

Forrester Consulting

MAKING LEADERS SUCCESSFUL EVERY DAY

A Commissioned Study Prepared For Network Appliance by Forrester Consulting

January 2008

The Continued Shift To Disk-Based Data Protection Solutions

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: Enterprises Shift To Disk As The First Line Of Protection.....	3
Study Methodology	3
Data Protection Challenges.....	4
Increasing Capacities.....	4
Backup And Recovery Budgets.....	5
Disk-Based Data Protection Continuum.....	6
Market Adoption: Disk-Based Data Protection.....	7
Current Satisfaction	8
Disk As Primary Backup Repository	9
Market Demand: Disk-Based Data Protection	10
Drivers Of Adoption.....	11
Barriers To Adoption.....	12
Deduplication: The Key To More Affordable Disk Data Protection	13
The Future Of Disk And Tape In The Enterprise	14
Optimizing Disk And Tape With Virtual Tape Libraries	15
Intelligent Disk	16
Integrating Data Protection For The Data Center And Remote Sites	18
Conclusions	20
Enterprise Recommendations	21
Endnotes.....	22

© 2008, Forrester Research, Inc. All rights reserved. Forrester, Forrester Wave, RoleView, Technographics, TechRadar, 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: Enterprises Shift To Disk As The First Line Of Protection

In the summer of 2006, Network Appliance commissioned a study to determine the adoption of and satisfaction with disk-based data protection technologies, the specific IT challenges fueling adoption of disk-based data protection, and the interplay of disk and tape in data protection. The term “disk-based data protection” encompasses all data protection solutions that rely on disk as the primary backup target or as the means by which to deliver advanced backup services. The study’s findings can be summarized as follows:

- Adoption of disk-based data protection solutions is already strong, and the market has already graduated from an early-adopter market phase to an expanding market phase.
- Firms that have adopted disk are “satisfied” to “very satisfied” with the performance of these technologies. This implies that disk-based solutions do achieve their advertised benefits (faster backups and restores).
- The key adoption driver of disk-based data protection adoption is increasing data capacities. Due to the avalanche of data, firms are unable to meet current backup windows or maintain current service levels with their existing tape-based solutions alone.
- Firms will continue to invest in tape but will shift more existing budget and net new investment disk. On average, enterprises predicted a 40/60 disk to tape ratio within two years.

Network Appliance has refreshed this study to determine how the adoption and satisfaction of disk-based data protection technologies has evolved from 2006 to 2007 but also to determine the following:

- The demand for advanced data protection functionality such as the ability to offer multiple recovery points, granular object recovery, etc.
- The role of capacity optimization techniques such as deduplication in the acceleration of disk-based data protection adoption.
- The role of application or data type in the preference for specific disk-based data protection technologies.
- The current state of consolidated remote office data protection.

Study Methodology

In the winter of 2007, Forrester Consulting conducted a telephone survey of 400 IT decision-makers and influencers across North America, Europe, and Asia. In this survey:

- 50% of respondents were from North America, 25% were from Europe, and 25% were from from Asia Pacific.¹

The Continued Shift To Disk-Based Data Protection Solutions

- 11% of respondents were from companies that had 500 to 999 employees, 41% had 1,000 to 4,999 employees, 29% had 5,000 to 19,999 employees, and 20% had 20,000 employees or more.²
- 33% of respondents were from companies with revenues of \$250 million to \$750 million, 33% were from companies with revenues of \$750 million to \$1.5 billion, and 35% were from companies with revenues greater than \$1.5 billion.³
- All respondents were knowledgeable of their company's backup and recovery environments and were also decision-makers in the planning and purchasing of backup and recovery solutions.
- Respondents were from a variety of industries.

The 2007 study methodology is consistent with the 2006 study methodology.

Data Protection Challenges

Increasing Capacities

By a wide margin, respondents rated completing backups within defined windows as the most important result to achieve at their company, followed by meeting service-level agreements for moving tape copies off-site. Respondents also ranked these two results as the most important in 2006. Despite the increased adoption of disk, firms still struggle with the essential objectives of backup, completing backups, and transporting the data off-site for disaster recovery purposes. This is mostly likely due to the fact that capacities have increased dramatically in one year. From 2006 to 2007, the average capacity increased from approximately 38 terabytes (TB) to 88 TB (see Table 1). This capacity increase far exceeded respondents' expected growth rates in 2006. As the volume of data that firms must protect increases, all focus is on completing backups within the allotted window and firms have little additional time and resources to determine how they can optimize the backup environment with advanced techniques and technologies.

Table 1: Capacities By Geographic Region

Region	2006 capacity*	2007 capacity	2007 expected 12-month growth
All	38	88	19%
North America	49	121	21%
Europe	21	49	17%
Asia Pacific	33	64	19%

Capacity measured in terabytes (TBs)

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Asia Pacific saw its capacity increase by 94%. In 2006, respondents from Asia Pacific had only anticipated a capacity increase of 18%. Clearly, these respondents had underestimated capacity growth in their environment. This year, respondents expect a 19% increase. It's not unreasonable to assume that this is also a gross underestimation of future capacity requirements.

The Continued Shift To Disk-Based Data Protection Solutions

In the current study, Forrester asked respondents to break down their capacity by file, email, and database applications. File and email capacity and expected growth exceed that of database applications (see Table 2). This is consistent across geographic region and company size. This isn't surprising, as email is the primary method of corporate communication and the backbone of a unified communication strategy, and firms consider it to be mission-critical; there is an avalanche of files due to the ongoing digitization of content such as multimedia files (graphics, photos, voice, video, etc.) and the creation of traditional office documents (word processing, spreadsheets, presentations, etc.).

