TANEJA GROUP
TECHNOLOGY ANALYSTS

T E C H N O L O G Y   I N   D E P T H

Wide Area Data Services: Optimizing the Branch

May, 2005

THE NET NET

In February 2004, Taneja Group defined the term Wide Area File Services (WAFS) to describe a group of bright young companies that were optimizing network file access across wide geographic areas. Since that term, we have continued to track the rapid evolution of various optimization challenges for the distributed enterprise, spending significant time talking with both customers and vendors. We reached the conclusion that WAFS must now be considered a subset of an even more encompassing category that we define here as Wide Area Data Services (WDS). The pioneering and defining firm in the WDS category is Riverbed Technology, based in San Francisco. WDS is about much more than file services across the WAN; it encompasses TCP optimization, data reduction for replication, email caching, and a range of application protocol optimizations. This technology is distinct and important because the converged WDS capabilities embodied in the Riverbed Steelhead solution provide us a path to the *truly* serverless, lean, network-efficient branch office. As WDS matures and gains wider market traction in coming months, Taneja Group firmly believes that the long-held vision of an efficiently distributed enterprise with completely centralized IT functionality is achievable. When this happens, WDS will be viewed as the enabling technology set. This white paper provides an overview of the constituent components of WDS, and brief look at the role Riverbed Technology is playing in bringing this technology to market.

Challenges of the Distributed Enterprise

Distributed organizations come in all sizes and shapes. Even small organizations may have multiple locations connected by data networks around the world. For the largest enterprises, even the simplest usable network is a complicated arrangement, while the most elaborate and sophisticated networks have unbelievable layers of replicated and interlocking functionality. Whether a distributed organization is small or large, whether its networks are simple or elaborate, each faces similar business challenges: making the distributed organization efficient and effective in its use of the Wide-Area Networks (WANs) that tie the organization together.

87 Elm Street, Suite 900   Hopkinton, MA 01748   Tel: 508-435-5040   Fax: 508-435-1530   www.tanejagroup.com

Often these business challenges show up as distinct IT and WAN issues at remote offices, satellite offices, or branch offices, which Taneja Group will consider for our purposes simply as *branches*. Poor application performance and poor control of information at branches are both common examples of ineffective use of WANs, which leads to reduced productivity and dangerous exposure to liability. Excessive spending on network bandwidth and high administrative costs for branches are both common examples of inefficient use of WAN resources: throwing away money that could be put to better use.

