

Forrester Consulting

HELPING BUSINESS THRIVE ON TECHNOLOGY CHANGE

February 2007

The Impact Of The WAN On Disaster Recovery Capabilities

A commissioned study conducted by Forrester Consulting on behalf of F5 Networks

FORRESTER®

FORRESTER®

Headquarters

Forrester Research, Inc., 400 Technology Square, Cambridge, MA 02139 USA
Tel: +1 617/613-6000 • Fax: +1 617/613-5000 • www.forrester.com

TABLE OF CONTENTS

Introduction: Upgrading Disaster Recovery Capabilities	3
Study Methodology	3
The State Of Disaster Recovery Preparedness	4
Distance Between Recovery Sites.....	6
Bandwidth, Transport, And Replication Between Recovery Sites	8
How Much Data Must Enterprises Protect?	11
What Are Enterprises' Current Recovery Capabilities?	12
Remote Sites Are A Critical Risk Exposure.....	13
The Cost And Impact Of The Wan On Recovery Objectives	14
The Drive To Improve Recovery Capabilities	16
Can WAN Acceleration Help?	18
Conclusions	20
Enterprise Recommendations.....	20

© 2007, Forrester Research, Inc. All rights reserved. Forrester, Forrester Wave, WholeView 2, Technographics, and Total Economic Impact are trademarks of Forrester Research, Inc. All other trademarks are the property of their respective companies. Forrester clients may make one attributed copy or slide of each figure contained herein. Additional reproduction is strictly prohibited. For additional reproduction rights and usage information, go to www.forrester.com. Information is based on best available resources. Opinions reflect judgment at the time and are subject to change.

Introduction: Upgrading Disaster Recovery Capabilities

Due to heightened risk, fiduciary responsibility, increased competition, and regulation, upgrading disaster recovery capabilities is a top priority for enterprises. Enterprises have historically relied on offsite tape vaulting for disaster recovery, but today a growing number of enterprises use more sophisticated disaster recovery solutions such as data replication between data centers and major sites. These enterprises now face a two-pronged challenge: They must determine how to optimize the performance of the replication solutions they already have in order to achieve a better recovery time and recovery point objective, as well as extend these solutions to cover critical remote sites at a reasonable cost. In many instances, issues related to wide-area network (WAN) connectivity between sites are key limitations to achieving these goals.

The study commissioned as part of this report indicates that, on average, bandwidth accounts for 29% of the total cost of a data replication solution for disaster recovery. In addition, a majority of enterprises state that WAN connectivity issues such as application latency, network reliability, limited bandwidth service options, and the cost of bandwidth itself significantly impact their ability to improve recovery time and recovery point objectives. With the amount of data requiring remote replication growing dramatically each year, enterprises will either need to spend more on bandwidth or determine ways to optimize their existing WAN connectivity in order to improve recovery objectives.

F5 Networks commissioned this study to determine:

- The current state of disaster recovery preparedness.
- The specific challenges and limitations of data replication solutions attributed to WAN connectivity.
- The importance of improving recovery time and recovery point objectives at data centers and remote sites.
- The market drivers fueling the need for improved recovery time and recovery point objectives.
- The market awareness of WAN acceleration appliances to improve the performance of data replication solutions.

Study Methodology

In December 2006 and January 2007, Forrester Consulting conducted an online survey of 504 IT decision-makers and influencers across North America and Europe. In this survey:

- Two hundred respondents were from North America (US and Canada), and 304 were from from Europe (UK, Germany, and France).
- Forty-seven percent of respondents were from enterprises that had 1,000 to 4,999 employees, 32% had 5,000 to 19,999 employees, and 20% had 20,000 employees or more.
- Thirty-one percent of respondents were from companies with revenues of less than \$1 billion, 39% were from companies with revenues of \$1 billion to \$5 billion, 16% were from

Workforce Continuity

companies with revenues of \$5 billion to \$10 billion, and 13% were from companies with revenues greater than \$10 billion.

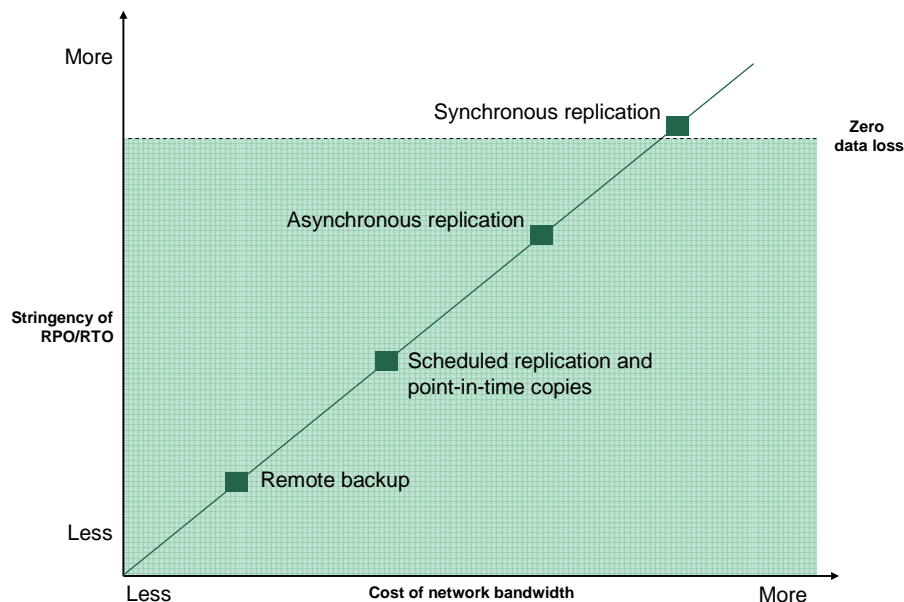
- All respondents were technology decision-makers or influencers for business continuity and disaster recovery. A majority of respondents had a title commensurate with director or above such as manager or director of IT, VP of IT, or CIO.
- All respondents had at least one backup data center (a data center or other site that acts as a failover site in the event of site failure of another data center).
- Respondents were from a broad range of industries. Telecommunication vendors were excluded from the survey.

