

How to Improve Disaster Recovery for the Enterprise: Advanced Replication Powered by WAN Optimization

January 2010



Today's distributed and dynamic enterprises rely increasingly on 24x7 access to a growing set of mission-critical business applications and sensitive data. These applications and data are more distributed than ever: they can reside in corporate datacenters, remote offices, and/or on user computers. In addition, overall data volumes are growing rapidly in every industry segment, and widespread virtualization means that servers and data are more mobile than ever before. Moreover, IT operations teams are struggling with flat or shrinking budgets in a tough economy. These combined challenges make disaster recovery (DR) planning more difficult than it has been in the past, but they also make it more important than ever.

In order to meet current demands for application and data availability, successful enterprises are increasingly relying on the wide-area network (WAN) as a storage transport resource for DR. This enables DR operations to be centralized—reducing redundancy and lowering overhead—and to leverage innovative disk-based backup and replication technologies offered by the leading storage vendors. Decentralized, tape-based DR strategies are simply too costly and labor-intensive. In practice, they fail to meet the recovery time and recovery point objectives (RTOs/RPOs) demanded by companies facing increasingly stringent customer service and regulatory requirements.

In this profile we examine the business and technology trends that complicate and increase the cost of enterprise-wide DR planning, and we summarize the proven benefits of disk-to-disk backup and replication technologies. We then dive deeper, and explore the critical role WAN optimization plays in unlocking DR efficiencies when deployed along with these data protection solutions. WAN optimization enables the enterprise to do more with its *current* network capacity—more frequent and faster backups and replication, plus faster recovery—while leveraging *new* capacity quickly and efficiently. We conclude that a WAN optimization solution, when combined with advanced replication technologies, delivers remarkable flexibility, performance, and cost benefits for multi-datacenter enterprise disaster recovery.

The Limitations of Traditional Disaster Recovery

Historically, disaster recovery has been designed to match traditional IT boundaries—physical servers, storage arrays, network devices, applications, etc.—and

primarily based on *over-provisioning* of resources. Most servers and data stores are backed up locally to tape, if possible, requiring local IT staff to manage backup software, schedules, tape libraries, and off-site archiving. When failure occurs, multiple, complex processes must be coordinated to

S O L U T I O N P R O F I L E

separately recover and reconfigure servers and data sets, often in multiple locations. As a result, recovery times are often too long and unpredictable.

Distributed, tape-based backup also suffers from geographic limitations: it can be prohibitively expensive to ship tapes long distances, and the farther they must be shipped, the longer it will take to recover in the event of disaster. This has led many firms to situate recovery sites too close to primary sites, significantly increasing the risk of catastrophic failure due to a major event (power grid failure, hurricane, etc.) affecting a large geographic area.

These workarounds are necessary when a DR strategy relies too heavily on backup; ideally, data should be *replicated* to hot or warm remote sites populated with a mirrored set of server and storage platforms. The most efficient DR architecture is based on datacenter-to-datacenter (or site-to-site) replication, eliminating the costs and delays imposed by backup media and handling. However, replication between datacenters—either synchronous or asynchronous—just shifts the cost burden away from the backup media and onto the network, where the high cost of bandwidth remains a prime barrier to deploying widespread replication for DR.

Current Industry Trends Further Complicate Disaster Recovery

As IT operations teams struggle to deliver ever-higher levels of availability, current trends present additional challenges:

More workloads are “mission-critical”: up to 60% of applications in US-based medium-to-large enterprises are considered

business-critical today (including email, collaboration, and intranet applications and data). This evolution demands that more systems, in more locations, that rely on more timely and sensitive data, be covered by DR planning, and requires that datacenter operations teams provide tier-1 application support and data protection for a growing percentage of applications.

