

# NetApp ROI Calculator User Guide

**To Accompany NetApp ROI Calculator Version 2.0**

**Specific to the Data Protection (Backup & Recovery, Disaster Recovery) Solution  
Segment**

November 2004

PROPRIETARY AND CONFIDENTIAL

# Table of Contents

[Disclaimer](#)

Questions

Using the Calculator .....	3
Figure 1) NetApp ROI Calculator 2.2	
How Inputs are Entered and Defined	
General Inputs Worksheet	
Figure 2) General Inputs Worksheet for Data Protection Solutions	
Backup Environment Details (Worksheet 1) .....	7
Figure 3) Backup Environment Worksheet (1)	
Backup Environment Details (Worksheet 2) .....	9
Figure 4) Backup Environment Worksheet (2)	
Figure 5) Backup Environment Worksheet detail	
Disaster Recovery Details Worksheet .....	11
Figure 6) Disaster Recovery Worksheet	
How Results Are Derived and Varied .....	14
<a href="#">The Upper Left Corner</a>	
Backups .....	14
Disaster Recovery .....	21

## Disclaimer

No portion of this calculator, or any of its output or results, may be reproduced, distributed, or published in any form, whether physical, electronic, virtual, or otherwise, except with the prior expressed, written permission of the copyright holder, Network Appliance Inc. This calculator is meant for use by NetApp sales, in direct connection with customer and prospect discussions only.

## Questions

For additional information about this guide or the NetApp ROI calculator, please contact June McDonald at [junem@netapp.com](mailto:junem@netapp.com) or 408-822-3238.

## Using the Calculator

The NetApp ROI Calculator is a Microsoft Excel-based proprietary tool designed to show customers and prospects the cost advantages of NetApp solutions. In its current form, the calculator offers the following ROI (return-on-investments) solution comparisons:

- Data Protection
  - Backup & Recovery
  - Disaster Recovery
- Storage Consolidation
  - File Server
  - SAN/IP SAN Storage




Figure 1) NetApp ROI Calculator 2.0

The calculator is designed to be simple and straightforward to use. To begin, simply open the calculator workbook and choose which of the ROI comparisons (Backup & Recovery, Disaster Recovery, File Server Consolidation or SAN/IP SAN Storage

Consolidation), to perform<sup>1</sup>. Step-by-step, the calculator guides the user through the necessary input worksheets until an ROI result is reached.

For Data Protection Solution comparisons (Backup & Recovery, Disaster Recovery), “**Back**”, “**Next**”, and “**Home**” buttons are available on each worksheet to allow for easy navigation. Print buttons are also available on selected sheets for properly formatted output to a printer. For Storage Consolidation Solutions comparisons (File Server, SAN/IP SAN Storage Consolidation) “Start”, “Questions”, “Assumptions”, “Summary”, “Report”, and “Benefits” tabs are available at the bottom of the workbook. For all solutions, the resulting ROI comparisons produced provide a valuable basis for discussion by using customer and prospect specific environment and requirement data.

This guide has been written to accompany version 2.0 of the ROI calculator, contains an explanation of all calculator inputs and results. The calculator attempts to achieve a balance between detail and simplicity. While producing a fairly sophisticated financial comparison, the number of inputs has been minimized in order to make the tool easier to use and the calculation formulas have been provided in order to reduce skepticism. Each input and result parameter is explained in detail in this document and the calculator Appendix A. If a user questions the calculator’s methodology, this approach allows full transparency into the tool’s assumptions and conclusions.

The calculator is divided into several input worksheets. For both solution sets (Data Protection and Storage Consolidation), there is a general input worksheet that collects data common across multiple areas. There are also additional specific worksheets for each ROI solution (specific to Backup and Recovery, Disaster Recovery, File Server and SAN/IP SAN Storage Consolidation).

While the series of quick inputs contribute to larger equations, the inputs are intended to require a minimum amount of time to answer. Each ROI comparison output contrasts various solution options with both numeric costs (and the intermediate values used to determine the final results) as well as a graphical display.

The following is a detailed explanation of the calculator inputs and outputs, specific to the Data Protection Solutions (Backup & Recovery, Disaster Recovery).

---

<sup>1</sup> If prompted by Microsoft Excel, workbook macros need to be enabled via the appropriate application security setting (go to application menu bar, select Tools, Options, Security, Macro Security, opt for Medium security, hit OK; save, close, and re-open workbook; enable macros).