Table 2: Capacities By Data Type

Data type	Capacity	2007 expected 12-month growth
Files	42	20%
Email	11	16%
Database	35	22%

Capacity measured in terabytes (TBs)

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

File, email, and databases have different protection requirements. For example, it's critical to have granular object-level recovery with files and email. Firms must be able to restore an individual file or email without restoring an entire file system or email store. Databases, on the other hand, do not typically require granular object-level recovery — restoring an individual tablespace, table, or row is dangerous and can corrupt the referential integrity of the database if not performed by someone without expert knowledge of the database structure. In the 2007 study, firms did not indicate a preference for a particular technology to protect a given data set, most likely because firms have not advanced their data protection operations to the point where they can take advantage of the distinctive capabilities of a particular technology (see Figure 1).

Figure 1: Backup Priorities

"On a scale of 1 to 4 where 1 = not important and 4 = very important, how important is it to achieve the following results at your company?"

	2006*	2007
Complete all backups within the defined windows	3.5	3.6
Meet service-level agreements for moving tape copies off-site	2.9	2.9
Store weeks of backup data online for fast restores	2.7	2.9
Reduce the IT staff required to manage backups	2.4	2.8
Encrypt data for physical tapes stored off-site	2.6	2.6
Reduce the annual spending on backup and recovery	2.7	2.5

Base: All respondents; 2006 N = 401; 2007 N = 400

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
 Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Backup And Recovery Budgets

When asked to estimate their backup and recovery, 47% of respondents reported that their worldwide annual budget for backup and recovery was less than \$500,000, and 21% reported that it

The Continued Shift To Disk-Based Data Protection Solutions

was between \$500,000 and \$1.5 million. This is an increase from last year, when more than 58% of respondents reported their backup and recovery budget was \$500,000 or less. In the previous survey, 39% of respondents planned to increase spending by approximately 19%. In this year's survey:

- 40% of respondents planned to increase spending during the next 12 months. The average increase in spending was 25%.
- 49% of respondents planned spending to remain the same during the next 12 months.
- Only 8% of respondents planned to decrease spending on backup and recovery during the next 12 months.

It's clear that backup and recovery budgets are growing, but the results are still a relatively small budget when you consider that enterprises must protect more than 88 TB of data. This could be attributed to how many firms don't have a budget specifically set aside for backup and recovery. IT often funds backup and recovery from multiple operational and capital budgets and can be difficult to estimate the total spend. Capacity optimization techniques such as deduplication will be critical to overcome the 3x capital acquisition cost of disk versus tape so that firms can expand the use of disk in data protection and to store more data on disk for longer periods of time. Expanded use of disk is critical to firms' ability to complete backups during allotted windows and increase the speed of recovery. Generally speaking, backup to a disk target is 30% faster than tape, and recovery from disk can be as much as 90% faster than tape.

Disk-Based Data Protection Continuum

Firms have multiple disk-based technologies to choose from to meet their recovery time and recovery point objectives. The following is a quick synopsis of the technologies covered in this study:

- **Backup to conventional disk subsystem.** In this solution, backups are streamed to a conventional disk subsystem — usually filled with low-cost SATA drives. The key benefit with this setup is that it can provide faster backups and restores than tape, though the disk subsystem itself has no particular features that optimize it as a backup target or a standalone backup solution.
- **Backup to a virtual tape library (VTL).** VTLs use tape emulation to appear as physical libraries to the backup application. As a result, they integrate more easily with existing backup applications, processes, and procedures. They can also integrate with physical tape and may offer capacity optimization techniques such as deduplication.
- **Continuous data protection.** In April 2005, the Storage Networking Industry Association (SNIA) defined CDP as “a methodology that continuously captures or tracks data modifications and stores changes independent of the primary data, enabling recovery from any point in the past.”⁴ Today there are several different data protection solutions in the marketplace that meet the basic requirements of this definition and refer to themselves as “CDP”. However, they differ in the methods they use to continuously track changes; their awareness and integration with applications, databases, and file systems; and their recovery object granularity.
- **Intelligent disk.** An intelligent disk is a disk array, system, or appliance that has advanced data backup techniques such as snapshots, clones, long-distance replication, capacity


optimization techniques such as deduplication, and software to manage it all. In last year's survey, we referred to this technology category as rotating snapshots with replication.

Market Adoption: Disk-Based Data Protection

Adoption of disk-based data protection technologies has remained steady over the past 12 months. There is strong adoption of backup to conventional disk and VTL targets, while there is less penetration of CDP and intelligent disk data protection. The higher adoption of backup to conventional disk and VTLs relative to other technologies can be attributed to two factors. First, these technologies have been available in the market much longer than other technologies. Early disk adopters have used backup to conventional disk in their backup environments for years, and there has been strong adoption of VTLs over the last two years. Second, most firms want to improve their existing backup environments by introducing some kind of disk target. CDP is typically used to protect the most mission-critical applications that require any point in time recovery. Not many general applications have this type of recovery point requirement; this requirement is more typical of, for example, a transaction processing application in the financial services industry. Likewise, the advanced techniques of an intelligent disk solution such as snapshots, clones, and replication are typically employed to protect mission-critical and business-critical applications. A decision to use an intelligent disk solution is often tied to your decision to replace your enterprise storage infrastructure or your enterprise backup application. As a result, adoption and commitment to an intelligent disk solution is a more strategic decision than the acquisition and deployment of a disk target for your current backup environment (see Figure 2-1). There was a decline in the adoption of "intelligent disk" from 2006 to 2007. This 10% decline can be attributed to the modified definition of the technology category in this year's study. In 2006, respondents were asked if they used "rotating snapshots and replication" for backup. In 2007, the category was defined as all disk systems that support advanced data protection functions such as snapshots, replication, deduplication, management functionality, but all disk systems that support these techniques.