Data volumes are expanding rapidly: Taneja Group research shows that over half of large US enterprises had 11 terabytes or more of unstructured data—business documents, virtual machine images, email, media files, etc.—in their environments, with annual growth rates hovering around 60%. This is compounded by a 20% or more annual growth rate for transactional data, historically the bulk of data processing. With remote office staffing levels in decline, IT’s ability to track and secure these growing data sets is in jeopardy.

Worker and data mobility: by the end of 2011, IDC expects mobile workers to make up 73 percent of the total US workforce. These workers increasingly demand access to current, comprehensive, and often sensitive data, while relying on smaller IT operations teams to provide it. This creates a significant challenge for DR planning: how to make the most data available to the greatest number of mobile employees, while maintaining availability and security levels.

Increasing regulatory oversight: as a result of recent natural disasters (Hurricane Katrina and others in the Gulf states, fires in Southern California) and acts of terrorism (the shutdown of Wall Street and much of the New York/New Jersey business districts following the attacks of 9/11), industry and

S O L U T I O N P R O F I L E

governmental regulations concerning the distance between DR sites and redundancy levels continue to tighten. In addition, high-profile customer data security breaches (TJX Cos., Heartland Payment Systems) have led to calls for stricter regulatory compliance controls across industries (Sarbanes-Oxley, HIPAA, PCI DSS, and European Union Privacy laws, to name a few).

New Technologies Present Additional Challenges

In addition to these business challenges, new technologies—virtualization and cloud computing, in particular—open new avenues for DR planning, but only if properly managed.

SERVER VIRTUALIZATION AND WORKLOAD MOBILITY

Server virtualization encapsulates servers and applications into mobile workloads, which is a boon for workload protection: any server capable of running a hypervisor can be a recovery target. Also, since virtual machines (VMs) are actually just sets of files, they can be protected and replicated as data. These features can aid disaster recovery planning, if understood and managed properly.

First, virtualization is not appropriate for all workloads today or for the foreseeable future, so existing physical-server DR is still required, alongside virtualization-aware DR. In addition, server virtualization delivers two key benefits that must be kept in balance: higher utilization through consolidation, and server mobility. Consolidation significantly reduces capital costs for server hardware, but raises the chance that a physical server

failure affects multiple applications. This higher failure risk can be mitigated by workload mobility within the local area network—for example, by using a live migration technology such as VMware's vMotion.

STORAGE VIRTUALIZATION AND DATA PROTECTION

It's clear that server virtualization has driven many advances in storage technology, leading to significant enhancements in storage flexibility and performance, and new data protection technologies.

Storage virtualization enables dynamically optimized storage tiering for greater choice and efficiency when choosing a recovery target. Virtual backup targets, including virtual tape libraries, allow easy migration from tape-based to disk-based backup. Thin provisioning reduces the disk footprint of VM images for efficient storage and transport for DR operations. Array-based VM snapshot and cloning moves server protection to the array and leverages economies of scale; multiple data types (including VMs) can be protected with a smaller set of backup targets. Finally, image and file-level recovery options on the array speed VM recovery as well as recovery of files within VM images.

These enhancements are further extended for DR by advanced data replication solutions, including EMC's Symmetrix Remote Data Facility/Asynchronous (SRDF/A), Hitachi Data Systems' TrueCopy and Universal Replicator, NetApp's SnapMirror, Brocade's FCIP Gateway and Dell EqualLogic's Auto-Replication, among others.

S O L U T I O N P R O F I L E

For virtual servers, many advanced data protection features are also offered by the hypervisor platform vendors themselves. Examples include VMware's Consolidated Backup, Site Recovery Manager, Storage vMotion & vStorage APIs and Citrix's StorageLink APIs. Together with replication, these virtualization-optimized enhancements deliver flexibility and choice, but also contribute to an increase in the overall volume of data requiring protection—an important consideration.