The State Of Disaster Recovery Preparedness

To achieve the desired recovery time objective (RTO) and recovery point objective (RPO) with the greatest possible distance between sites, enterprises often design multisite disaster recovery configurations that combine synchronous, asynchronous, and even batch or scheduled replication technologies. Site separation is crucial because enterprises need to locate their recovery site far enough to escape the likeliest of local and regional threats such as natural and man-made disasters.

Synchronous replication ensures zero data loss because it does not return a write acknowledgment to the application until the data has been written to the recovery site. As a result, synchronous replication requires high bandwidth and low latency between data centers in a metro area — ideally no more than 240 kilometers (km) or 150 miles (mi) apart. Asynchronous solutions are typically deployed for long distances. With asynchronous replication solutions, the primary and secondary sites will be slightly out of sync, and there is some chance of data loss in the event of a disaster or business disruption. However, applications are not forced to wait for the remote site to confirm a write acknowledgment before processing can continue (see Figure 1).

Figure 1 The Data Replication Continuum

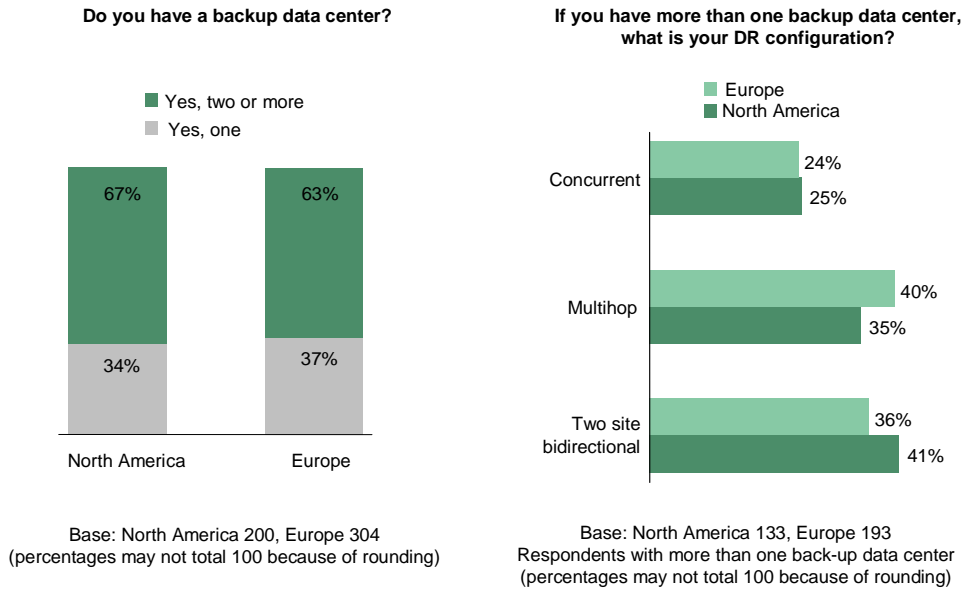


Workforce Continuity

In fact, multisite disaster recovery configurations are the norm. According to the survey, 67% of North American respondents and 63% of European enterprises have two or more backup data centers or recovery sites (see Figure 2). There are three major types of multisite disaster recovery configurations:

- **Two site bidirectional.** In this configuration, a firm replicates data between two of its data centers in both directions. This configuration is helpful for firms that want to take advantage of two existing production data centers — each data center acts as the recovery site for the other.
- **Two site multihop.** In this configuration, firms replicate data synchronously from the production data center to a bunker site that is within a campus or metro distance and then replicate data asynchronously from the bunker site to a recovery site outside the region — potentially thousands of miles away. This configuration helps firms achieve the greatest possible distance between data centers without any data loss.
- **Two site concurrent:** In this configuration, a firm replicates data synchronously from the production data center to a data center within a campus or metro distance; at the same time, data is asynchronously replicated to a recovery site outside the region. Like the multihop, a two site concurrent configuration helps achieve the greatest possible distance between data centers without any data loss.