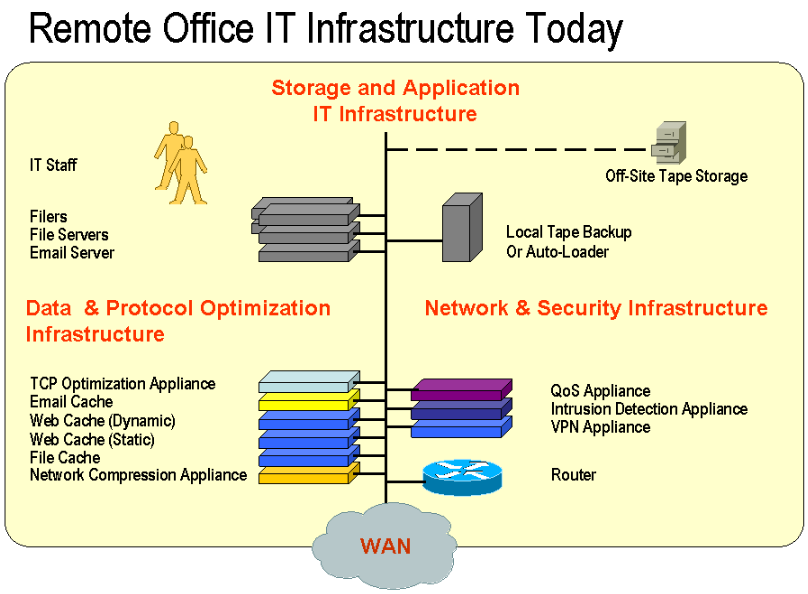

What are the underlying causes of these challenges? The network of a distributed organization typically consists of an identifiable richly-connected *core* and an identifiable weakly-connected *fringe*. The core typically includes most servers, many clients, and most of the organization's available bandwidth. In contrast, the fringe typically includes most of the organization's remote locations. Operations between clients and servers in the core occur at the speeds typical of Local-Area Networks (LANs), whereas operations involving clients in the fringe happen at speeds typical of WANs. Typical WAN bandwidth at the fringe is only 1% or less of the typical LAN bandwidth, while latency (the time required to take a single end-to-end round trip) is often 100 times longer or more on a WAN than on a LAN. Because the fringe has network capacity that is 100 times *narrower* (bandwidth) and 100 times *longer* (latency) than the core, it is not surprising that performance is often a problem.

With this problem in mind, we can consider how to solve the effectiveness and efficiency challenges at branches by making the WAN perform more like a LAN.

Previous attempts to address these challenges have led to the deployment of a variety of local servers, Wide Area File Services (WAFS) for file content, and an assortment of compression or TCP-optimization devices. These discrete technologies can be considered a subset of a larger technology category that encompasses all of them. Namely, Wide Area Data Services (WDS).

The Wide-Area Data Services Solution

Making the WAN behave more like a LAN has two components: 1) delivering dramatic performance improvements, especially for those applications or protocols that show the worst degradation when running over a WAN, and 2) contending with a wide spectrum of protocols that are crossing the networks of a distributed organization.

Solutions that deliver this broad-spectrum improvement of WANs go by the name of Wide-area Data Services (WDS). *Wide* refers not only to the WAN but also to the wide applicability of WDS, and the wide spectrum of bottlenecks improved by WDS. WDS is not only a solution to attack problems of ineffectiveness or inefficiency, but also a means to realize goals that would otherwise be impossible:

- Centralizing distributed infrastructure like file servers, mail servers, Network Attached Storage (NAS), and remote office backup systems – without affecting remote users
- Sharing large files among colleagues on different continents – as if they were in the same building

  Performing backup and replication over long distance WAN links – and completing them during backup windows that were unachievable just a year ago
- Delivering significantly more services on existing WANs – without upgrading the bandwidth

If desired, an organization can eliminate branch servers and devices, replacing them with only network services and WDS. Alternatively, a branch without local servers can dramatically improve its productivity by adding WDS.