CLOUD COMPUTING AND STORAGE ELASTICITY

The emergence of new utility-priced outsourcing options is leading many enterprises to explore cloud computing as an adjunct to existing DR strategies. Low-cost cloud-based storage offerings are maturing, and can provide an effective alternative to building dedicated recovery facilities. Cloud-based storage provides an elastic pool of trusted storage, if several key requirements are met. The provider platform must be able to scale quickly, on demand, and to large capacities. It must provide clear multi-tenancy policies to isolate one firm's data from all others, and must enable secure access to data at any time. Cloud storage providers must also provide rapid, robust data recovery, which demands a storage infrastructure with very high MTBF and MTDDL (mean-time-to-data-loss).

We expect DR to be one of the most compelling use cases for cloud-based storage, which has matured from a poorly-understood technology to a sought-after solution in roughly two years. In fact, over 40% of enterprise datacenter managers recently reported that they either have or plan to

deploy some form of cloud storage by the end of 2011. Most business-critical applications are protected by some amount of off-site backup today, but many firms require more advanced DR capabilities in order to satisfy industry compliance, risk mitigation, or business partner requirements—they struggle, however, to justify the capital and operating costs required to implement them. Digital content serving, video surveillance, and medical image archiving are some common data use cases for cloud storage which demand strong protection and security, and cloud storage vendors are quickly adding more advanced functionality for additional data types.

However cloud storage is implemented, it demands a reliable, efficient, and high-volume WAN interface. Any DR strategy which includes cloud-based storage will clearly benefit from maximizing capacity and performance of the WAN infrastructure.

The Wide-Area Network Becomes a Core Element of Disaster Recovery

Taken together, these industry trends and technology developments reveal a consistent underlying infrastructure requirement: the wide-area network (WAN), more than ever, is an *essential* storage transport resource for DR, and will increasingly play a critical role in unlocking DR operational efficiency.

However, if the WAN is to become a core DR infrastructure element, it's imperative to resolve the most challenging aspects of network performance: contention for bandwidth and protocol latency.

S O L U T I O N P R O F I L E

The Current State of Data Protection in the Enterprise

The search for solutions to optimize DR is a top priority in the enterprise, as revealed by Taneja Group research. The firms leading the way are often the same ones who have invested heavily in virtualization. Our research reveals that, today, over 50% of medium to large US enterprises rely on the

same backup processes for virtual servers as they do for physical servers. 55% still use tape as their primary backup/DR medium, but most also employ at least one additional disk-to-disk method (Figure 1). The most common of these are tier 1 SAN backup (backing up to a primary disk on another SAN) and tier 2 backup to a low-cost disk or virtual tape library (VTL) at the local site.