Figure 2 Disaster Recovery Configurations



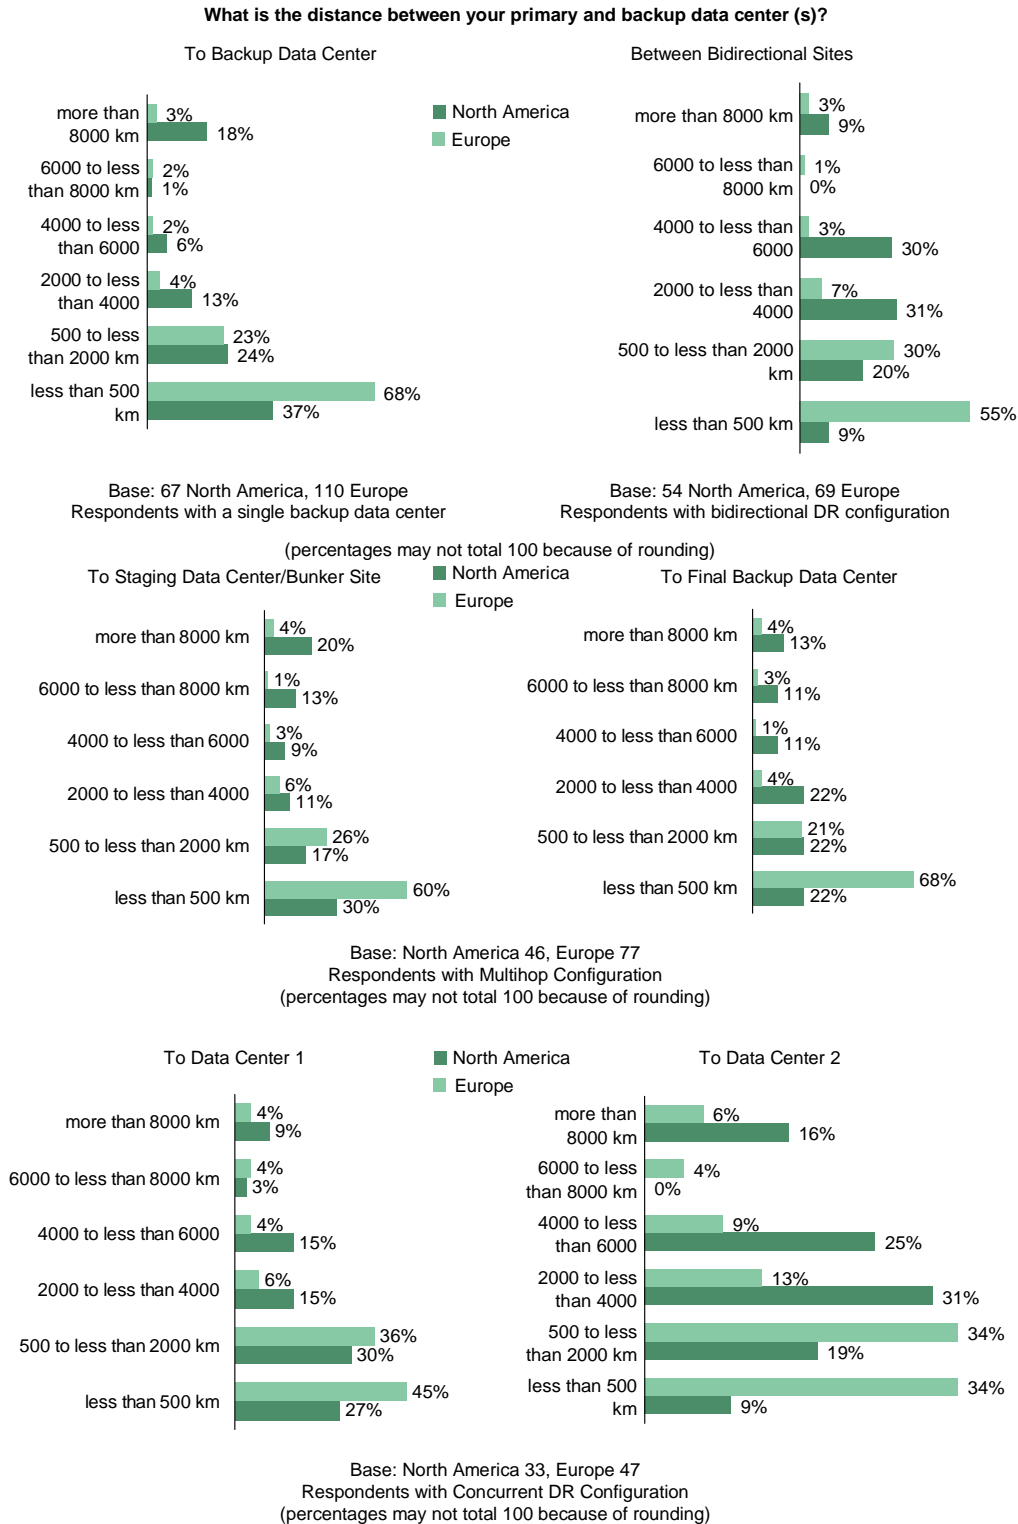
Source: Disaster recovery and data replication study conducted by Forrester Consulting and commissioned by F5 Networks, January 2007

Distance Between Recovery Sites

Enterprises with a single backup data center favor much shorter distances between their sites. Thirty-seven percent of North American respondents and 68% of European respondents said that the distance from their primary data center and their backup data center was less than 500 km (311 mi). Enterprises with bidirectional configurations seemed to favor much greater distances: 61% of North American respondents had site distances between 2000 km and 6000 km (1243 mi and 3728 mi), and 27% of European respondents had site distances between 500 km and 2000 km (311 mi and 1243 mi). A similar trend emerged among enterprises with multihop and concurrent disaster recovery configurations.

The distance between recovery sites differs significantly between North America and Europe. Overall, European enterprises favor much closer distances between sites. In almost any disaster recovery configuration, there is a significant percentage of European respondents with distances of less than 500 km (311 mi) and only a small percentage that exceed 2000 km (1243 mi) (see Figure 3).

Figure 3 Distances Between Recovery Sites



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

Bandwidth, Transport, And Replication Between Recovery Sites

The cost of bandwidth is a significant component of any disaster recovery solution that uses replication between sites. The amount of bandwidth required between sites is a function of:

1. The amount of data that must be replicated
2. The distance between sites
3. The enterprise's recovery objectives

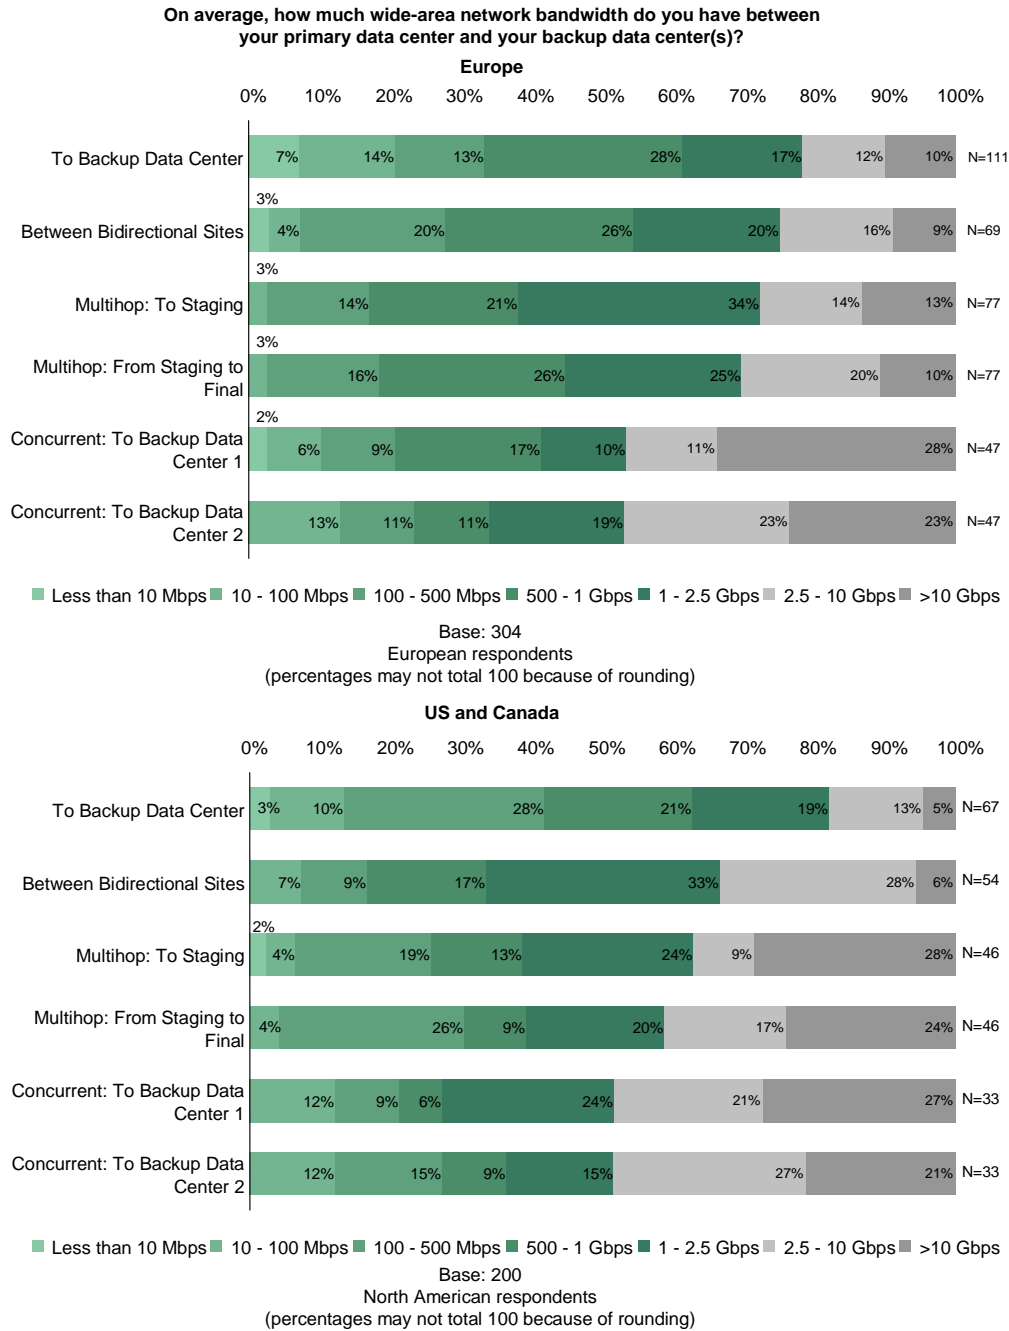
Enterprises that use multisite disaster recovery configurations are attempting to achieve the greatest possible distances between sites without sacrificing their recovery objectives. It's therefore not surprising that enterprises in both North America and Europe with multihop or concurrent disaster recovery configurations use more bandwidth between sites than enterprises with a single backup data center or bidirectional disaster recovery configuration. There were some key regional differences: North American enterprises use significantly more bandwidth in bidirectional configurations and in multihop configurations than European enterprises (see Figure 4-1).

Both North American and European enterprises rely predominantly on Ethernet as the WAN transport between sites. European enterprises also take advantage of Internet connectivity with additional security. In the past, replication between sites required high bandwidth transport options such as Fibre Channel over dedicated fiber, WDM, or SONET/SDH that only the largest enterprises could afford. However, with the advent of more affordable, yet reliable Ethernet services and advances in replication and remote backup technology, more enterprises are able to affordably link their data centers and sites and support replication (see Figure 4-2).

Fewer and fewer enterprises rely on offsite tape vaulting as their disaster recovery method and more and more rely on remote backup and replication between sites. In North America, a majority of respondents rely on synchronous and asynchronous replication technologies in most disaster recovery configurations. North America's heavy use of synchronous and asynchronous replication explains why the region requires more bandwidth between its sites. In Europe, there is a strong use of synchronous and asynchronous replication, but remote backup over the WAN is also a very popular method of copying data between sites (see Figure 4-3).