Figure 2-1: Data Protection Technology — 2006 To 2007 Adoption Trends

"Which of the following backup and recovery solutions do you use in your environment today?"



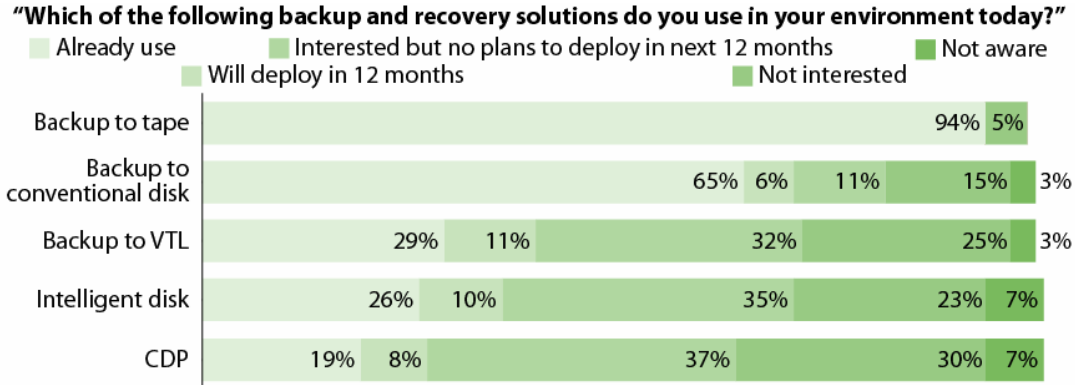
Base: All respondents; 2006 N = 401; 2007 N = 400
(multiple responses accepted)

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

While tape is entrenched in most firms, there is little interest in new deployments in the next 12 months. More firms plan to deploy VTLs followed by intelligent disk solutions. More firms are interested in intelligent disk solutions than in any other technology — they just don't have immediate plans to deploy. When an opportunity arises to replace incumbent solutions, firms are more likely to

adopt an intelligent disk solution such as an enterprise storage array with advanced protection techniques or an integrated backup appliance (see Figure 2-2).

Figure 2-2: Data Protection Technology — 2007 Current And Planned Adoption



(percentages may not total 100 because of rounding)

Base: All respondents; 2006 N = 401; 2007 N = 400

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Current Satisfaction

Firms continue to report high satisfaction with disk-based technologies. In fact, the vast majority of firms report that they are “satisfied” to “very satisfied” with all the disk-based technologies. In addition, from 2006 to 2007, the number of firms reporting that they were “very dissatisfied” to “dissatisfied” with their disk-based technologies decreased (see Figure 3).

Given that in both 2006 and 2007, firms rated completing backups within defined windows as the most important result to achieve, it’s safe to assume that the high degree of satisfaction with disk-based technologies is due to the fact that the introduction of disk into the backup environment has indeed helped firms meet their backup windows. Fifty-five percent of firms that currently use or plan to deploy “intelligent disk” report that they were “very satisfied” with the technology, the highest degree of satisfaction of any technology.

Figure 3: Overall Satisfaction



Base: Respondents who are already using each backup and recovery solution

† In 2006 intelligent disk was referred to as “rotating snapshots with replication.”

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Disk As Primary Backup Repository

According to last year’s survey, firms reported that they could store backup data on disk an average of 3.4 weeks or almost 1.5 months before vaulting to tape. In this year’s study, enterprises report storing data on disk for almost 3.0 months.

While all regions increase the length of time they stored backup data on disk, the largest increase was in North America. North American respondents reported an increase in the length of time they store backup data on disk from less than one month in 2006 to almost three and a half months in 2007 (see Table 3). Some of this increase can be attributed to adoption of disk technologies with capacity optimization techniques such as deduplication, but it doesn’t explain the increase entirely since there wasn’t a significant difference in the adoption of deduplication across the regions. In both 2006 and 2007, North America had the highest adoption of all disk-based technologies with the exception of CDP. It’s therefore likely that since North American firms were the early adopters of disk-based data protection and they’ve had an opportunity to see the results and develop expertise, they are at a point where they feel more comfortable storing data for longer periods of time on disk.

Table 3: Data Retention On Disk

Region	2006 disk retention*	2007 disk retention
All	1.5 months	3.0 months
North America	.9 months	3.3 months
Europe	.8 months	1.6 months
Asia Pacific	1.5 months	3.7 months