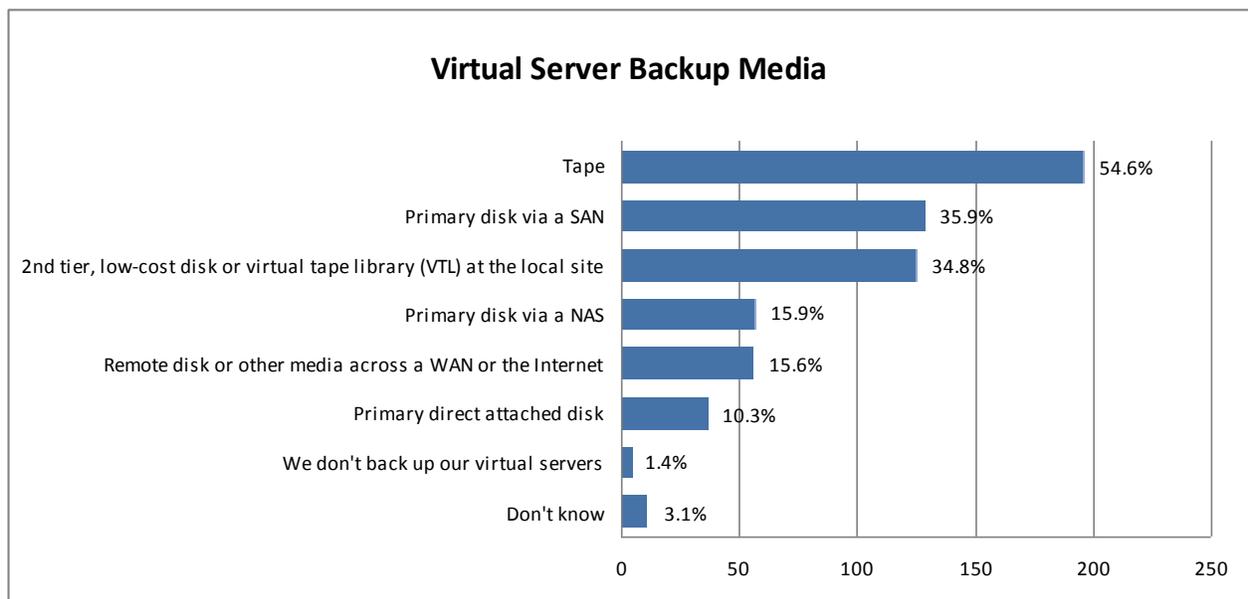


Fig. 1: Backup media used in virtual server environments (source: Taneja Group, 2009)

However, change is already well underway: the same research also reveals that virtual server data protection is seen as a major issue requiring improvement by a majority of users at every stage of virtualization.

Though most firms have not yet adopted virtual server technology for their key DR practices, one virtualization-aware DR approach is widely used: migration of virtual machines between servers, using VMware

VMotion or an equivalent. As shown in Figure 2, well over 50% of respondents are employing manual and/or automated VM migration for DR today.

Our research also verifies that replication is another popular DR mechanism, increasingly popular in the most virtualized datacenters. *Between 45 and 50% of surveyed users employ data replication in their virtual infrastructures today.*

S O L U T I O N P R O F I L E

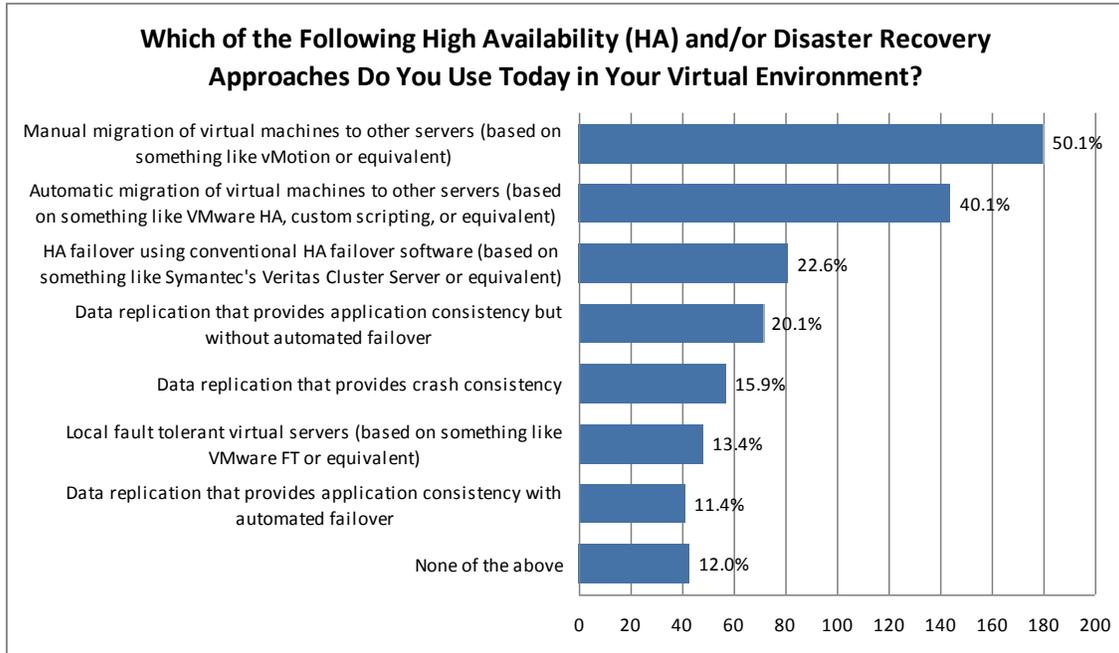


Fig. 2: HA/DR strategies deployed in virtual server environments (source: Taneja Group, 2009)