Workforce Continuity

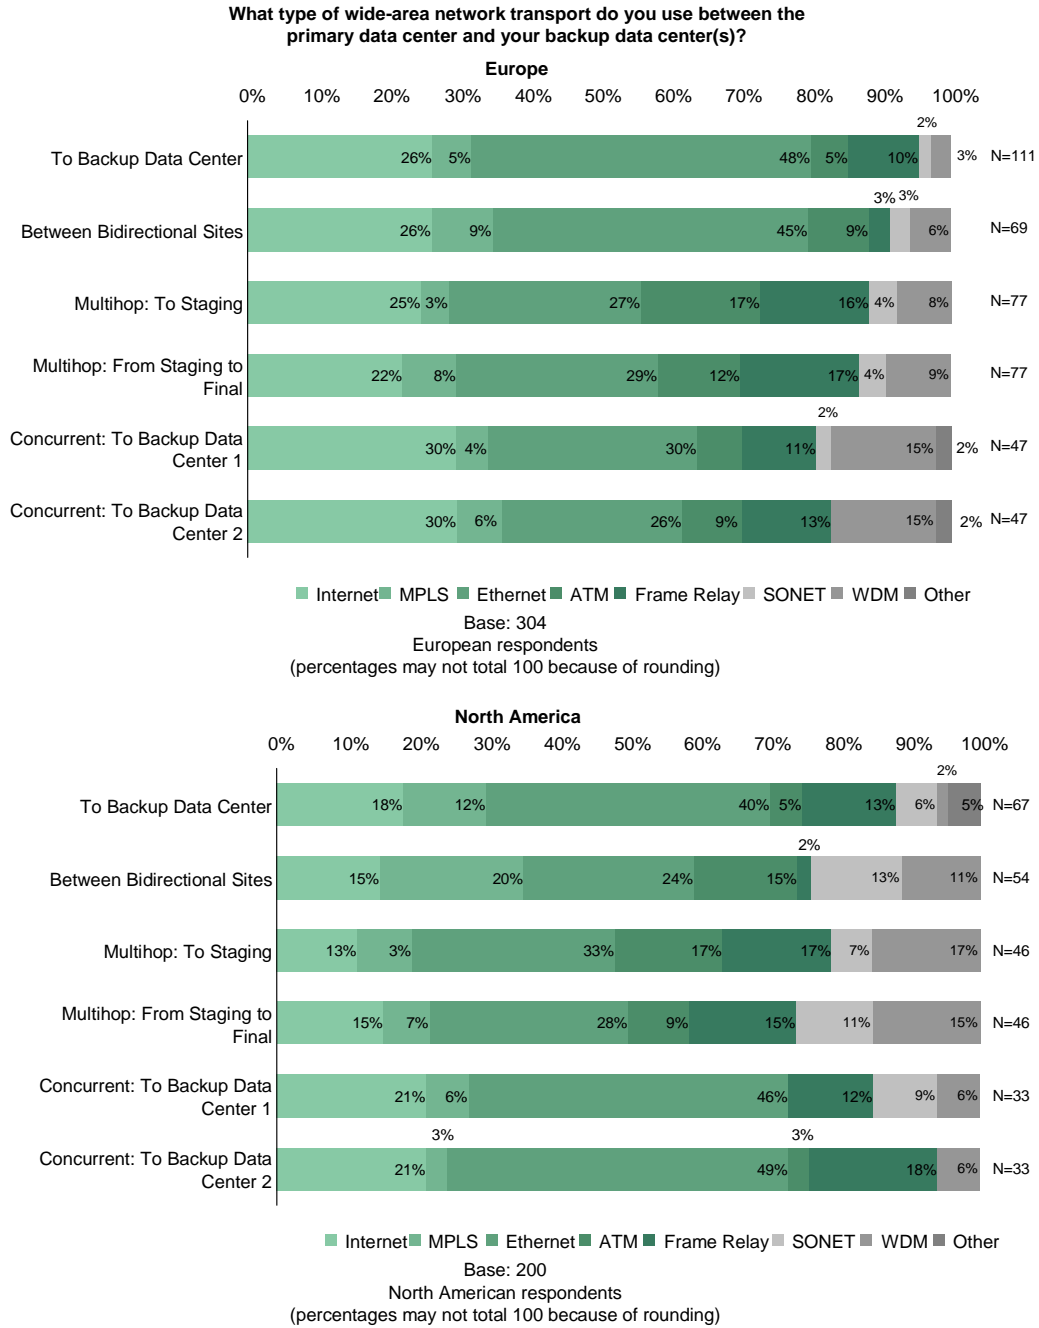
Figure 4-1 Bandwidth And Transport Between Recovery Sites



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

Workforce Continuity

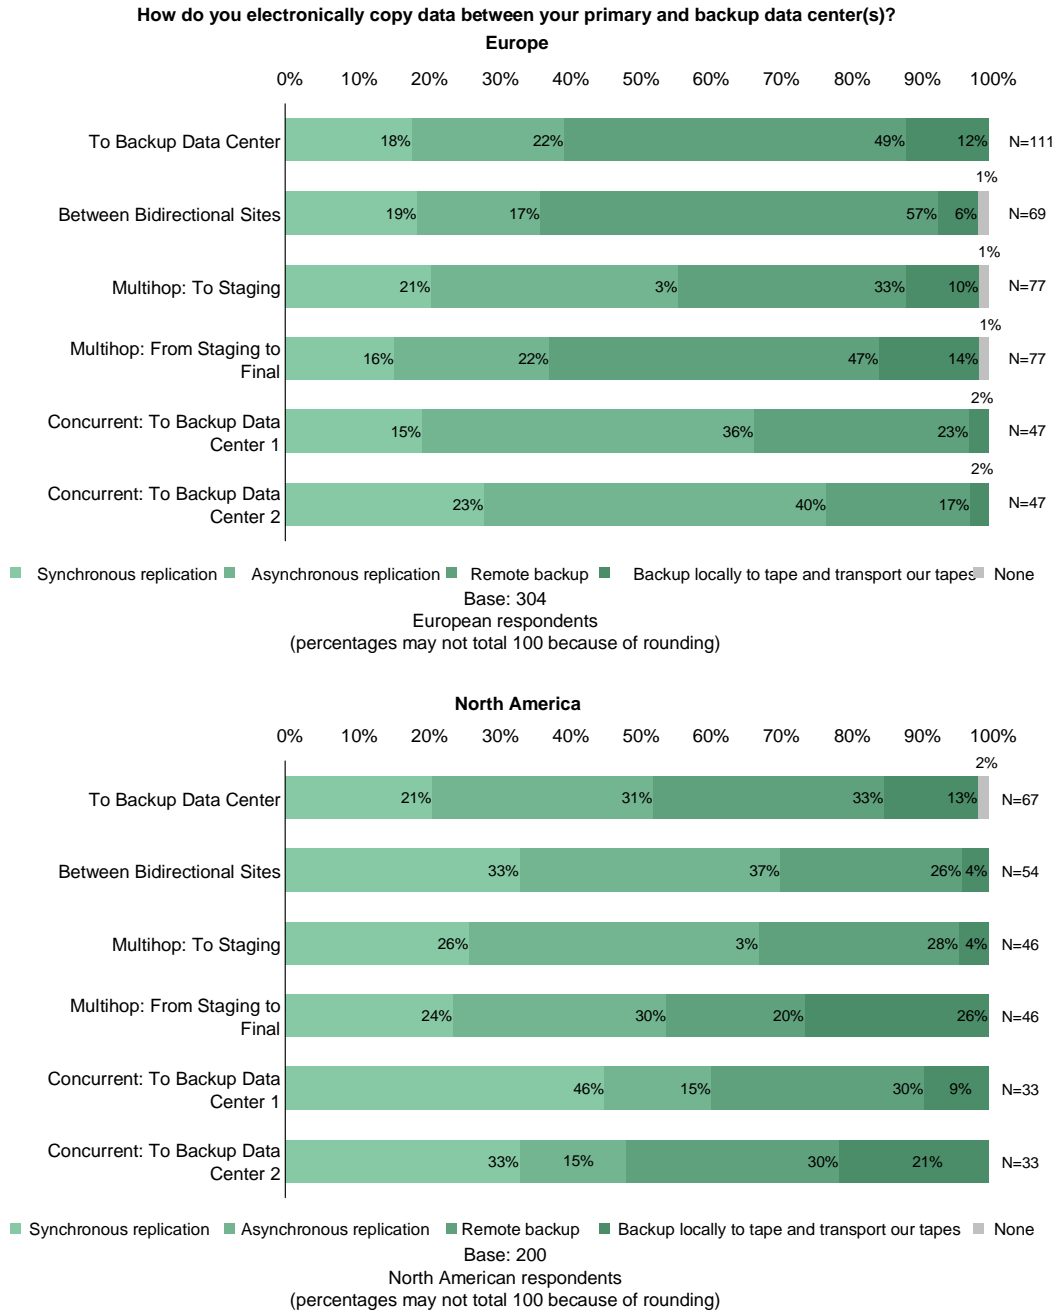
Figure 4-2 Ethernet Is The Dominant Transport Between Backup Sites