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

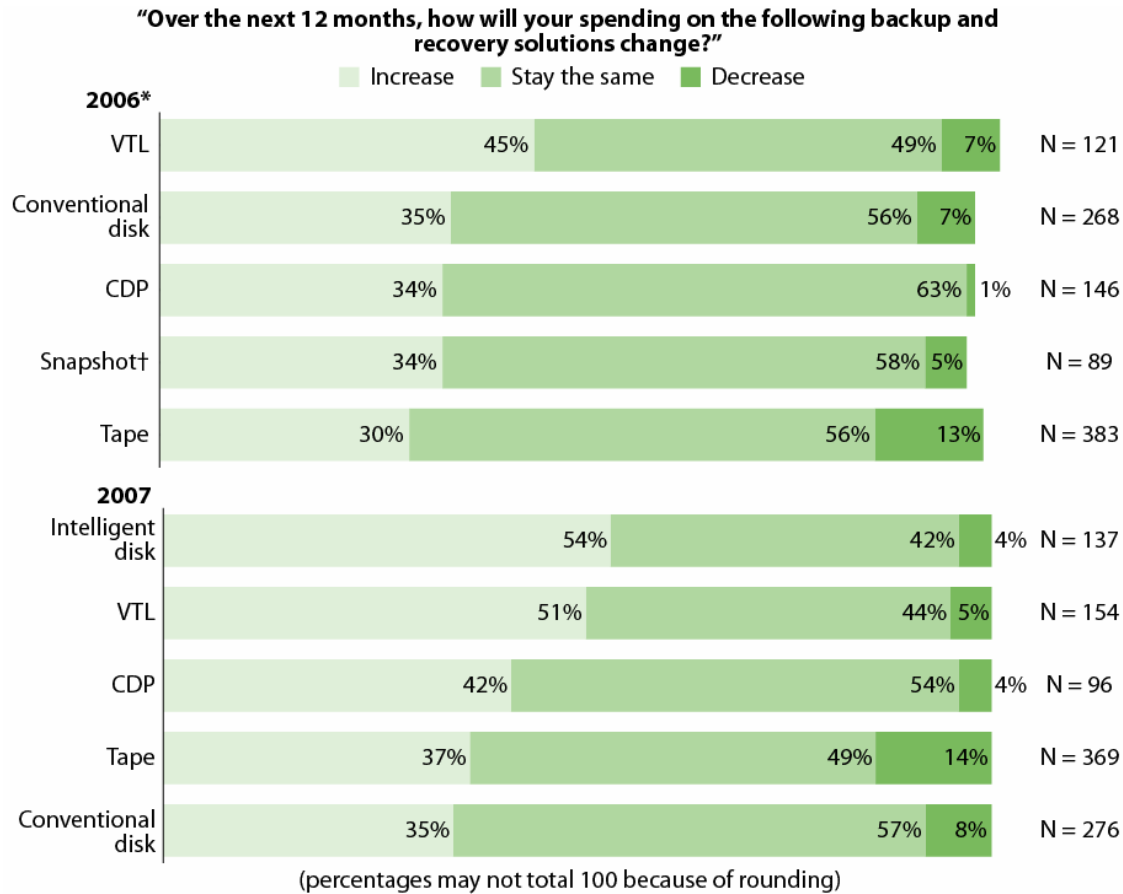
Market Demand: Disk-Based Data Protection

Just as the total budget for backup and recovery is increasing, spending on individual data protection technologies is increasing across the board for both disk and tape technologies, but firms report that they will spend more on disk (see Figure 4).

In 2006, 45% of firms that currently use or planned to deploy a VTL reported that they planned to increase spending on VTLs. In 2007, this increased to 51%. It's clear that there is strong spending intent for VTLs. However, actual VTL adoption remained flat over the past year, most likely because backup budgets increased only slightly while capacity skyrocketed. It's often easier to put a band-aid on your existing solution than to build a business case for new capital investment to senior management. This would account for the 7% increase in firms planning to increase spending on tape in 2007. However, given that 45% of respondents in 2006 and 51% of respondents in 2007 planned to increase their spending on VTL, it's expected that VTL adoption will increase over the next several years.

Another interesting trend in spending plans is the almost 20% increase in the number of firms planning to spend more on intelligent disk solutions. In 2007, more firms are reporting strong interest in intelligent disk, and now that they are reporting strong intent to invest in the technology, it's clear that firms are very interested in the range of recovery points and deduplication capabilities offered by these solutions.

Figure 4: Future Spending Plans On Data Protection



Base: Respondents who are already using or will deploy in 12 months each backup and recovery solution

† In 2006 intelligent disk was referred to as “rotating snapshots with replication.”

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Drivers Of Adoption

We asked enterprises that currently deploy or are planning to deploy disk-based data protection solutions to rank the importance of specific factors in their consideration of disk-based data protection solutions. Reliability is the most important consideration, followed by improved recovery time and speed of backups — these considerations also ranked the highest in 2006 (see Figure 5).

It’s likely that reliability will always remain the top consideration; no firm would deploy an emerging or even mainstream technology without proof of its reliability. Even with a mature technology such as tape, firms are reluctant to introduce new tape capabilities without customer references and documented best practices. And of course, until firms have more accurate capacity forecasts and the time and resources to optimize their backup operations, meeting backup windows and improving recovery capabilities — the fundamentals of backup — will always rank as important considerations.

Figure 5: Most Important Disk Considerations

“On a scale of 1 to 4 where 1 = not important and 4 = very important, how important are the following in your consideration of disk-based backup technologies?”

	2006*	2007
Reliability	3.7	3.8
Improved recovery time and recover point objectives	3.5	3.6
Speed of backups	3.5	3.5
Integration with current backup software and business apps	3.3	3.3
Economics (total cost, ROI)	3.3	3.2
Reduction of management complexity	3.4	3.1
Improved backup and recovery at remote sites	3.1	3.1

Base: Respondents that currently use or will deploy in the next 12 months VTLs, conventional disk, or intelligent disk; 2006 N = 347; 2007 N = 338


*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
 Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Barriers To Adoption

As the 2006 study proved, the market for disk-based data protection has already transitioned from an emerging market phase to an expanding market phase (based on a three-phase technology adoption cycle of emerging, expanding, and mature). In the expanding market phase, the market understands the technology, but customers are still waiting for the cost to drop a bit more. In the 2007 study, firms that currently do not use disk were either satisfied with tape or they were stymied by the initial acquisition cost of disk (see Figure 6).

Some firms are simply satisfied with tape; perhaps they have smaller storage capacities to protect or they have the luxury of long maintenance windows or they don't have stringent recovery time and recovery point objectives. Whatever the reason, these firms are likely the market laggards — they'll wait until the cost of disk technology drops to an extreme low or tape becomes obsolete before they transition. For everyone else, capacity optimization techniques, such as deduplication, that will help lower the total costs of ownership of disk and help IT operations build the business case for the initial investment to senior management will help drive new adoption of disk.

Figure 6: Barriers To Adoption



Base: Respondents who currently have 0% backup capacity as disk; N = 68 (multiple responses accepted)

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007


Deduplication: The Key To More Affordable Disk Data Protection

The Storage Networking Industry Association (SNIA) defines deduplication as "the replacement of duplicate data with references to a shared copy in order to save storage space".⁵ Deduplication will help firms store more backup data on fewer physical disks, as well as store this data on disk for longer periods of time.

