



## Optimizing Data Replication for Disaster Recovery: Using F5 WANJet to Accelerate EMC SRDF®

**Overview** This white paper addresses the performance needs of data replication for disaster recovery, and the common barriers to achieving success. It also addresses the dramatic performance gains that can be achieved by combining the F5 WANJet application acceleration solution with the EMC Symmetrix Remote Data Facility® (SRDF®) family of remote storage replication solutions. Finally, this paper details how the F5 WANJet is specifically designed to optimize SRDF traffic and how storage managers can achieve 5 to 10x performance improvements in SRDF replication, significantly reducing bandwidth requirements and lowering IT costs.

### **Challenge** Factors That Affect Disaster Recovery Success

Disaster Recovery (DR) plans are becoming a key part of a company's overall IT planning process to ensure continuous availability of the company's critical infrastructure at all times. A major component of these plans involves protecting business-critical data through backups and data replication.

Such replication and backup processes may occur between data centers, branch and home offices, or primary and backup sites. A successful business continuity/DR plan has two key components at its core: a solid replication product to manage replication processes, and an effective and efficient Wide Area Network (WAN) that enables those processes to be accomplished successfully.

Two of the critical metrics used in measuring the success of a disaster recovery plan are recovery point objectives (RPO) and recovery time objectives (RTO). These two metrics measure the amount of data lost during a disaster and the time required to restore to normal operations.

IT managers must counterbalance the lowest RTO and RPO possible with factors such as:

- Increasing data storage requirements from increased usage and regulatory archival requirements
- Limited bandwidth between primary and backup locations
- The expense of adding additional bandwidth between the DR locations
- Variable factors that can affect the performance of the DR solution over the WAN (e.g. WAN latency, reliability and service options).

According to a recent analyst survey<sup>1</sup> these factors were significant barriers to improving RPO and RTO objectives to companies around the world (Figure 1).

---

<sup>1</sup> "The Impact Of The WAN On Disaster Recovery Capabilities", February 2007, Forrester Consulting. Available at [http://www.f5.com/reports/Forrester\\_Consulting.pdf](http://www.f5.com/reports/Forrester_Consulting.pdf)



Please rate the impact of the following to your ability to improve your time to recovery and to limit data loss (On a scale of 1 to 4, where 1 equals "no impact" and 4 equals "very strong impact")

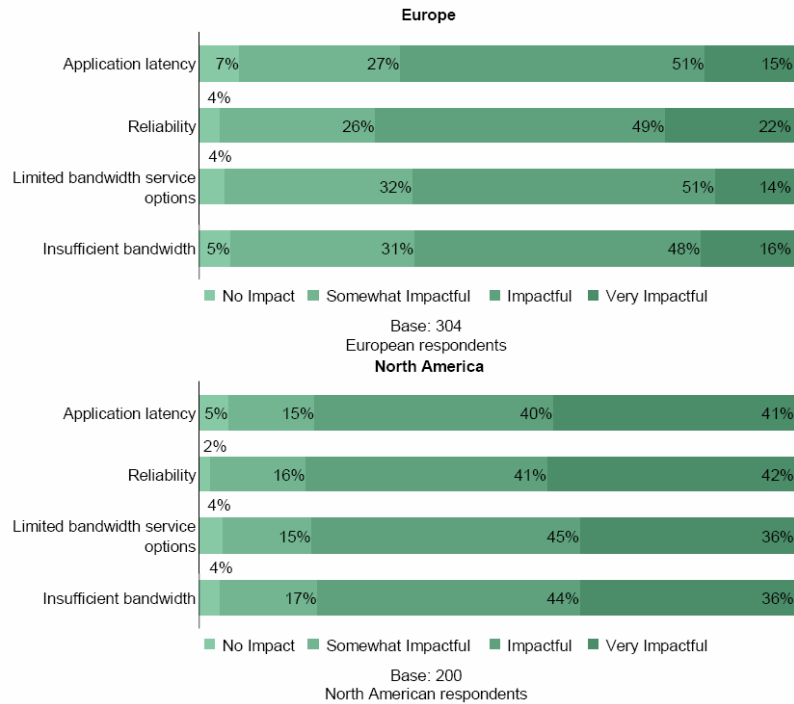



Figure 1: WAN impact on RPO and RTO objectives

One of the most common barriers to the effective deployment of any high-performance data replication solution is the performance of the solution over the WAN between the DR sites. Storage teams, when sizing the bandwidth requirements, often find that their initial sizing estimates are insufficient to meet the performance requirements of a DR solution.

In practice, true WAN performance is rarely given much thought until the organization ramps up their production replication system and realizes that the WAN bandwidth they have does not provide the expected throughput. Suddenly, the RPOs and RTOs they expected to meet are no longer realistic. According to this same survey, 25% of North American enterprises and 26% of European enterprises reported that the cost of bandwidth represented between 20% and 30% of the total cost of data replication (see figure 2).



