT Infrastructure Consolidation

A Focus on Government Organizations

THE 3 BARRIERS TO IT INFRASTRUCTURE CONSOLIDATION: A FOCUS ON GOVERNMENT ORGANIZATIONS

Introduction

Federal, state, and local government agencies are making the strategic choice to consolidate remote site IT infrastructure into central data centers. They are compelled to move some or all remote file servers, email servers, backup, and other servers because through such site consolidation they can jointly address the need to reduce remote site operating costs and mandates for more rigorous security and compliance.

The stumbling block to consolidation, however, is the severe impact on application performance as seen by remote users. Relocating local servers to a data center and connecting them across a wide area network (WAN) link often results in order-of-magnitude slowdowns to response times and data transfer rates. At these levels of delay business processes are impacted forcing site consolidation efforts to be stalled.

CIOs and CTOs often discover that upgrading bandwidth to remote sites has little or no effect on application performance. The problem lies instead with the way the applications interact with the server across the WAN. Microsoft Windows file systems, Microsoft Exchange®, NAS, backup applications, CAD applications, and many others were developed with the idea that the client and server were local. Across the WAN, however, where congestion, resource contention, diverse routing conditions, and high latencies exist, these applications grind to a crawl.

Riverbed Steelhead products use the powerful Riverbed Optimization System (RiOS) and its combination of patented mechanisms to achieve application acceleration. These mechanisms include Data Streamlining, Transport Streamlining, and Application Streamlining to deliver orders of magnitude increases in application response time and throughput.

Because Riverbed systematically addresses each of the issues affecting application performance over the WAN Riverbed helps government agencies to consolidate remote server infrastructure and deliver consistent end-to-end application performance without resorting to expensive and often ineffective upgrades to WAN bandwidth.

A Short Case Study on Consolidation

Remote site server consolidation is a clear win in terms of reducing operating costs and improving data security. However, there were compelling reasons for distributing server infrastructure in the first place.

The reason many companies chose to place servers at remote sites has been to deliver consistent application performance to remote users working with local data sets. Microsoft Exchange servers, for example, have commonly been deployed at remote sites with only 20-30 users because above those numbers, most Exchange messages end up being between local users.

Provisioning of servers at remote sites, however, often leads to low resource utilization and high costs. Since Exchange servers are typically resourced for a capacity of several thousand users, deploying a dedicated server for a few dozen means inefficient use of server resources. This same issue exists for file servers, and web servers. Worse, all those servers have to be managed, backed up, repaired, and patched.

Centralizing servers at a data center means greater resource utilization, and fewer servers to backup and patch. Since complexity is reduced, such consolidation also means lower IT staff requirements, less chance for errors, and better system security.

Because of the clear benefits, companies are trying to consolidate infrastructure as much as possible, yet many are surprised at how difficult it is to actually complete a successful site consolidation project. They find they can't deliver consistent end-to-end application performance even with significant upgrades to WAN bandwidth.

Consolidation of remote site infrastructure offers significant benefits:

- ✓ Reduces cost and complexity
- ✓ Improves compliance
- ✓ Improves data and network security
- ✓ Improves resource utilization
- ✓ Eliminates need for costly WAN bandwidth upgrades
- ✓ Eliminates write consistency issues associated with caching
- ✓ Frees up WAN capacity for VoIP and video applications

For example, a large division of a Federal civilian agency that focuses on fraud prevention wanted to ensure that its email communications were as secure as possible. In order to do that, they needed to eliminate exchange servers from over 100 locations around the United States. By doing so, they would benefit in three ways:

- There is significantly less likelihood that a thief, hacker, or unsuspecting individual could walk away with important data and private information.
- The cost of deploying the infrastructure to support a nationwide Exchange user base is dramatically decreased.
- The cost of managing, patching, and upgrading the Exchange deployment is decreased significantly.

Of course, the agency wanted to complete this project with no impact to the performance of Exchange in remote locations – the real challenge associated with consolidation. The rest of this paper will discuss the technical challenges in-depth and present the Riverbed-based solution that this agency used to ensure a successful consolidation project.