According to the study, 9% of respondents already use deduplication in all of their backups, 17% of respondents use it to support some backups, and 4% are testing it (see Figure 7-1). Adoption is high for functionality that is still emerging. Today, deduplication is supported in some VTLs and intelligent disk solutions. Within six months to one year, the functionality will be nearly ubiquitous in VTLs and more widely available in intelligent disk solutions. In addition to the high adoption, there is a high degree of awareness: Only 8% of respondents reported that they were not aware of the technology. It's clear that the promise of deduplication to reduce capacity requirements is resonating strongly in the market.

Figure 7-1: Deduplication — Current Adoption

“Are you currently utilizing capacity optimization technologies such as deduplication in your backup-to-disk environment?”



Base: All respondents; N = 400

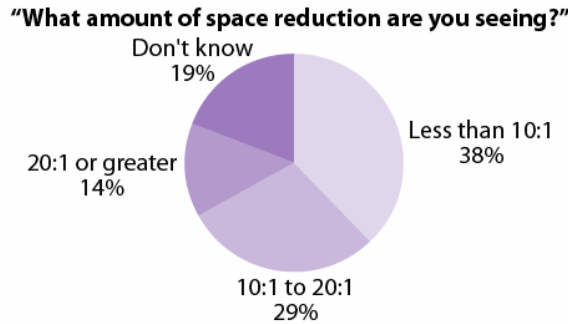
Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

More widespread availability of deduplication functionality will certainly increase adoption. But for adoption to really take off, firms will need help understanding the advantages and disadvantages of the various deduplication methods and guidance on the expected space savings. Many vendors are

claiming data reduction ratios such as 50:1 and sometimes even 100:1. The reality is that reduction ratios will vary greatly from firm to firm. Reduction ratios are dependent on data type and the number of previous backups you intend to store.

The good news is that early adopters are reducing their capacity requirements, but the reality is that they are achieving much more modest reduction ratios. According to the study, 38% of firms achieve reduction ratios of less than 10:1, and 29% achieve between 10:1 and 20:1 (see Figure 7-2).

Figure 7-2: Deduplication —Space Reduction Ratios




Base: Respondents who are currently utilizing optimizing technologies for at least some backups; N = 104
Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Firms report strong future adoption of deduplication. According to the study, approximately 36% of firms plan to expand the use of deduplication in their backup environment. This is healthy planned adoption for functionality that is still considered emerging (see Figure 7-3).

Figure 7-3: Deduplication — Planned Adoption

“What are your plans for data deduplication in your backup to disk environment in the next year?”



(percentages may not total 100 because of rounding)
Base: All respondents; N = 400

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

The Future Of Disk And Tape In The Enterprise

In both the 2006 and 2007 study, it was clear that most firms will continue to take a tiered disk-and-tape approach to data protection. However, in 2007, there is a greater percentage of respondents who want to ultimately eliminate tape (see Table 3-1). And these respondents believe that they can eliminate tape in less than 30 months.

Table 3-1: Expected Mix Of Disk And Tape

	2006*	2007
Current mix	29% disk 71% tape	33% disk 67% tape
Expected mix in two years	43% disk 57% tape	46% disk 54% tape
% who want to eliminate tape entirely	20%	26%
Expected time to tape elimination	28 months	30 months

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
 Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

There are some interesting regional differences in the current and planned mix of disk and tape. European firms have a much lower ratio of disk to tape and while firms in this region plan to increase spending on disk, even two years from now they still expect to have a ratio of 35% disk to 65% disk (see Figure 3-2). They also have the lowest percentage of respondents wanting to eliminate tape entirely. It's clear that tape will remain strategic to most firms in Europe. Because of this, it's logical to assume that disk-based data protection technologies such as VTLs that integrate and optimize with physical tape will be more preferable.

Table 3-2: Expected Mix Of Disk And Tape

	North America	Europe	Asia Pacific
Current mix	33% disk 67% tape	22% disk 78% tape	42% disk 58% tape
Expected mix in two years	48% disk 52% tape	35% disk 65% tape	52% disk 48% tape
% who want to eliminate tape entirely	31%	18%	23%
Expected time to tape elimination	32 months	19 months	32 months

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Optimizing Disk And Tape With Virtual Tape Libraries

As is evident in both the 2006 and 2007 study, tape is still very much entrenched in backup environments. And while more firms plan to spend more on disk than tape, tape is still important to backup — particular in long-term retention of backups. For environments that have a large investment in tape or simply want to continue use both disk and tape, VTLs are very appealing. Because VTLs are easy to introduce into existing backup environments, they are easy to manage and flexible and they integrate with and optimize physical tape.

The study validates this value proposition. When VTL users were asked to rate the importance of certain VTL features, performance, integration with existing backup applications, and manageability were the most important features (see Figure 8-1).

Figure 8-1: VTL Satisfaction

“On a scale of 1 to 4 where 1 = not important and 4 = very important, how do you rate the following features?”

Performance	3.6
Ease of integration into existing backup apps, processes, and procedures	3.4
Manageability	3.4
Ability to vault data from the VTL to tape	3.2
Flexibility to share VTL with more apps and servers	3.1
Efficient regrouping and de-staging to physical tape	3.0
Remote replication between VTLs	3.0

Base: Respondents who currently use VTLs or will deploy them in 12 months; N = 159

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

In both 2006 and 2007, most VTL users agreed that VTLs provide the benefits of disk-based data protection with the least amount of disruption to existing tape environments, and that VTLs provide an intuitive way to address tape limitations (see Figure 8-2).