What percent of the total cost of your replication or remote backup solution can you attribute to the cost of bandwidth?



Base: North America 200, Europe 304  
(percentages may not total 100 because of rounding)

Source: Disaster recovery and data replication study conducted by Forrester Consulting and commissioned by F5 Networks, January 2007

**Figure 2: How much does bandwidth contribute to the cost of Replication?**

WANs have several inherent characteristics that are the source of missed expectations within replication scenarios:

- Latency (caused by limits to the speed of light over distance and the number of network hops between the DR sites)
- Packet Loss (caused by signal degradation over the network medium, oversaturated network links, corrupted packets rejected in-transit, or faulty networking hardware)
- Network congestion (excess of data on the network slows overall transmission speeds)
- Actual bandwidth does not match expected bandwidth (often due to a combination of the factors listed above)
- Expense (large pipes can incur significant monthly leasing charges)

Unfortunately, these factors often cripple what was originally a good backup/DR plan. Moreover, when the DR application shares the WAN links with non-replication traffic, file transfers, and even possibly other migration or recovery activities, the RPOs and RTOs that were previously attainable can become completely unobtainable.

One common fix attempted by storage teams is to replicate only the most critical data and hence reduce the amount of data replicated. The other option frequently exercised is to increase the amount of bandwidth leased. Neither option is optimal because they both have significant drawbacks such as increased business risk or increased budget demands.

**Solution** F5's WANJet® appliance uses compression and acceleration technologies to dramatically improve the speed of application traffic over WANs. WANJet accelerates many varieties of application traffic including replication traffic, file transfers, e-mail, client-server applications, etc.

One of the data replication applications that WANJet is able to provide acceleration benefits to is the EMC Symmetrix Remote Data Facility® (SRDF®) family of remote storage replication solutions. SRDF is the most widely deployed suite of remote storage replication solutions on the market, and is installed in tens of thousands of demanding environments worldwide.

SRDF family base products include:

- SRDF/Synchronous (SRDF/S) enables high-performance, host-independent, real-time synchronous remote replication from one Symmetrix to one or more Symmetrix systems.
- SRDF/Asynchronous (SRDF/A) enables high-performance extended distance asynchronous replication using a Delta Set architecture for reduced bandwidth requirements and no host performance impact.
- SRDF/Data Mobility (SRDF/DM) enables rapid transfer of data from source volumes to remote volumes anywhere in the world, permitting information to be shared, content to be distributed, and information to be consolidated for centralized parallel processing activities.

SRDF does an excellent job at meeting the sophisticated replication needs of organizations worldwide, but like all replication solutions, it's performance depends heavily on the WAN. The benefits of using WANJet with SRDF include:

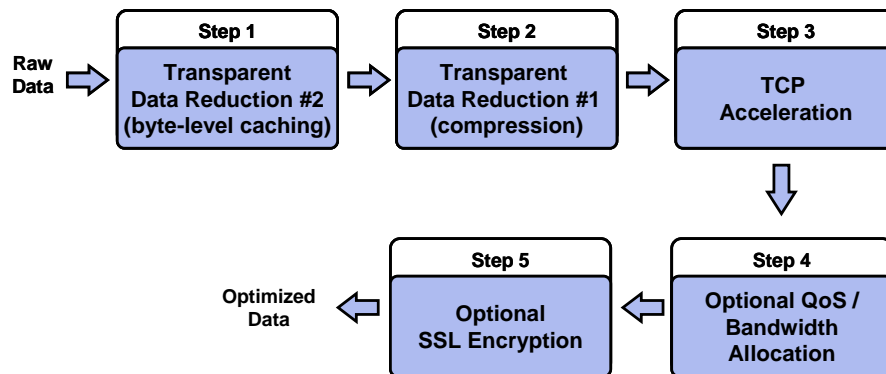
- Fewer dropped SRDF groups
- Lowered RPOs and RTOs
- Less susceptibility to uncontrollable network congestion peaks, whether from within the storage infrastructure, or from outside
- Less susceptibility to latency variations

Ultimately, that boils down to less business risk and lower cost of insuring against those risks. Organizations can meet their SLA (i.e. RPO/RTO) targets more easily, and within budget. That means less risk to equity investors as well. Imagine investing in a company whose order management system at any given point in time is at risk of going down for over two hours, and losing the last thirty minutes of orders placed. This is a significant financial risk to the company and thus to the investors. In a worst case scenario, it might go down at peak Christmas ordering time for instance. Reducing that risk (for example, going down for no more than five minutes, and losing no more than thirty seconds of orders placed), is a valuable risk reduction for everyone involved.