Three Barriers to Site Consolidation

When WANs are involved, client-server applications that worked fantastically on LANs break down and work poorly, or not at all. This is caused by 3 factors:

1. Constrained WAN bandwidth
2. TCP throughput drop-off with latency
3. Application chattiness multiplies the effect of latency

Constrained WAN bandwidth

WAN bandwidth is often orders of magnitude less than local area network (LAN) bandwidth. A typical remote office has between 64 kbps and T1 bandwidth (1.544 Mbps) or E1 (2 Mbps). Compared to modern LANs that have 100 Mbps to 1,000 Mbps, a remote site typically relies on less than 1% of the bandwidth to access data as a result of server consolidation.

From a pure bits-per-second perspective it's easy to see why moving a large file across a WAN link should take more time than over the LAN. However, it's often the other two constraints, not bandwidth, that result in low application performance.

TCP throughput drops off with latency

All applications rely on underlying communications protocols which, for reliable transport across the network is almost always TCP. TCP sends data in "windows". A window defines the maximum amount of data a sender can transmit before receiving acknowledgement from the receiver. Since it takes a round-trip time to receive the acknowledgement from the receiver the maximum throughput is the amount of data in a window divided by the round trip time. TCP slow start and congestion control features designed to increase reliability also make the throughput problem worse.

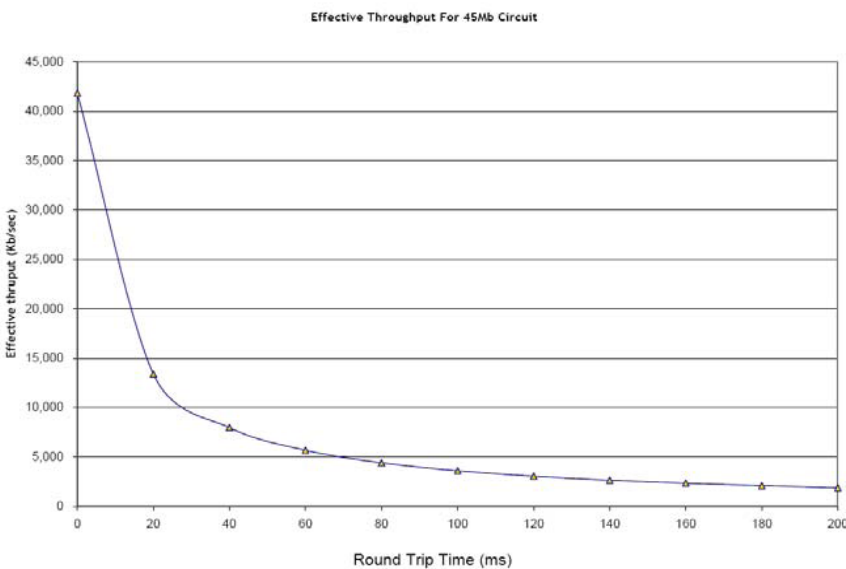


Figure 1: TCP throughput drop-off with latency on a 45Mbps (T3) link

Application chattiness multiplies the effect of latency

On top of TCP, applications have their own communications protocols. For example, Microsoft Windows® uses CIFS, the Common Internet File System. Microsoft Exchange® uses MAPI, the Messaging Application Programming Interface protocol. Web based applications rely on HTTP, and so forth.

Some protocols (application or transport) are extremely “chatty”, which means they generate hundreds or thousands of round trips from client to server, even to accomplish seemingly simple tasks. For example, dragging and dropping a 1 MB file in Windows can trigger over 4,000 WAN round trips. On a LAN, when the latency between client and server is often less than a tenth of a millisecond, those thousands of round trips are completed virtually instantaneously. When the same operation is done on a WAN, the latency is usually in the range of 50 ms to 250 ms, or even more when satellites are involved.

The significant difference shown in Figure 2 (at right), in which completion time goes from 0.4 seconds to almost seven minutes, is why “just moving the servers” doesn’t work – it is very noticeable to users. Application protocols also have a limited amount of data they can transmit on each round trip. So the problem of many round trips is worse for large files. If the application protocol has a “transfer size” of 16 KB, then 16 MB file will require 1,000 trips, just to deliver the data, plus lots of additional round trips generated by the application to manage the data transfer, file system operations, or whatever other operations are required.