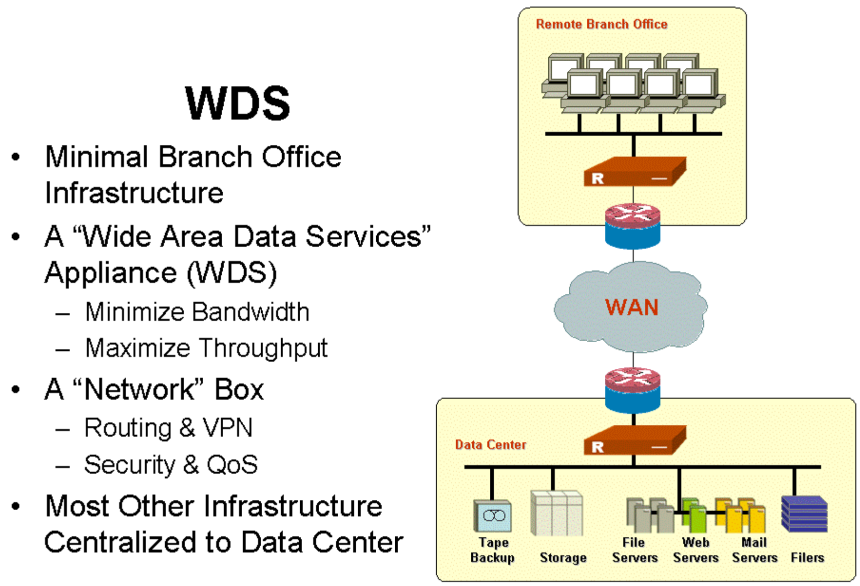

Bandwidth and Latency Bottlenecks Limit WAN Performance

Let's consider the bottlenecks in some more detail. WAN connections typically have lower bandwidth and higher latency than LAN links, but how do those constraints actually affect application performance? There are four distinct bottlenecks, one relating to bandwidth and three relating to latency. The bandwidth bottleneck is straightforward: no application can send data faster than the available bandwidth. However, there are several forms of latency bottlenecks that may impact applications, even when the bandwidth appears to be plentiful.

Latency Bottleneck #1

The first latency bottleneck is caused by the end-to-end acknowledgement behavior of TCP. TCP has a *window* of packets that can be in flight from one end to the other (i.e. between client and server). After the window is full, the sender cannot send additional packets until the destination acknowledges receipt of at least some of what has already been sent. If the maximum window is too small, the throughput of the link will be limited by the rate at which each full window can be sent to the other side and acknowledged.

In theory, this bottleneck should be rare, because well-specified mechanisms exist to allow TCP to use large windows, and most recent operating systems implement those mechanisms. However, settings on both clients and servers are usually more attuned to LANs than WANs, and it is unusual to find clients or servers with TCP stacks matched to WAN latencies.

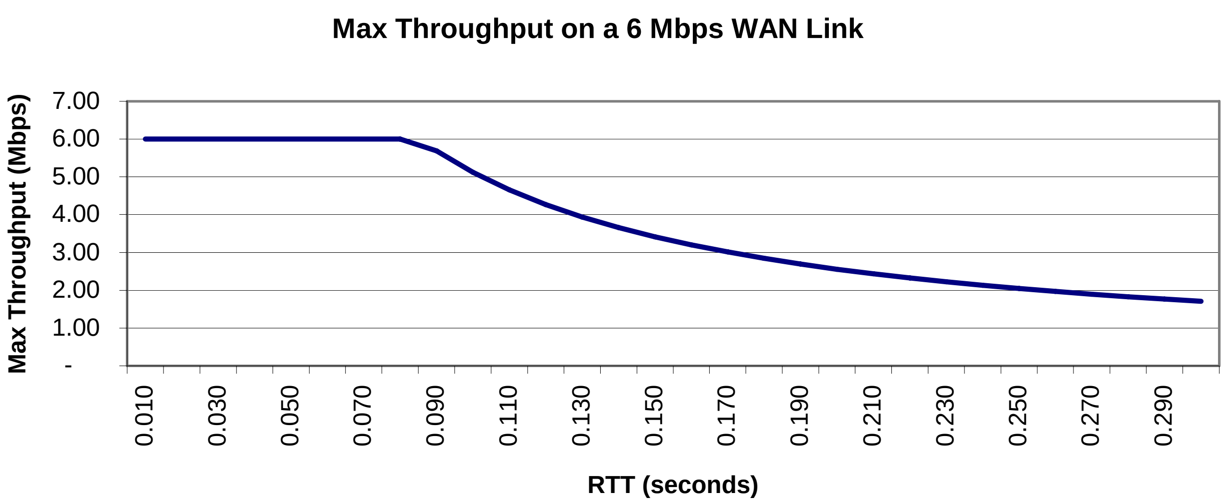

***Figure 1**: Figure 1 illustrates the effective throughput of a 6 Mbps link (i.e. 4 x T1 lines) for a TCP connection with a 64 Kbyte maximum window and increasing latency. For low latencies, the link reaches its bandwidth-determined throughput, but for latencies larger than about 80 ms the first latency bottleneck is narrower than the bandwidth bottleneck.*