In our opinion, the key DR take-away from this data is that virtualization-savvy companies are eager to use virtualization to simplify and accelerate DR. Snapshots, clones, and on-demand conversions (physical-to-virtual and vice-versa) enable near real-time replication of virtual workloads wherever there is a hypervisor, with availability levels limited only by *how often and how quickly workloads and data can be moved between locations*.

Clearly, network speed and capacity are the primary factors limiting multi-site DR performance. WAN optimization solutions that deliver more of both on existing networks will be increasingly popular. We fully expect higher demand for WAN optimization as recognize that it enables dramatically higher levels of replication and, in turn, overall datacenter availability.

The Power of WAN Optimization for Disaster Recovery

Enterprises should be investigating solutions to optimize network capacity and utilization *today*, to ensure that WAN bandwidth and latency are not limiting factors in current or future DR planning. Unlocking existing—but untapped—WAN bandwidth early in any DR planning effort will help avoid overspending on other DR solution components (storage or tape stackers, for instance) in the future.

WHO SHOULD EXPLORE WAN OPTIMIZATION?

WAN optimization breaks down operational barriers, relieves geographic limitations, and enhances emerging storage technologies for DR. You should explore WAN optimization if any of the follow statements is true:

S O L U T I O N P R O F I L E

“I can’t reliably meet my DR service-level agreements (SLAs).”

If you only replicate a subset of critical data or workloads within datacenters and on LANs, and rely on tape backups or portable disks for disaster recovery operations between sites, your ability to predict recovery times and/or deliver adequate recovery points is likely variable and difficult to test. WAN optimization can deliver near-LAN performance over WAN links, enabling more datacenter-to-datacenter replication for tighter RPO/RTO SLAs. Also, with automated network quality of service (QoS) management, WAN optimization allows you to test your DR processes more often and test them on live networks, without affecting application traffic.

“It’s difficult to protect multiple, geographically dispersed locations.”

WAN optimization allows you to overcome the bandwidth mismatch between your LAN and WAN environment, enabling you to back up more often and to more locations: datacenter-to-datacenter, branch-office-to-datacenter, remote-worker-to-datacenter, etc. By overcoming protocol latency, WAN optimization allows you to transfer backup and recovery data sets over greater distances without performance degradation. Also, WAN optimization allows you to continuously replicate or move sensitive data held in multiple locations to a single secure location for improved compliance auditing.

“I want to centralize my backup operations to reduce redundancy.”

WAN optimization can eliminate local tape backups and the staff required to manage them. Centralizing data protection leverages the datacenter staff and skills you have in

place today and reduces dependencies on branch office contractors and other third-parties. WAN optimization benefits every datacenter discipline: the storage team can deploy advanced, bandwidth-hungry replication technologies between sites and arrays, the server team can take VM snapshots more frequently and back them up to multiple locations, and the network team can deliver higher performance returns for every dollar spent on WAN capacity.

WHAT DOES WAN OPTIMIZATION DELIVER?