There are numerous specific technical advantages of using WANJet to accelerate all modes of SRDF traffic, and they fall into three main categories:

**1) Reducing RPOs and RTOs without upgrading bandwidth or replication infrastructure**

- a) WANJet is able to accelerate all modes of SRDF traffic, from 5-10 times faster, depending on the type of data being replicated. It accomplishes this through a six step process that includes byte-level caching, compression, TCP acceleration, bandwidth allocation and encryption.



- b) **Enabling the network to adapt dynamically to network congestion levels.** All of the processes within WANJet are adaptive to network conditions. For example, in the case of

compression, when it makes sense to compress heavily, WANJet will do so using a series of advanced algorithms custom-built for this purpose. If network conditions change such that it is no longer optimal to compress as heavily, WANJet will back off the compression levels automatically. This ensures that overall throughput is maximized for the application.

- c) **Guaranteeing bandwidth for important and critical replication traffic over less important traffic.** WANJet has the ability to guarantee a minimum amount of bandwidth to be available to certain high priority applications over other applications. For example, the administrator could set 80% of bandwidth to be made available to SRDF if and when SRDF needs it. Otherwise that bandwidth would be made available to any other applications that need it.
- d) **Prioritizing SRDF traffic over non-SRDF traffic.** WANJet has the ability to prioritize traffic from certain applications over traffic from other applications at all times. This is configurable by the administrator, however many replication deployments call for assigning high priority to the replication application, and low priority to non-replication types of traffic.
- e) **Mitigates the effects of latency.** Adding WANJet to an SRDF deployment makes it less susceptible to unwanted variations in WAN performance due to the combination of technologies listed above. This is achieved because:
  - i) Less bandwidth is being used so the WAN is not on the verge of being over-utilized.
  - ii) Traffic is allocated a certain amount of bandwidth depending on its importance, so in the event that WAN performance degrades, a tradeoff occurs whereby the less important traffic suffers the effects more than the important traffic.
  - iii) Important traffic is prioritized by setting higher Type-of-Service (ToS) bits, meaning the routers downstream help important traffic overtake less important traffic.
- f) **Providing more control of WAN resources allocated to Storage or DR needs.** Another significant advantage to the storage team is the ability to gain control over WAN resources which are allocated to it. Without WANJet, most storage departments have no control whatsoever over the WAN, and how it impacts data replication applications. This is especially true if the WAN is shared with non-replication traffic. WANJet however provides an easy and intelligent way for the storage team to control how, when and where those valuable WAN resources get used, and they are able to ensure that mission-critical applications like replication don't suffer due to non-mission-critical applications.

## 2) Reducing the ongoing cost of meeting RPO's and RTO's

- a) **Using less bandwidth to replicate the same or more amounts of data.** Clearly the ability to reduce by a factor of 5-10X the bandwidth needed for SRDF replication can result in significant cost savings through downsizing, or through delaying the need to upgrade WAN connections by years.
- b) **Reducing the tangible and intangible costs associated with troubleshooting.** Customers who have deployed WANJet with large-scale replication solutions continually point to the hard and soft cost savings of not having to continually troubleshoot the system when it is relying on a WAN link that has reached near maximal usage.
- c) **Provides a comprehensive view of WAN performance metrics and bottlenecks.** Because WANJet sits at the edge and aggregates all forms of traffic going over the WAN, it is able to provide valuable data that helps identify and eliminate performance problems and bottlenecks.



### 3) Securing replication traffic through SSL encryption

- a) WANJet's ability to encrypt traffic using SSL provides a valuable and cost effective way to increase security and meet certain governmental regulations that require encryption of data traveling over the WAN.

#### WANJet Performance with SRDF

SRDF replication scenarios being accelerated by WANJet often have the following requirements:

- Must be fully redundant with no single point of failure
- Must achieve maximal throughput on the existing connection
- Must prioritize SRDF over non-SRDF traffic
- Must guarantee a minimum of X% of bandwidth to SRDF traffic
- Must achieve a minimum of Y% acceleration

Based on our experience, the typical SRDF deployment without any, or with minimal compression in place, can be expected to see benefits in the range of 5 to 10 times faster performance over the WAN. In addition, each WANJet 500 unit can be expected to deliver approximately 400 Mbps of LAN-side throughput (which means approximately 40-80 Mbps of output on the WAN-side).

Multiple WANJets can be load balanced to achieve failover redundancy, or simply higher overall throughput of the system. Using F5's BIG-IP LTM application delivery networking technology in conjunction with WANJet yields n+1 scalability which is limited only by the size of the WAN connection, which could be as large as an OC-192 (10 Gbps).