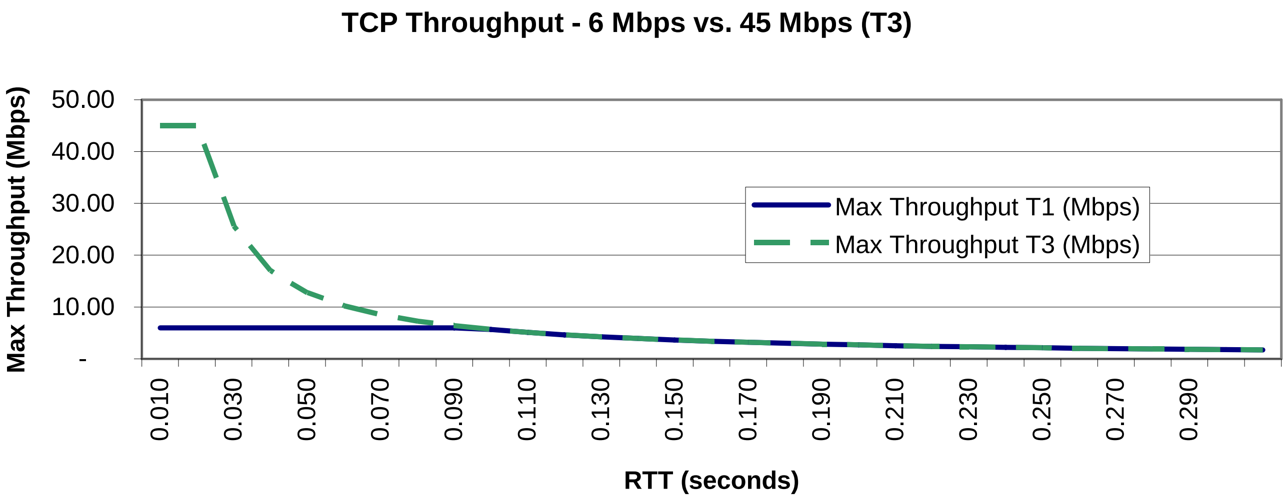

Figure 2: *Figure 2 illustrates the same data as Figure 1, but has added a similar curve for a 45 Mbps T3. On this scale of graph, the 6 Mbps link's throughput looks flat – but more alarming is that the T3 rapidly declines to a level indistinguishable from the 6 Mbps link as latency increases. At a latency of 90 ms or above and these small TCP windows, a T3 will not offer a single connection any more capacity than a 6 Mbps.*

Latency Bottleneck #2

The second latency bottleneck is caused by the slow-start and congestion-control behaviors of TCP. The first latency bottleneck, explained above, is a limit based on the maximum window possible. This second latency bottleneck is caused by TCP not even running at that (probably inadequate) maximum window size all the time. Instead, TCP gradually ramps up its window size when transmission appears to be successful and sharply cuts back its window size when transmission appears to be unsuccessful. In networks with both high bandwidth and high latency, this behavior leads to extended periods in which available bandwidth goes unused. However, this bottleneck is primarily an issue for users trying to fill long fat networks (LFNs), not for our example of a 6 Mbps connection.

Latency Bottleneck #3

The third latency bottleneck is caused by so-called application protocols that are running on top of TCP. Recall that with the first latency bottleneck, the availability of bandwidth didn't matter if TCP was limited by the size of a window of data and the need to acknowledge that data. Analogously, the availability of bandwidth and the avoidance of the first and second latency bottlenecks (at the TCP layer) are irrelevant if the application is limited by the size of application messages it must acknowledge. Application protocols that were originally designed for wide-area environments – such as HTTP and FTP – do not encounter this third latency bottleneck. However, application protocols originally designed for use on LANs – such as Microsoft Windows file sharing via CIFS -- are often severely affected by this third latency bottleneck.