WAN optimization immediately increases the ROI for your current network infrastructure, and delivers compelling operational benefits:

- **Greater Flexibility:** more backup and replication to more sites, in more configurations (1:1, 1:N, N:1, P2V, V2P, Mobile-to-Datacenter, etc.).
- **Higher Cost Efficiency:** unlocking additional capacity extracts value from what you already own, and helps delay the need to purchase additional bandwidth.
- **Lower Risk:** consolidating backup reduces complexity, redundancy, and the potential for data loss; fewer locations leads to simpler and more successful compliance audits.
- **Higher Availability:** protecting more data and workloads and taking advantage of more replication raises your overall enterprise availability levels.
- **Shorter RPO:** replicating more often—while protecting service quality for other network traffic—ensures that the most

S O L U T I O N P R O F I L E

current data possible is available—as close as possible to where it needs to be.

- **Shorter RTO:** overcoming protocol inefficiencies and reducing data on the wire allows faster backups and more frequent replication.

HOW DOES WAN OPTIMIZATION WORK?

In our view, a comprehensive WAN optimization solution should include the following essential capabilities in order to provide maximum benefit for DR:

DR Objective	Solution Capability
Increase WAN Utilization & Efficiency	Lowers network inefficiencies via TCP optimization, allowing maximum use of rated capacity for existing networks and enabling more DR data sets to be sent more often
Overcome Application Protocol Latencies	Overcomes inherent application protocol latencies to allow backup/replication to multiple datacenters and remote sites at greater distances
Reduce Data on the Wire	Lowers the amount of duplicate data sent over the WAN to shorten both backup and recovery times and enable near-synchronous replication
Optimize Performance for DR Workloads	Automatically detects and tunes network performance for bursty traffic and large datasets, providing additional optimization specifically for DR workloads
Manage Contention and Quality of Service	Brokers contention for shared WAN resources to maintain required throughput for application workloads during DR operations or test runs

Evaluating WAN Optimization Solutions for Disaster Recovery

How should you identify the best WAN optimization solution for your chosen backup and replication technologies? First, explore the relationship between all vendors: confirm that they have tested the joint solution fully, certified and documented it to both vendors’ standards, and provide published benchmarks. Once this baseline is established, look for these technology differentiators:

TCP Optimization: the solution should maximize the amount of data per payload on each TCP round-trip; remove redundant WAN data; and offer high-speed TCP options (HS-TCP, MX-TCP) to fill the network ‘pipe’ most effectively. Leading solutions deliver 60-95% reduction in WAN utilization.

Data Optimization: the solution should employ advanced data de-duplication algorithms and optimize the use of on-board processing power and memory. De-duplication should be at the byte level and

S O L U T I O N P R O F I L E

produce data reduction ratios that exceed and enhance block-based data differencing technologies provided by storage vendors. Data reduction ratios should range from at least 4:1 to 8:1 or greater.

Workload Optimization: the solution should provide DR-specific workload identification and acceleration; automatically detect DR-specific traffic (requests for large datasets, for example), intercept requests, and segment the data for maximum wire efficiency using adaptive algorithms. Optimization should yield at least 5x and up to 50x improvement in WAN-based backup and recovery times.

Quality of Service (QoS) Management: the solution should broker contention for and allocate WAN bandwidth by application

and/or port, while independently prioritizing packets based on individual latency sensitivity for both TCP and UDP traffic. Leading solutions provide innovative QoS tuning algorithms that allocate the minimum bandwidth required for each application, avoiding overcompensation.

Transparency and Availability: the solution should be essentially transparent to your existing network infrastructure from various vantage points (user, deployment, reporting, port- and address-level); it should enable centralized management of multiple systems; and it should offer advanced availability, scalability and deployment features, such as efficient use of on-board memory, solid-state drives for caching, and clustered configurations.