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

Workforce Continuity

Figure 4-3 Replication Between Recovery Sites



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


How Much Data Must Enterprises Protect?

Multiterabyte (TB) storage environments have become commonplace. In fact, many enterprises in financial services, life sciences, and managed services measure their storage capacity in petabytes (1,000 TBs). While many enterprises still have less than 25 TB of storage capacity (33% of North American enterprises and 37% of European enterprises), a growing percentage of enterprises have significant amounts of storage capacity to contend with. Twenty-six percent of North American respondents and 24% of European respondents have storage environments of 500 TBs or more (see Figure 5).

The more data that must be replicated, the more bandwidth is required. Because storage capacities grow dramatically each year, enterprises that lack enough bandwidth today to support their disaster recovery solutions will only find this problem compounding over time.

Figure 5 How Much Data Must Enterprises Protect?

Please estimate the amount of data (in terabytes) that must be protected at the primary site



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

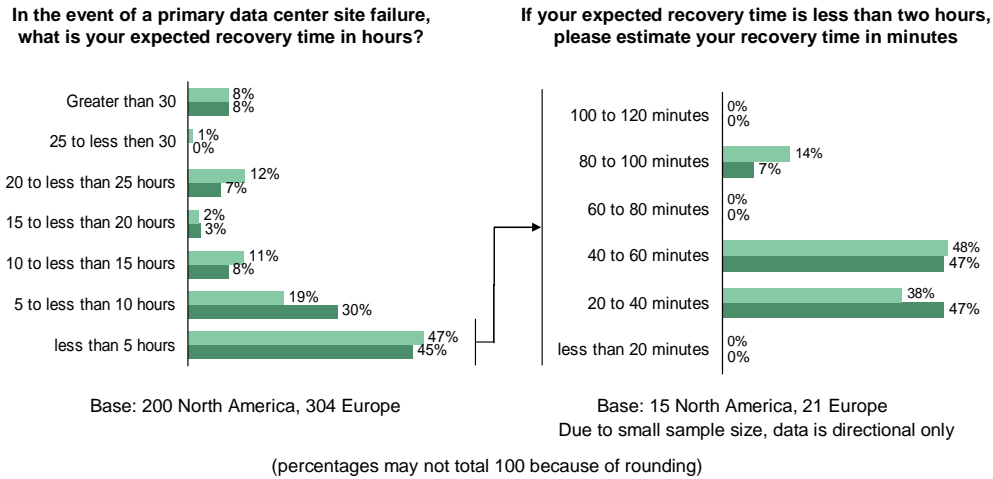
What Are Enterprises' Current Recovery Capabilities?

Enterprises that invest in alternate recovery sites and use data replication or remote backup technologies can measure their RTO and RPO in hours (as opposed to days with offsite tape vaulting).

- In the event of a primary data center site failure, 45% of North American respondents and 47% of European respondents could recover in five hours or less. However, only a small percentage of respondents, 8% of North American respondents and 7% of European respondents, could measure their recovery time in 120 minutes or less (see Figure 6-1).
- In the event of primary data center site failure, 55% of North American respondents and 59% of European respondents would lose five hours of data or less. Twenty-eight percent of North American respondents and 27% of European respondents could measure their data loss in 120 minutes or less (see Figure 6-2).


The challenge for enterprises now is to determine how they can optimize their existing disaster recovery solutions to the point that their recovery time and recovery point capabilities are measured in minutes, not hours.

Figure 6-1 Recovery Time Capabilities



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

Figure 6-2 Recovery Point Capabilities



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

Remote Sites Are A Critical Risk Exposure

Remote sites can be as small as local branches and sales offices and as large as regional data centers and headquarters. According to the survey, all North American respondents and 98% of European respondents have at least one remote site. Twenty-nine percent of North American respondents have between one and 20 remote sites; another 27% of North American respondents have between 21 and 40 remote sites. Thirty-seven percent of European respondents have between one and 20 remote sites; another 16% of North American respondents have between 21 and 40 remote sites.


A majority of enterprises report that many of their remote sites are not protected from any local disaster through remote replication or remote backup to a central data center or other facility. In fact, only 16% of North American respondents and 15% of European respondents report that 80% or more of their remote sites are protected with centralized data replication or remote backup

technologies. This means that enterprises are either relying on local backups and offsite tape vaulting for disaster recovery, or there is no disaster recovery solution in place at all. Either situation is not good for enterprises. Offsite tape vaulting is error-prone and cumbersome, and it exposes the enterprise to the risk of lost or stolen tapes. Having no solution at all means that critical data could be lost in the event of a local business disruption (see Figure 7). As enterprises become increasingly distributed, more and more critical corporate data may actually be located in remote sites, not necessarily at corporate headquarters or large data centers. This is particularly true in industries that execute on local projects, such as consulting firms, law firms, construction, etc.