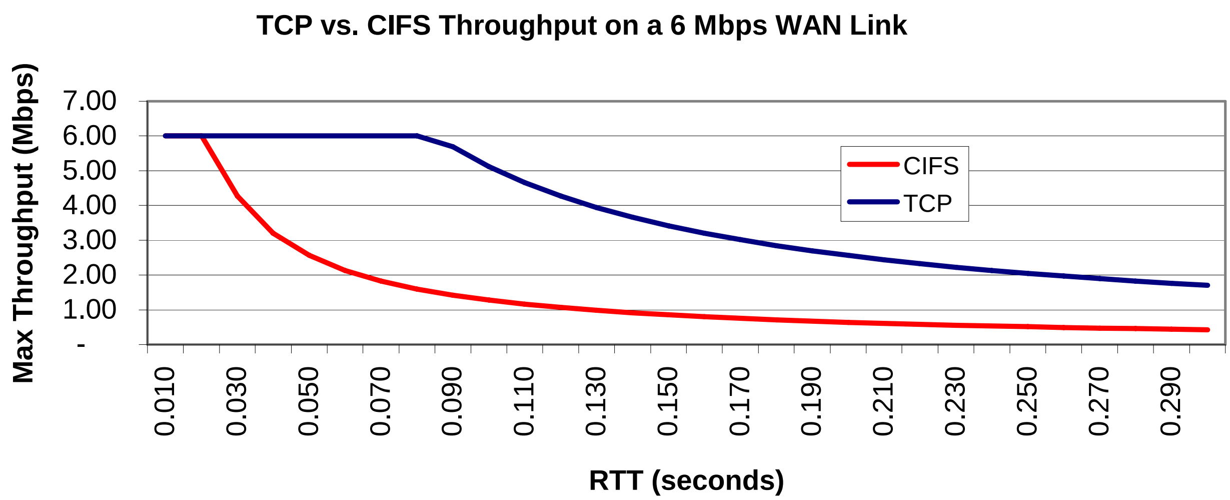

Figure 3: *Figure 3 shows a similar curve for both CIFS and TCP running across a 6 Mbps link. Again, we see a sharp decline of capacity with increasing latency – but notice how CIFS is the dominant bottleneck (because it falls off more sharply). For applications with better behavior on the WAN, TCP latency may be the dominant effect. But for Windows file sharing over CIFS, even perfectly optimized TCP and plentiful bandwidth are not enough to overcome the dire effects of high latency.*

WDS: Wide Applicability and Wide Area

Historically, the various point approaches to WAN improvement have addressed only specific bottlenecks, or have offered improvement for only a narrow set of protocols. Some examples of those approaches are summarized below. Note that WAFS technology, while excellent for alleviating bandwidth issues and application latency, does not address TCP latency. Conversely. TCP Optimization software addresses the TCP latency challenge, but does not have applicability to bandwidth and application issues.

Approach	Bandwidth?	TCP Latency?	Application Latency?
WAFS	Yes		Yes
WAN Optimization	Yes		
Web Caching	Yes		Yes
Data Reduction	Yes		
Email Caching	Yes		Yes
Block-Replication	Yes		
Dynamic Caching	Yes		Yes
TCP Optimization		Yes	
QoS			
SSL Acceleration			Yes

This table illustrates that prior to the conception of a unified WDS approach, no single point technology can address a broad spectrum of protocols challenges. WDS combines four distinct threads that have been present in these constituent approaches, but always as separate elements:

- High-performance, disk-based data reduction
- TCP optimization
- Application- or protocol-specific optimizations
- Caching or local servers

By categorical definition, WDS is multi-protocol, multi-configuration, multi-application, and targets multiple bottlenecks simultaneously. The table below illustrates the components that together comprise a WDS solution.