Figure 8-2: VTL Benefits

“On a scale of 1 to 4 where 1 = strongly disagree and 4 = strongly agree, rate your level of agreement with the following statements about VTLs.”

	2006*	2007
Benefits of disk, yet nondisruptive to existing tape environment	3.1	3.2
Provides an intuitive way to address some of the limitations of tape-only backup environments	3.2	3.2
Easy to manage	2.7	3.0
Maximize existing tape limitations	2.9	2.9
Perform better than conventional disk arrays	2.9	2.9

Base: Respondents who currently use VTLs or will deploy them in 12 months; 2006 N = 161; 2007 N = 159

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

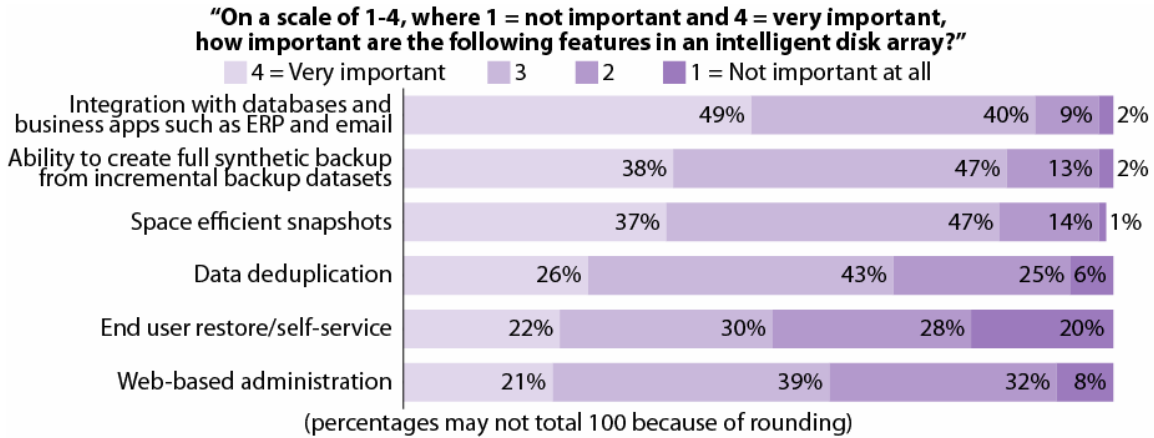
Intelligent Disk

In the 2006 study, “intelligent disk” was referred to as disk-only solutions that supported snapshots and replication and had other advanced data protection techniques such as granular recovery points and recovery objects (e.g., the ability to restore an individual file or email) and deduplication. In the 2007 study, this category was referred as “intelligent disk.” It not only refers to enterprise disk systems that support these advanced features but also purpose-built disk systems for data protection that integrate these features with backup application software into an appliance.

When current and planned users of intelligent disk solutions were asked to rank the most important features of an intelligent disk solution, they ranked integration with applications, the ability to create a full synthetic backup, and space efficient snapshots as the most important (see Figure 9-1).

Space-efficient snapshots will help firms set up a rotating snapshot strategy that offers multiple points of recovery throughout the day or potentially over several days. Synthetic backups are important because they allow firms to reduce backup windows by performing only incremental backups yet restore a full backup when necessary. Integration with applications is critical because to ensure that a snapshot or a backup represents a consistent point in time, applications and file systems need to be taken offline or quiesced before the snap or backup is performed. In the past, this process required scripting, but increasingly, intelligent disk solutions are “application-aware” and no longer require scripting.

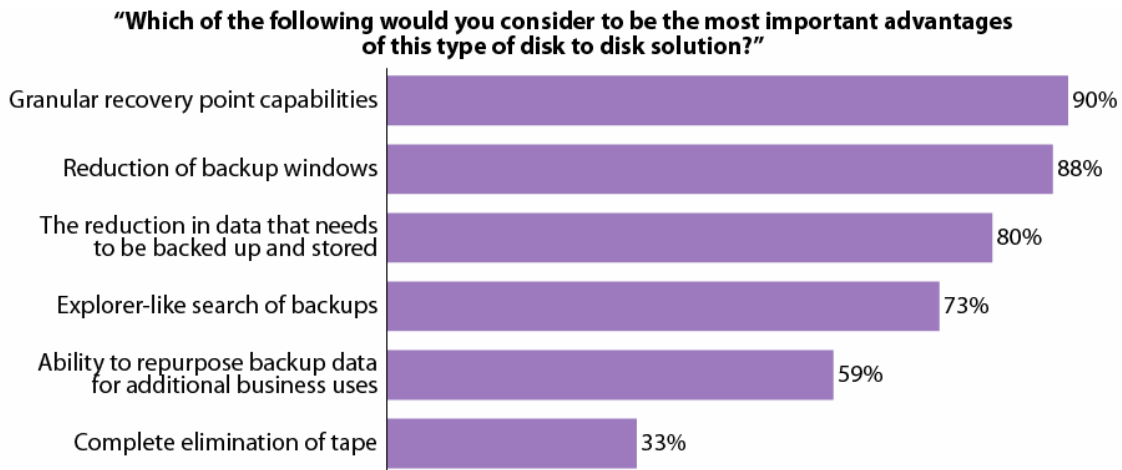
Figure 9-1: Intelligent Disk Feature Requirements



Base: Respondents who currently use intelligent disk or will deploy them in 12 months; N = 145
 Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

When current and planned users of intelligent disk solutions were asked to identify the most important advantages of an intelligent disk solution, more respondents selected reduction of backup windows and granular recovery points (see Figure 9-2). These advantages correspond to the most important considerations for disk-based data protection identified in Figure 5.