LAN vs. WAN Time to Complete

	LAN	WAN
Latency (in ms)	0.10	100.00
Number of Round Trips	4,000	4,000
Time to Complete (in ms)	400	400,000
Seconds / Minutes	0.4 / 0.01	400 / 6.67

Figure 2: A 1 MB file drag & drop in Windows generates 4000 transactions between the client and the server

A similar chattiness issue applies to the TCP layer, which affects web-based business apps, as well as applications like Notes, FTP and other mission critical applications.

Learning from Past Mistakes – Exposing the Myths

Over the past few years, vendors have created a number of products to accelerate application performance. These solutions, often categorized as WAN Optimization or WAFS, have fallen into three categories:

1. TCP optimization
2. Compression
3. Caching

IT professionals have learned that these solutions are either insufficient to address performance across a wide range of applications, introduce additional complexity, or both. While each of these solutions can solve specific issues with application performance across the WAN, the myth still exists that they are general solutions to application performance.

Myth #1: You can solve application performance with TCP Optimization alone

Many IT professionals are aware that TCP as originally defined has a maximum window size of 64 KB (the typical amount of data that can be carried in each TCP round-trip), and that the limit can be modified with some work. In many cases, the configured maximum TCP window size is even smaller – 16 KB or 32 KB, which makes the problem even worse. Even companies who elect to go the route of modifying TCP find that fixing or improving TCP does not help application performance if the application protocol is less efficient than TCP.

For example, in Microsoft Exchange 2003, the window size was increased from 8 or 16 KB to 64 KB. This helps to reduce the number of round trips generated by the application when sending large amounts of data, but does nothing to accelerate operations such as calendaring operations that are bottlenecked by the huge number of application level (MAPI) client-server transactions.

A consolidation project on the brink of failure

A division of state government wanted to reduce the cost associated with its file sharing infrastructure, without making changes to the way users typically interact with their file sharing tools. The agency decided that they should consolidate the file sharing infrastructure into its main data center. In addition, they decided that in the long run they would like to consolidate other applications as well.

Despite the fact that branch offices were a maximum of just 200 miles from the headquarters, this agency saw noticeably poorer application response time once they consolidated the server infrastructure. This poor performance was negatively impacting user productivity in all locations.

The agency’s IT managers now faced a tough choice – do they cancel the consolidation project entirely, or should they invest in technology that could eliminate this performance problem? Read on to find out more about the solution this agency chose to deploy.

Hardware solutions exist to modify TCP's behavior in other ways across the WAN to increase its throughput, but modifying the TCP layer in the networking stack does nothing to improve performance issues caused by higher layer protocols. For many applications like Windows file sharing or Exchange, the application protocol (CIFS and MAPI, respectively) are much chattier and less efficient than TCP itself. Thus, making TCP more efficient can be helpful, but in many cases this approach alone is insufficient.

Myth #2: You can solve application performance with compression

Agencies that attribute application performance issues to lack of bandwidth often reach the conclusion that they can solve the problem by adding compression appliances. This is equivalent to adding more bandwidth. More bandwidth is helpful, but again it's insufficient to solve the problem. Adding bandwidth does not help alleviate the chattiness of the application, which means that all those round trips still have to take place. No matter how much bandwidth you buy, once the initial congestion has been alleviated, the application performance will not be materially affected.

Myth #3: You can solve Application performance with Caching

Some government organizations have investigated caching appliances as a way to enable site consolidation. That approach can work for single data types, but will not provide a general solution and often is used just to hide the underlying performance problem. For Exchange, there are special purpose mail caching appliances available, but they are not a general purpose solution to the problem.

Caching is an application-specific technology: File caching works for file systems, web caching works for web pages, mail caching works for email, and so on. So while adding an Exchange mail cache will help by storing attachments locally, it adds complexity and only affects the perceived performance of Exchange.

Another issue for file caching is write consistency. Often caching products will implement elaborate file locking mechanisms to prevent two users from writing the same file. But in the event of network outages and/or box failure, these mechanisms can fail leading to catastrophic results.