Centralizing remote site data protection and extending disaster recovery to remote sites is critical. However, enterprises will need to determine how to extend replication or remote backup from a remote site to a central facility over limited bandwidth. According to our survey, most enterprises have less than one gigabyte per second (Gbps) of bandwidth to their remote site and the transport is Ethernet, MPLS, or Internet connectivity.

Figure 7 How Many Remote Sites Are Unprotected?

What percentage of remote sites replicate or backup data to data center or central facility?



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


The Impact Of The WAN On Recovery Objectives

A comprehensive disaster recovery solution requires investment in multiple hardened recovery sites, duplicate IT assets such as servers and storage arrays, networking equipment at these sites, replication software, and the necessary bandwidth between sites. The cost of bandwidth is often a significant component of the cost of a disaster recovery solution that relies on data replication between sites. According to the study, 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 8).

The amount of bandwidth and the type of network transport selected (e.g., wavelength, SONET, Ethernet, or IP) are keys to achieving desired recovery objectives, limiting the impact of latency to business applications, and increasing the distance between sites. According to our survey, WAN connectivity issues such as latency, reliability, limited service options, and limited bandwidth each made a significant impact in improving recovery objectives. A majority of both North American and European enterprises rated issues related to WAN connectivity as either having an “impact” or “very strong impact” on their ability to improve recovery objectives. A majority of North American and European enterprises also “agreed or strongly agreed” that their current bandwidth prevented them from increasing the geographic distance between sites, extending replication protection to more business applications within the organization, and extending remote replication protection to remote sites. The majority of respondents from both regions also “agreed or strongly agreed” that it would be important to improve RTO and RPO without increasing bandwidth (see Figure 9 and Figure 10).

Workforce Continuity


Figure 8 How Much Does Bandwidth Contribute To The Cost Of Replication?



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

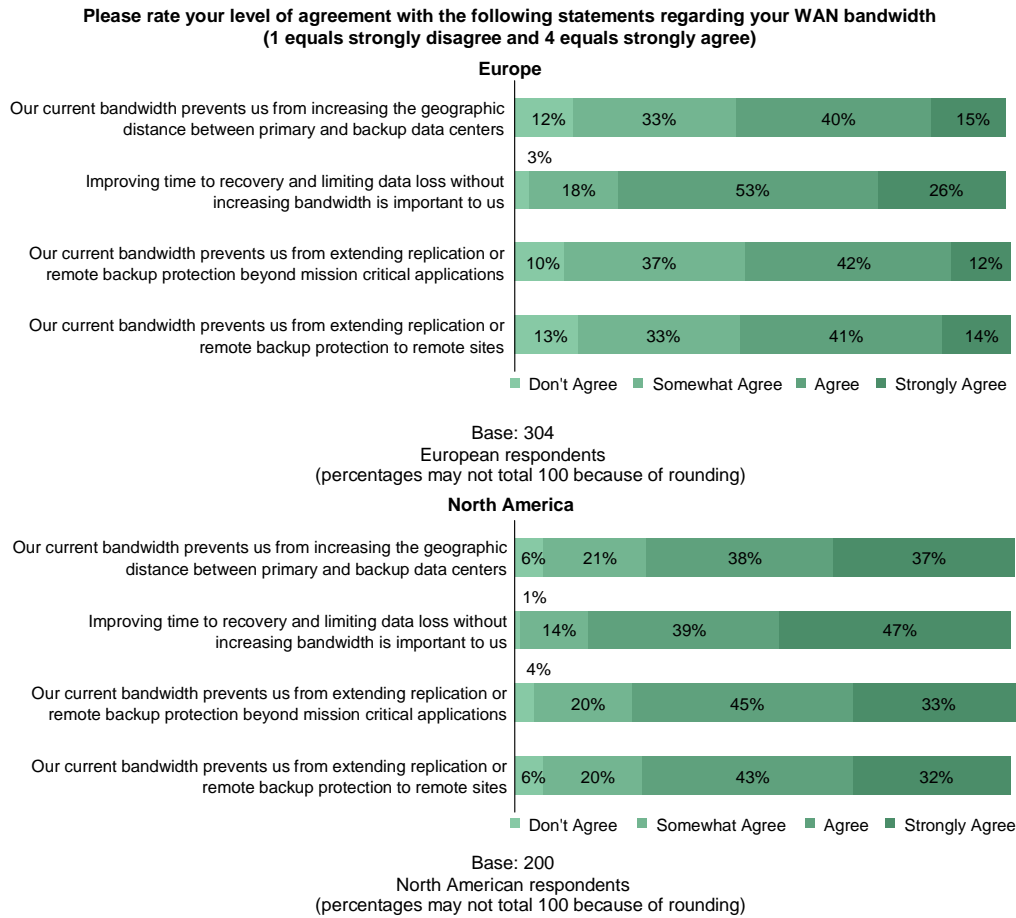
Figure 9 The Impact of the WAN On Recovery Time And Recovery Point Objectives

**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")**



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

Figure 10 The WAN Affects More Than Just Recovery Objectives



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


The Drive To Improve Recovery Capabilities

Given that most enterprises can recover business operations in less than five hours and can also limit data loss to less than five hours, how important is it to further improve recovery capabilities? Fifty percent of North American enterprises and 37% of European enterprises ranked having the ability to improve time to recover and to limit data loss as “very critical.” What drives this need for improved recovery? In both North America and Europe, more enterprises ranked “the cost of downtime” as one of their top three drivers. The cost of downtime is more than just the lost revenue directly attributed to the business disruption, it is also the cost associated with permanent customer loss and the ability of competitors to gain market share (see Figure 11).

Aside from the cost of downtime, the drivers fueling the need to improve recovery capabilities differ between North American and European enterprises. North American enterprises next ranked “increased risk” and “fiduciary responsibility” as their top drivers, while European enterprises next ranked “competitiveness” and “regulatory/legal” concerns as their top drivers (see Figure 12).


Figure 11 Criticality Of Improving Recovery Objectives

How important is it to improve your time to recovery at the backup data center and to limit data loss?
 (Please answer on a scale of 1-4, where 1 is not critical and 4 is very critical)



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


How important is it to improve recovery times and limit data loss for remote sites?
 (Please answer on a scale of 1-4, where 1 is not critical and 4 is very critical)



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 12 What's Fueling The Need To Improve Recovery Objectives?