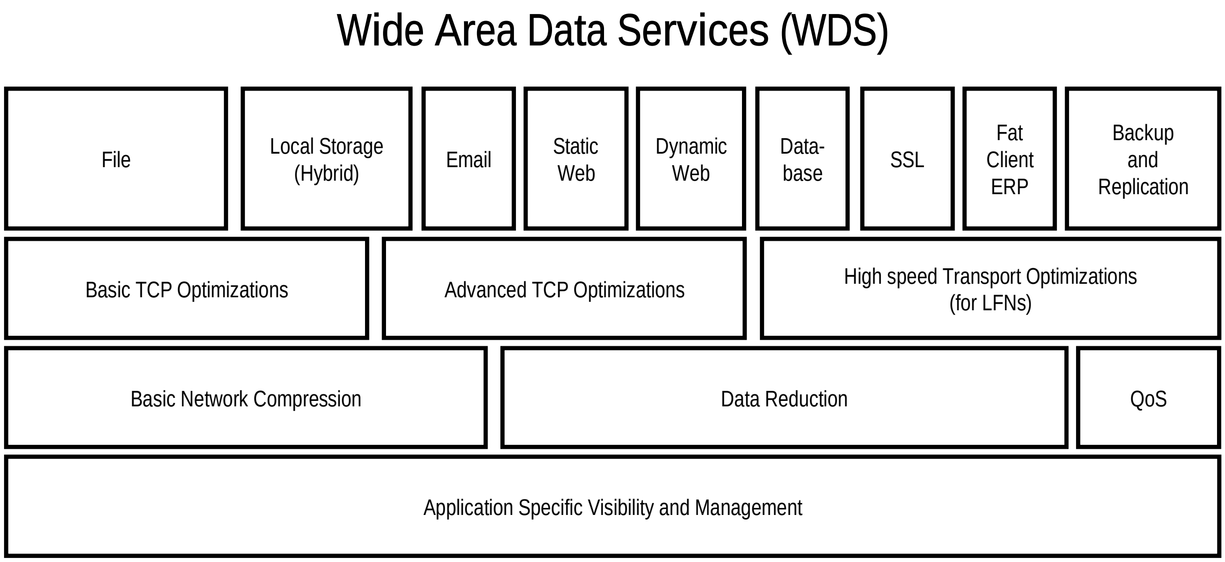

Riverbed: Pioneering WDS

For the past three years, Taneja Group has been tracking Riverbed Technology, a pioneer in the field of wide area optimization. The company was founded with the goal of solving the problem of distributed performance for all applications used on WANs. The solution had to be application-independent, transparent to clients and servers, and most importantly had to overcome the issues of high network latency inherent in WANs.

From a categorical definition standpoint, it became clear to us that Riverbed was definitely more than a Wide Area File Services (WAFS) vendor, and more than a WAN network optimization software provider. We came to the conclusion that Riverbed architected a new approach that collapsed historically disparate solutions together into one. In short, the company had pioneered the evolution of Wide Area Data Services.

Riverbed's WDS Architecture

Riverbed's Steelhead appliance architecture has several key elements that differentiate it from other approaches.

Disk-Based System

Riverbed's Steelhead solutions are built on a disk-based architecture – a bold choice for a general TCP-processing system in 2002, when Riverbed began development. Using a disk to store network traffic affords a vast capability to go *back in time* to find old repeated data patterns, even when the data in question last traversed the network days or even months earlier. Devices using only RAM are easily overrun by typical traffic levels and file sizes, leading to much lower performance. With Riverbed's dramatic success, other vendors are also rushing to add disks; but Taneja Group believes that Riverbed offers the most mature implementation available today.

Application-Independent Foundation

Unlike a caching approach, Steelhead appliances are built on two key application-independent pieces of technology: Scalable Data Referencing (SDR) and Virtual Window Expansion (VWE) which remove all redundant TCP traffic and reduce TCP round trips, respectively. The benefit of this approach is that any application running on TCP sees a significant reduction in WAN traffic and an increase in throughput. Unlike a compression appliance, Riverbed's TCP optimization accounts for the effect of high latency on TCP-based applications, which when combined with SDR, can have a dramatic improvement on applications like Lotus Notes, FTP, backup and replication traffic as well as on web-based applications.