With Exchange 2003 and Outlook 2003 Microsoft introduced integrated client-side caching to address performance across the WAN. This hides the delay in getting email from servers to clients by not displaying any new email headers until the entire email and any attachments are fully delivered. Thus by the time a user is notified that a new message has arrived, the entire message (and any attachments) have already been cached in the client. This does nothing to improve the actual time to deliver messages or the amount of time to download one's inbox. Moreover, to get any benefit from client-side caching, you must deploy both the server (Exchange 2003) and the client (Outlook 2003). Client-side caching can improve the perceived user experience but may cause much heavier traffic across the WAN, since emails are delivered to the client that might have been deleted without being read if only the headers were seen.

Riverbed Breaks the Barriers to Site Consolidation

Riverbed Steelhead appliances and Steelhead Mobile software accelerate the performance of all applications running on WANs over TCP by addressing all three bottlenecks to WAN performance. Steelhead products are the first to address all the issues affecting application performance over the WAN to deliver the highest performance improvements to the widest range of applications.

Based on the Riverbed Optimization System (RiOS), Steelhead products combine patent-pending data reduction, TCP optimization, application-level latency optimizations, and remote office file and management functionality. Together, these technologies provide a comprehensive solution for enterprise wide-area data services, scaling across a range of applications and network topologies to accelerate applications up to 100 times or more. With this kind of LAN-like performance, site consolidation projects can proceed without impacting end users.

Bandwidth Optimization – Data Streamlining

RiOS Data Streamlining works across all TCP applications to reduce bandwidth consumption by 60% to 95%. Data Streamlining works across Windows file sharing (including MS Office), Email (including MS Exchange and Lotus Notes), CAD, ERP, databases, and all other applications that use TCP, to ensure the same data is never sent more than once over the WAN. Data Streamlining also supports rules-based policy administration of optimization classes and packet marking for QoS and route control. Riverbed combines compression with patented Scalable Data Referencing (SDR) bandwidth optimization technology dramatically to reduce the amount of data that is sent across the WAN. SDR replicates data across the network in a new and

unique protocol-independent format to reduce subsequent transmissions of the same data. Rather than attempt to replicate data blocks from a disk volume, files from a file system, e-mail messages, or Web content from application servers, Steelhead appliances represent and store data in a protocol and application-independent format.

As data is sent through a Steelhead appliance, SDR segments the data into variable size segments and creates a short "references" to each segment. As data and their accompanying references are created on one side, they are compressed and sent to the Steelhead on the other side. Thus, once the Steelhead appliances have seen the data once, they need only send references to the data to the other end. Moreover, these references are hierarchical; references can point to groups of references such that a single reference can represent megabytes of data that has moved across the network before.

The elegance of the approach is that the Steelhead products are transparent to the client and server. There are no consistency issues to be tackled even though the data segments may exist in multiple locations. The client-server transactions always occur directly between the client and server, preserving the protocol semantics, even though very little actual data is transferred across the link.

By addressing all areas of WAN performance, Riverbed offers several key advantages over file caching or compression-only approaches for site consolidation:

- ✓ **Broad Applicability** – Steelhead products optimize all TCP traffic covering a broad range of applications. Unlike specific file caching or email caching approaches, Steelhead products deliver performance and bandwidth savings whether a company is centralizing Exchange servers, Notes servers, file servers, NAS, tape backup or any combination of these.
- ✓ **Time and Bandwidth Savings** – Steelhead products provide response time improvement, in addition to compression and bandwidth savings. In contrast, WAN optimization devices that just compress data will reduce the amount of data within a packet but, because they do not terminate TCP, they send compressed data in the same number of round trips as it would take to send uncompressed data.
- ✓ **Better Optimization** – File caches only give you a "hit" when a user requests an identical file to one that's been requested before. Compression devices rarely delivery more than 2-3x improvement. Steelhead products in many cases offer more than 100x and will deliver improvement even on new versions of old files, with different file names, and even different applications using similar data.
- ✓ **Easier Deployment** – Steelhead products require no client or server re-configuration, which means the rollout is much simpler and quicker. In fact, Steelhead appliances can be deployed in as little as 15 minutes, and managed centrally from an optional central management console.