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

Can WAN Acceleration Help?


As Figure 10 shows, 86% of North American enterprises and 79% of European enterprises “agreed or strongly agreed” that improving time to recover and limiting data loss without increasing bandwidth is important. If enterprises are desperate to improve recovery objectives but don’t want to increase bandwidth, what are their alternatives? One alternative is to consider deploying a WAN acceleration technology. WAN acceleration technologies help to improve the throughput and mitigate the latency of existing networks through such techniques as compression, data reduction, and transport protocol optimization. WAN acceleration technologies are offered as standalone appliances, built in as part of routers or switches, software, or in some cases, even as a service. Enterprises seeking to improve the performance of their existing WAN infrastructure investments typically turn to standalone appliances or software. Often, the cost of deploying a WAN acceleration appliance at each end of the link is less expensive than the cost of increasing bandwidth. For enterprises with existing replication solutions present between data centers, these appliances can help in the following ways:

- **Mitigate application latency.** In situations where the WAN is the bottleneck, WAN acceleration can improve the performance of synchronous and asynchronous replication solutions, which in turn can mitigate application latency and its performance impact on applications.
- **Support the replication of more data with existing bandwidth.** Traditional business applications like ERP, SCM, and CRM, as well as messaging and collaboration applications, such as email, continue to grow steadily each year. These are often the very applications that are supported with remote replication. WAN acceleration can help enterprises support the continued replication of these applications with existing bandwidth.
- **Extend replication to other applications.** Due to the cost of replication, most enterprises are very selective about which applications they replicate and which they don't — usually it's only mission critical applications. Today, there is no longer a one-to-one relationship between a business process and an application. Business processes rely on multiple applications, and to restore the entire process means the coordinated recovery of multiple applications. Applications that were once deemed only business critical, as opposed to mission critical, also require replication to another site.

For enterprises that want to use replication or remote backup between remote sites with limited bandwidth to the corporate data center, WAN acceleration software is an appealing approach because the software can be installed on existing servers; the enterprise does not have to invest in a standalone appliance at each site.

WAN acceleration not only helps to improve the RTO and RPO of individual applications or sites, but when taken together, it improves the overall disaster recovery preparedness of the entire enterprise. Not surprisingly, the awareness among the BC/DR decision-makers and influencers surveyed for this research report of WAN acceleration appliances as a means by which to improve the performance of replication solutions was high in North America and Europe. 82% of North American respondents were "aware or very aware" of the benefits of the technology, while 70% of European respondents were "aware or very aware" of the technology.

Figure 13 Awareness Of WAN Acceleration Appliances



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

Conclusions

The objectives of this study were to develop an understanding of the current state of disaster recovery preparedness, the impact of the WAN on recovery capabilities, the enterprise demand for improved recovery, the drivers behind this demand, and the market awareness of WAN acceleration appliances in improving the performance of data replication or remote backup technologies used in disaster recovery. The findings of the study reveal that:

- With about 45% of enterprises able to recover from a primary data center site failure within five hours, disaster recovery preparedness is generally good, albeit with some room for improvement. Remote office locations — which the majority of responding enterprises said do not have remote backup — represent the greatest area for improvement. Enterprises have made significant efforts to deploy multisite disaster recovery configurations that balance the need for geographic distance to escape regional threats with the ability to rapidly recover operations and limit data loss to five hours or less. However, only about 8% of enterprises can measure their time to recovery and data loss in minutes.
- WAN connectivity has a significant impact on the ability of enterprises to improve their time to recover and limit data loss. A majority of enterprises agree that issues such as latency, reliability, availability of service options, and the cost of bandwidth itself all negatively affect recovery capabilities.
- While overall disaster recovery preparedness is good, enterprises are still driven to improve recovery time and recovery point capabilities; 82% of enterprises surveyed rated this as either “critical or very critical.” Bandwidth is a significant constraint in these efforts: 82% of responding enterprises “agreed or strongly agreed” that improving recovery objectives without increasing bandwidth is important to them. Furthermore, 63% “agreed or strongly agreed” that their current bandwidth prevents them from extending replication or remote backup protection beyond mission critical applications.
- Cost of downtime remains the main driving force behind the need to improve RTO and RPO. However, regional difference exists between North America and Europe on secondary drivers. After the need to improve RTO and RPO, North America next ranks “increased risk” and “fiduciary responsibility” to key stakeholders. In contrast, European firms’ second highest drivers are remaining competitive 24x7 in a global environment and meeting regulatory requirements.
- Awareness of WAN acceleration is high among BC/DR decisions-makers and influencers, with 75% of respondents claiming to be “aware or very aware” of the technologies. This is a good indicator that the cost of the WAN and its impact on replication has come to the attention of enterprises.

Enterprise Recommendations

- Undertake a business impact analysis to map the dependencies between critical business operations and the people, resources, applications, and physical IT assets they rely upon. In today’s environments, business processes rely on multiple integrated applications, databases, storage systems, etc. In order to restore a business process in the event of a disruption, you must be certain you’re replicating and coordinating the recovery of all the dependent applications. Defining RTO and RPO application by application and selectively replicating data might mean you can only partially restore a business process.
- Protect your remote sites. Data loss at remote sites is currently a huge risk exposure for most enterprises. Leverage existing investments in existing data centers and recovery sites

Workforce Continuity

to offer consolidated backup solutions from remote sites to a central facility. Consolidation will help ensure that backups are run regularly and successfully, improve manageability (central administrators will have visibility into remote site data protection and potentially the ability to remotely configure and manage protection), and enable the remote office to recover from a localized disaster.

- Before investing in additional bandwidth to support remote sites, improve the performance of existing replication technologies or expand replication to other applications; consider WAN acceleration offerings. When evaluating WAN acceleration appliances, focus on the vendors that have made the time and investment to test the interoperability of their appliance with independent software vendors, storage vendors, and storage networking vendors. Also look for case studies and/or customer references that prove its capabilities and intended benefits.