#### Proven Solution

The WANJet and SRDF have been successfully tested and deployed for large SRDF installations. In one case, the customer was replicating with SRDF/A between two data centers over 1,000 miles apart. They had a very small connection (18Mbps) but required a larger connection (90Mbps). This lack of bandwidth was causing a number of problems with consistent achievement of their RPO and RTO. Installation and configuration took under a week, and because of WANJet they were able to achieve a consistent production acceleration rate of 5-times faster. Therefore they avoided the need for any upgrades in bandwidth. While each customer deployment is unique, these results are not at all uncommon.

WANJet is a safe solution for SRDF because it is transparent to the Symmetrix devices. It acts as an inline accelerator and does not alter or redirect SRDF traffic in any way. In addition, in the event of a physical failure such as power, the WANJet fails-to-wire, which means all traffic will continue to pass through unaffected.

Acceleration performance varies of course, and some of the factors that affect this result include:

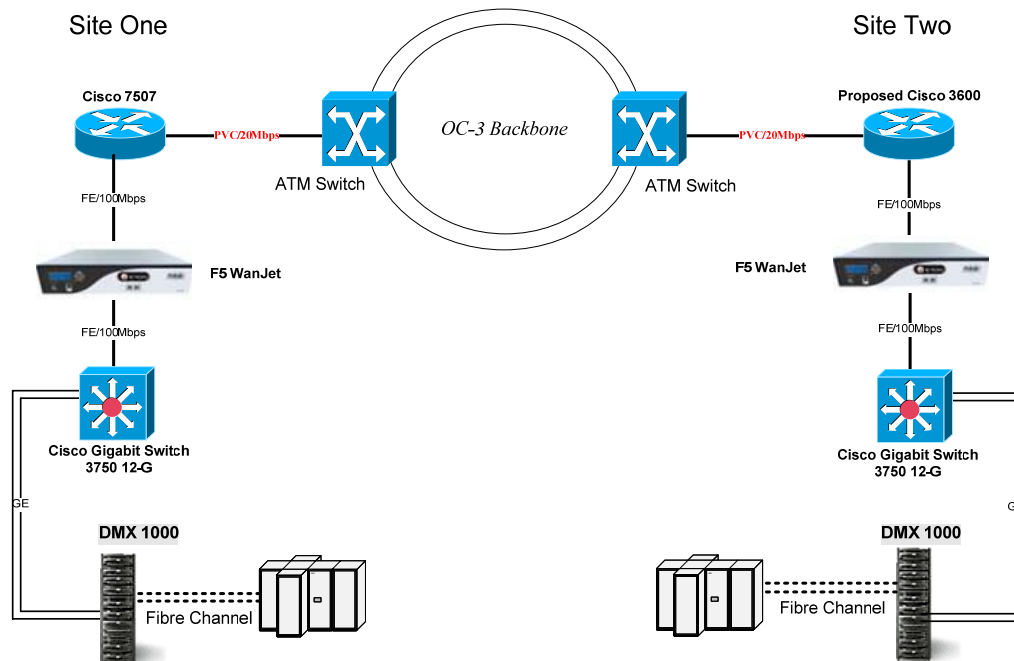
- The amount of redundant data traversing the WAN. More redundancy yields higher compression rates.
- Natural "Compress-ability" of the data (e.g. text is easily compressible, images are typically not)
- Traffic mix over the WAN links (this requires WANJet to begin enforcing bandwidth guarantees which can significantly improve performance of the important traffic, at the expense of the less important traffic)
- Traffic volume and link utilization (Congestion on the WAN links are also affected by the change in traffic volume over the course of a day. Peak load times during which a replication process previously ground to a halt can now be prevented using bandwidth allocation).

**Overview of Test Scenario**

The network configuration (see diagram below) consisted of a OC-3 (155 Mbps) Backbone with 20Mbps PVC dedicated for the SRDF traffic on the backbone. The customer needed to find a way to reduce the 90Mbps connection requirement to fit into the smaller connection allocation of less than 20Mbps. Testing results has shown up to an average of 8:1 performance improvements, all that was needed was a 5:1 ratio which the WANJets easily and consistently achieved.

The two sites were configured with a WANJet 500 at each location in a symmetrical configuration (inline) in between the Cisco 3750 and the Cisco 7507 WAN Router for each location. A standard WANJet optimization configuration was applied for the SRDF traffic with TCP Optimization, TDR-1 and TDR-2 turned on.

This was a simple deployment. More sophisticated deployments can be configured for high availability and redundancy.



**Conclusion** The combination of F5 Networks’ WANJet acceleration appliances with EMC SRDF software data protection and recovery solutions offer significant performance gains to the storage team managing replication. The combination delivers cost savings and improved RPOs and RTOs for the customer’s business. The end result is reduced risk and lower costs.

**About F5** F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5’s extensible architecture intelligently integrates application optimization, protects the application and the network, and delivers application reliability—all on one universal platform. Over 10,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to [www.f5.com](http://www.f5.com).