Figure 9-2: Intelligent Disk — Key Benefits



Base: Respondents who are currently use intelligent disk or will deploy them in 12 months; N = 145 (multiple responses accepted)
 Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Integrating Data Protection For The Data Center And Remote Sites

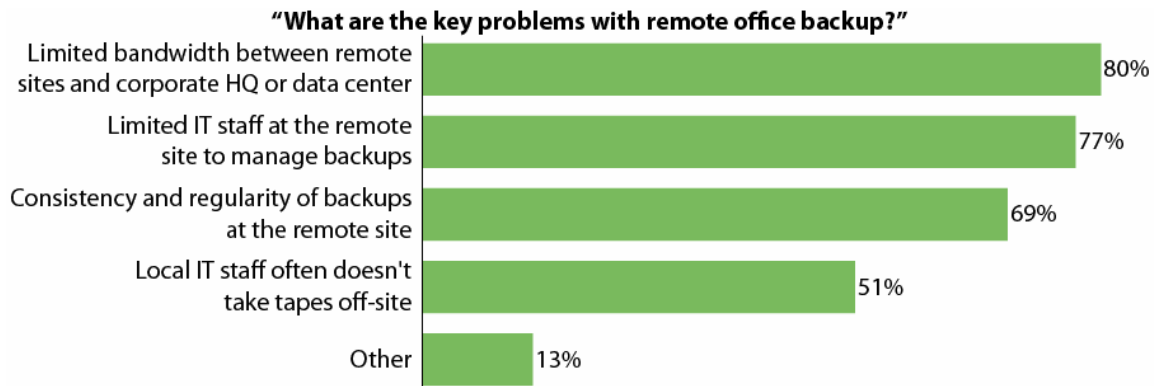
Remote sites include a firm's major data centers, main and regional headquarters, branches, offices, etc. Firms are increasingly distributed (less than 8% of respondents reported that they did not have a remote site) and as such they need to be concerned not only with the protection of data at their main data center and headquarters but also the data generated and stored at remote sites such as branches and offices. In some cases, mission and business critical data maybe stored at these remote offices.

Remote site backups are very challenging because they require local hardware and software and local staff to manage it, and corporate IT has little insight into the regularity and success of local backups and whether physical tapes are taken off-site according to schedule. The physical transport of tapes is risky in and of itself, unless the tapes are encrypted before they are transported off-site. Lost or stolen backup tapes with personal data can cost a firm millions of dollars in compliance and damage its reputation. According to the 2007 study:

- Firms will spend approximately 16% of their backup and recovery on solutions to solve remote office challenges. Asia Pacific spends the most on the remote office: 21% of its backup and recovery budget.
- Respondents reported that remote office backup is more problematic in 2007 than it was in 2006. In the current study, 39% of respondents reported that remote office backup was "problematic" to "very problematic" as compared with 34% of respondents in 2006.
- When respondents were asked how they primarily back up data today at remote offices, 72% selected local backup to tape and 44% selected remote backup to a central facility. Respondents could select multiple approaches. There is some adoption of remote backup solutions, but local backup to tape is still the dominant approach.

When respondents were act to identify the key problems of remote office backup, 80% of respondents identified limited bandwidth between remote sites and headquarters, and 77% identified limited IT staff at the remote office (see Figure 10-1).

Figure 10-1: What's Problematic About Remote Office Backup

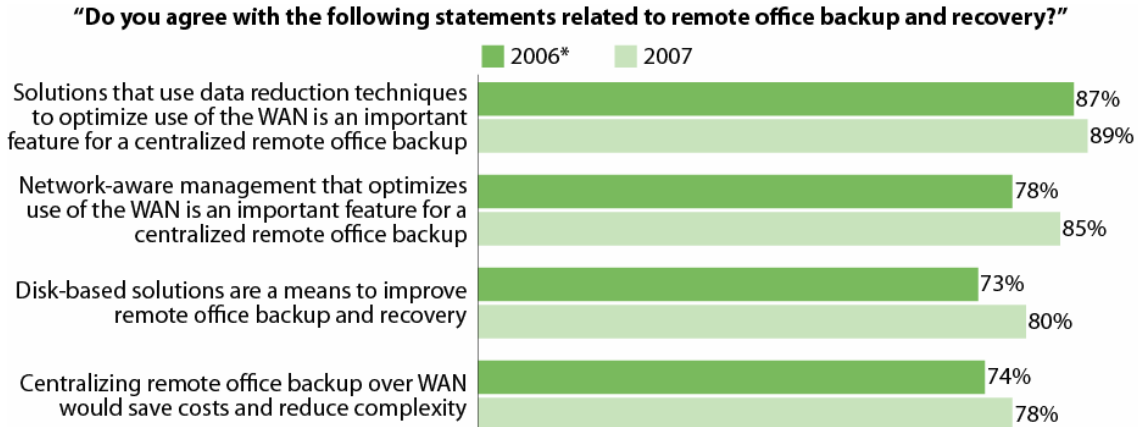


Base: Respondents who find remote office backup and recovery problematic; N = 142 (multiple responses accepted)

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Not surprisingly, more firms agree that data reduction techniques and network aware management are important features of a remote office backup solution. If limited bandwidth between the remote office and headquarters is a key problem of remote office backup, then it's clear that firms will have a preference for a solution that reduces the amount of data that must be transmitted over the wide-area network (WAN) and hence the amount of required bandwidth and ensures that backups can be scheduled during off-peak times or can be configured to consume less of the available bandwidth so as not to compete with regular corporate traffic (see Figure 10-2).