Application Specific Latency Optimizations

On top of the application-independent foundation (SDR + VWE), Riverbed has built a set of application-specific optimizations, including elements specialized for HTTP and FTP. The most important application-specific optimizations are latency optimizations called Transaction Prediction. The first Steelhead appliances came with Transaction Prediction modules for CIFS (Windows) and MAPI (Exchange); additional Transaction Prediction modules are planned for other widely-used protocols. These Transaction Prediction modules offer incremental optimizations for those applications beyond those provided by SDR and VWE.

87 Elm Street, Suite 900   Hopkinton, MA 01748   Tel: 508-435-5040   Fax: 508-435-1530   www.tanejagroup.com

WDS in Action

WDS can be applied to many areas affecting enterprise networks today. Because WANs are such an integral part of a distributed organization's infrastructure, they affect many critical business processes and in some cases stand as a real obstacle to the plans and goals of the organization. Here are a few of the key usage models that WDS solution such as Riverbed's can provide:

Application Acceleration

Many business processes are dependent on applications deployed on WANs. The application can be as simple as sharing Windows files, or as complex as a custom-built application. In either case, if a WAN is in the middle, it's almost guaranteed that the application won't work as well as it was intended. Riverbed's Steelhead appliances can accelerate many applications, like Windows, Exchange, FTP, backup and others by as much as 100 times. With LAN-like performance, users can work together no matter which office they are located in.

Remote Office Data Backup

Backing up remote office servers requires the transfer of what normally wouldn't be too much data. But, when it needs to be sent over a low bandwidth, high latency WAN link, the throughput drops dramatically as we've seen, and what was a quick job can easily stretch out to take longer than the available backup window. For this reason, most IT managers rely on local tape auto-loaders or other backup schemes in the branch office. Of course, local backup is fraught with poor execution, equipment failure, and operational difficulty. With a WDS solution, backups can be completed in a fraction of the time they currently take, which lets you take a different approach to protecting your company's data.

Data Replication

Whether data needs to be replicated to support data replication plans, or to mirror data so it's available for users around the world, data replication is critical. WDS solutions can accelerate data replication processes by a factor of 10 or more. A WDS appliance will remove all the redundant traffic from the WAN, and optimize TCP. The two together make a huge difference in the time required to complete a replication task.

Site Consolidation

WDS appliances allow successful consolidation of file servers, email servers, NAS, and local tape backup. Deployments can start by consolidating only a single type of infrastructure with incremental consolidation of other infrastructure thereafter. Alternatively, it is possible to consolidate everything and achieve the *serverless office.*

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

In-sourcing

Distributed organizations doing knowledge work often find that they have spare capacity in some locations and insufficient capacity in others. By enabling more flexible sharing of information and applications across geographic boundaries, WDS make it possible to use idle internal resources to assist in situations that might otherwise require additional local consultants or other temporary help. This in-sourcing saves money and helps lower the volatility of work life by spreading work more evenly across locations.

Bandwidth Optimization

Sometimes the goal is as simple as avoiding a WAN upgrade, and an effective WDS appliance can help with that too. Riverbed's experience indicates a reduction in WAN traffic by 60% to 98%, which means that an existing WAN can often support many more users, new applications like VOIP can be rolled out, and an expensive WAN upgrade can be delayed or avoided.

Taneja Group Opinion

Taneja Group believes that WDS appliances like Riverbed Steelhead will become indispensable tools for distributed organizations over the course of the next 24 months. The early market success and extremely high growth that Riverbed currently enjoys indicates that the company's pioneering work in the WDS category has resonated deeply across a wide range of enterprise customers and usage cases; a prerequisite to eventual market ubiquity for any emerging technology. Riverbed is such compelling technology because it provides a single solution that can reduce WAN traffic, ensure high application performance, enable successful site consolidation projects and ensure effective data protection. There is a wide range of customer value that can be generated from such a robust technology base. To date, we have found no other product offering in the market the presents such a comprehensive set of optimizations for the enterprise customer. For companies exploring answers to the range of optimization challenges highlighted in this paper, Taneja Group endorses an evaluation of Riverbed Technology's Steelhead, without reservation.

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com