Transport Streamlining

RiOS Transport Streamlining reduces the number of TCP packets required to transfer data by 65% to 98%. Transport Streamlining overcomes TCP limitations by adapting transmission characteristics such as window scale, loss handling, congestion notification, and more. RiOS Transport Streamlining also enables greater utilization of high bandwidth, high latency connections with High-Speed TCP capabilities.

For all TCP based applications, Steelhead products utilize RiOS Transport Streamlining minimize the time it takes to send data across the WAN. One of the key features of Transport Streamlining is Virtual Window Expansion (VWE), which multiplies the effective TCP window size.

Most TCP implementations including Windows 2000 and XP by default send no more than 64KB in round trip across the network. It is often difficult to change these defaults across all hosts and resize buffers in the network elements to accommodate the change. Steelheads implement window scaling across the WAN correctly without host reconfiguration and without requiring larger network buffers for LAN-directed traffic. But beyond window scaling, Steelhead products terminate TCP and repack TCP payloads by substituting references to arbitrarily large amounts of data in combination with Riverbed's SDR technology. This technique "virtually" expands TCP windows beyond the expansion delivered by window scaling because the amount of data that is represented by a reference can be 1 MB, 10 MB or more.

By virtually expanding the TCP window size, the number of round trips is minimized which in turn increases throughput. All of this is done without changing the underlying TCP protocol or changing the client server interaction. In contrast, WAN optimization devices that just compress data on a per-packet basis will reduce the amount of data within a packet but, because they do not terminate TCP, they send compressed data in the same number of round trips as it would take to send uncompressed data.

Transport Streamlining also includes High-Speed TCP functionality to immediately increase the capacity of high bandwidth, high latency links between data centers used for data replication and disaster recovery. On these high-speed WAN links normal TCP can fail to ramp up to full capacity even though plenty of bandwidth is available. This leaves data replication and mirroring applications starved for throughput, thwarting site consolidation efforts with an insufficient data protection.

Riverbed has implemented numerous Internet Engineering Task Force (IETF)-specified congestion control mechanisms in the Steelhead appliance that enable TCP performance to scale to hundreds of Mbps over significant latencies (>100ms RT). Riverbed customers with high-speed WAN links can now achieve full utilization of their investment in network bandwidth without losing or compromising any of the familiar and essential characteristics and benefits of TCP. This includes safe congestion control, even when high-speed TCP connections share WAN links with normal TCP connections.

Application Streamlining

RiOS Application Streamlining provides additional order-of-magnitude application performance improvements by reducing application protocol chattiness up to 98% and minimizing application overhead. By minimizing application demands on the network such as application protocol round trips and required network connections, RiOS can provide massive throughput increases to applications including Windows file sharing (CIFS), Exchange (MAPI), Web (HTTP), and Database (MS-SQL). RiOS also includes important features for maximizing branch office productivity, such as file server capabilities and transparent pre-population of popular data.

To address application chattiness, Riverbed has developed a set of Application Streamlining algorithms, which further minimize the number of round trips taken across the WAN without interfering with the client-server semantic. Application Streamlining works in combination with Data Streamlining and Transport Streamlining in order to provide even higher levels of performance to the most common enterprise applications.

With specific knowledge of application protocols like CIFS, MAPI, SQL, HTTP, and others, Steelhead appliances are able to predict upcoming client requests in advance, inject requests to the server on behalf of the client, and then “bundle” the results of the server interaction into a few round trips. Each round trip avoided saves a discrete amount of time, independent of how much bandwidth is available. When thousands of round trips are avoided, the time saved can be measured in minutes or even hours, depending on the workload.

Application Streamlining also enables Proxy File Services (PFS). This functionality enables an organization to provide disconnected operations to remote users in the event of WAN failures. Proxy File Services gives the IT manager the flexibility to decide what files should be available to particular locations, and maintains current file locking and permissions in the event of WAN failure.

Summary

Site consolidation has a tremendous ROI, as long as user application performance can be preserved. Most enterprises have a range of remote office IT infrastructure: File servers, Exchange® servers, web applications, app servers, Notes® servers, NAS, tape backup, and more. The more IT infrastructure that can be centralized or consolidated to the data center, the higher the ROI for the IT department.

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 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