Figure 10-2: What Firms Want In A Remote Office Backup Solution




Base: Respondents with at least one remote site in their organization; 2006 N = 321; 2007 N = 369 (multiple responses accepted)

*Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, August 2006
 Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Completely eliminating local backups is not always possible, no matter how WAN-efficient the backup technology. For larger remote offices, the initial full backup cannot be performed over the WAN. Instead, tapes or removable disks must be shipped to the corporate office to "seed" the first backup. All other backups will only transmit incremental data over the WAN, and in cases where deduplication is employed, only unique data segments are transmitted. While backups can be very WAN efficient, restores of large data sets can be very painful over the WAN. Many firms will want to continue to backup locally to disk but then vault a second copy of the data to a corporate facility. The local copy gives the remote office fast local restores. When firms were asked to identify the approach to remote office backup they would likely deploy in two years, a majority of firms across regions and company size selected backup locally to disk (see Figure 10-3).

Figure 10-3: What Solutions Will Firms Use To Solve Remote Office Backup

“Which of the following approaches to remote office backup are you most likely to implement?”



Base: Respondents with at least one remote site in their organization; N = 369

Source: A commissioned study conducted by Forrester Consulting on behalf of Network Appliance, Winter 2007

Conclusions

The objectives of the study were to understand how the adoption and satisfaction of disk-based has evolved from 2006 to 2007, the demand for advanced functionality, the role of deduplication in adoption and the state of remote office data protection. The study's findings can be summarized as follows:

- The disk-based data protection adoption remained steady from 2006 to 2007. There were not major increases or decreases of any technology. The market remains in an expanding market phase where there are some early adopters, but a greater number of firms are still waiting for prices to decline. Yet interest and planned spending on disk is very strong. Firms still expect to achieve a 46% to 54% ratio of disk to tape in two years, and 26% of firms plan to completely eliminate tape.
- There is strong evidence to suggest that firms that have already deployed disk-based data protection technologies are seeing value from these technologies and gaining expertise. Satisfaction increased from 2006 to the 2007, and the average length of time data is stored in disk increased from a month and a half to three months.
- There is strong interest in intelligent disk solutions that support multiple recovery points, granular object recovery, data deduplication, snapshots, replication, and application awareness. The most important features of an intelligent disk solution are application awareness and integration, synthetic backups, and space-efficient snapshots.
- Aside from those firms for whom tape meets recovery objectives, the biggest barrier to disk adoption is still cost. Thirty-four percent of respondents indicated that they were waiting for capacity optimization techniques to help reduce the total cost of disk. Another 43% indicated that the capital acquisition cost of disk was still too high. It's clear that deduplication, when it is proven and widely available in multiple data protection technologies, has the potential to rapidly accelerate the adoption of disk-based data protection adoption.
- The study found that there is already healthy adoption of deduplication: 9% of respondents employ it in all their backups, and 17% employ it in some of their backups. The study also found that most firms are achieving deduplication ratios of 20:1 or less.

- Despite some clear differences in the ability of various technologies to provide application awareness or granular recovery objects, firms did not show a preference in the use of specific data protection technologies for a given data set.
- From 2006 to 2007, remote office data protection has actually become more problematic. There is some use of remote backup to solve some of the challenges, but local backup to tape is still the dominant approach to remote office backup. Limited bandwidth between sites and lack of staff at the remote office are the key problems. Firms do agree that remote office backup solutions that can be centrally managed, are WAN-efficient, and are based on disk will help solve remote office backup challenges. In two years, firms expect to solve their remote office backup challenge with local backup to disk, likely with the vaulting of second copy to a corporate facility.

Enterprise Recommendations

Firms have multiple disk-based data protection technologies to choose from. In fact, most data protection strategies will incorporate several tape and disk technologies. The right combination of technologies and approaches depends on multiple criteria, including: 1) specific recovery objectives for what needs to be protected (database, messaging and collaboration application, or unstructured file data); 2) the restore granularity requirements (i.e., image or object level); 3) the disruptiveness of the technology to the existing backup environment; 4) the manageability of the technology (is it a standalone technology? Or is it managed from within existing backup applications?); and 4) the total cost of the technology, including any licensing costs, capital acquisition costs, manageability, etc. In addition to developing a comprehensive strategy that incorporates multiple technologies, firms should:

- **Have more realistic and accurate capacity forecasts.** Firms do not have accurate capacity forecasts. In last year's study, respondents reported an expected capacity growth of 19%. It turns out that the one year increase in capacity is close to 100%. This underestimation of capacity growth will keep firms forever playing catch-up when it comes to data protection, not to mention essential storage management and storage operations.
- **Consider deduplication an essential feature of disk-based data protection.** The next disk-based protection technology you evaluate should support deduplication. Deduplication will allow you store and retain more data on less physical disk capacity for longer periods of time. In the case of remote office data protection, deduplication will help reduce the amount of data that must be transmitted over the WAN.
- **Get to know your applications and your application owners.** Based on their recovery requirements and recovery granularity, different disk-based data protection technologies are more or less appropriate for databases versus messaging and collaboration applications versus unstructured file data. It doesn't appear that firms use different data protection techniques to meet the needs of specific data sets or applications. Firms need to work with application owners to understand their requirements and discuss how the various disk-based technologies can augment existing backup operations.
- **Prioritize remote office backup.** Data loss at remote sites remains a huge risk exposure for most firms. Firms agree on the key features of a remote office backup solution — central backup administration and management, disk, deduplication, and network-aware management. There are several solutions and approaches in the market that meet these requirements. It's time for firms to decide on an approach and develop a shortlist of vendors and technologies with which to move forward.

Endnotes

¹ Within Europe, Forrester interviewed 54 enterprises from the UK and 46 enterprises from Germany. Within Asia Pacific, Forrester interviewed 33 enterprises from Australia, 34 enterprises from India, and 33 enterprises from Japan.

² Percentages may not total 100 because of rounding.

³ Percentages may not total 100 because of rounding.

⁴ Source: SNIA DMF CDP SIG working definition, April 2005

⁵ Source: SNIA Dictionary <http://www.snia.org/education/dictionary/>