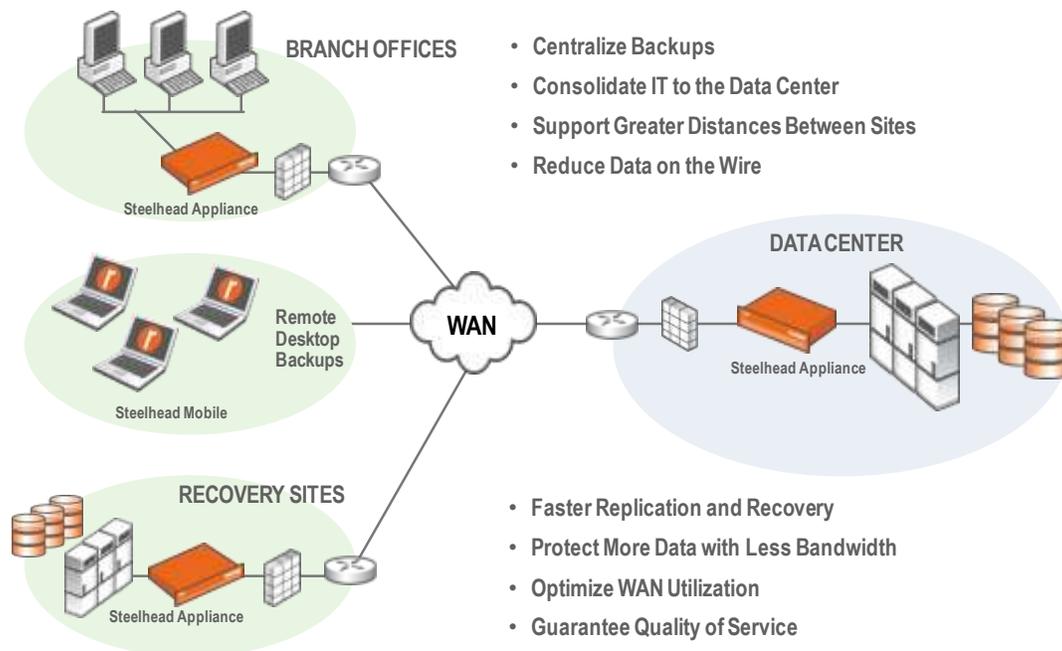


Fig. 3: WAN-Optimized Disaster Recovery Reference Architecture (based on Riverbed's Steelhead family). Appliances are deployed at branch offices, remote desktops, disaster recovery sites, and in the primary datacenter, working transparently over the existing WAN to accelerate performance and maximize throughput.



S O L U T I O N P R O F I L E

Taneja Group Opinion

The datacenter is rapidly evolving. As a result, most data protection and disaster recovery (DR) plans were developed for an IT environment that bears little resemblance to today's dynamic, highly distributed, and increasingly virtualized datacenters. Traditional DR planning assumed that remote locations had dedicated IT staff, that applications and data were generally static, that local data protection was most efficient, and that wide-area networks between locations were strictly limited and expensive.

These assumptions are being challenged by economic pressures, rapidly-growing data volumes, widespread virtualization, and more powerful storage array capabilities. It's no longer cost-effective to maintain dedicated IT staff in smaller remote offices, or to rely on redundant tape-based backups in the event of disaster. An increasingly mobile workforce and growing regulatory oversight drive ever-increasing demands for greater DR efficiency, including workload and data replication between multiple sites and across broad geographies. And, virtualization itself offers new levels of application mobility and enables innovative virtualization-aware backup and replication strategies at the storage array level. However, our capacity to replicate data and workloads is rapidly outpacing our ability to place them where needed.

A new approach to disaster recovery is required, one that addresses the key constraint on enterprise DR scalability and efficiency: limited network capacity. WAN optimization removes this "choke point" by unlocking existing network bandwidth and optimizing the WAN for the unique transport requirements of backup and recovery workloads. The proven benefits of WAN optimization include dramatic data reduction ratios and massive gains in network throughput, and these benefits are even further magnified when combined with storage replication solutions from the premier array vendors.

This powerful joint solution makes long-distance, multi-site replication feasible, simpler, and cost-effective, and in our opinion provides the clearest path to reducing enterprise dependence on expensive and error-prone tape-based disaster recovery.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are provided without warranty of any kind. All product names mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.