## How Inputs are Entered and Defined

### General Inputs Worksheet

General Environment Parameters	
<b>Site Information</b>	
Amount of data to be protected (in TB)	<input type="text" value="60"/>
Number of years to amortize	<input type="text" value="3"/>
<b>Staff Costs</b>	
Number of operations staff	<input type="text" value="3"/>
Annual Cost Per Operations Staff for Data Protection	<input type="text" value="\$90,000"/>
<b>Storage Costs</b>	
Average Primary Storage cost/TB	<input type="text" value="\$55,000"/>
Cost of NearStore	<input type="text" value="\$1,300,000"/>
Cost of Centera	<input type="text" value="\$0"/>
Cost of SnapMirror	<input type="text" value="\$35,000"/>
Cost of SnapVault	<input type="text" value="\$100,000"/>
VTL Software Cost	<input type="text" value="\$30,000"/>
<input type="button" value="BACK"/> <input type="button" value="HOME"/> <input type="button" value="NEXT"/> <input type="button" value="Print Sheet"/>	

Figure 2) General Inputs Worksheet for Backup and Recovery

The General Inputs worksheet gathers environment information necessary for multiple ROI calculations, regardless of the specific solution. For the Backup and Recovery solutions, the inputs are as follows:

- **Amount of data to be protected (in TB)** – The number of terabytes (TB) to be protected by either a backup or disaster recovery strategy.
- **Number of years to amortize** – Most companies calculate capital costs (such as hardware purchases) over the useful life of the product. Choosing a value of “1” here simply will use the actual cost of any hardware or software purchases. Any other value will assume straight-line depreciation (i.e., any capital cost will be divided by the number of years when calculating annual cost).
- **Number of operations staff** – The number of administrative and/or professional staff (sometimes referred to as “operators” herein) working on data protection in the particular environment.

- **Annual cost Per Operations staff for Data Protection** – The salary cost of each operator or administrator working on data protection. Note that often times, operators and/or administrators spend only a portion of their time on data protection solutions such as backup or data replication. In these cases, the salary cost should be adjusted down to compensate.
- **Average Primary Storage Cost/TB** – The purchase cost of a TB of primary storage.
- **Cost of NearStore** – The hardware purchase cost of the NearStore storage solution.
- **Cost of Centera** – This input is only relevant when comparing NetApp solutions to EMC Centera-based solutions. If Centera is not being compared, this input can be left at “0” (the default value). Otherwise, enter the cost of the Centera storage solution.
- **Cost of SnapMirror** – Enter the cost of the SnapMirror software license. This does not include any hardware costs. If SnapMirror is not part of the desired comparison, a value of “0” can be entered here.
- **Cost of SnapVault** – Enter the cost of the SnapVault software license. This does not include any hardware costs. If SnapVault is not part of the desired comparison, a value of “0” can be entered here.
- **VTL Software Cost** – Enter the cost of virtual tape library software. This does not include any hardware costs. If VTL is not part of the desired comparison, a value of “0” can be entered here.

## Backup Environment Details Worksheet (1)

Backup Environment Details	
Number of servers (backup clients) to backup	100
Incremental backup window (in hours)	9
Full backup window (in hours)	40
Backup Window for Nearline to Tape Backup (in hours)	72
Number of days between Full Backups	30
Full backups run during incremental backups	<input checked="" type="checkbox"/> TRUE
Relative size of incremental vs full (file level backup)	19.5%
Block level change in data (per day)	2.3%
Average backup failure rate	2%
Average number of restores/week	39
Cost of Data Unavailability (per hour, in dollars)	\$5,000 <small>Sample Figures</small>
Average data retention for incrs (in days)	90
Average data retention for fulls (in days)	90
Average data retention for archive copies (in days)	365
Frequency of archive copies (times per year)	0
SnapVault Snapshot Retention (days)	90
Offsite copy of backups?	None

Figure 3) Backup Environment Worksheet (1)

- **Number of servers (backup clients) to backup** – The number of systems being backed up. More clients will lead to higher software licensing costs for all solutions other than SnapVault and SnapMirror. Thus, a larger number of clients can offer a significant advantage when considering these solutions.
- **Incremental backup window (in hours)** – The number of hours available to successfully perform an incremental backup. In many companies, different servers will have different backup windows. In this case, it is recommended to use an average value.
- **Full backup window (in hours)** – The number of hours available to perform a full backup.
- **Backup Window for Nearline to Tape Backup (in hours)** – The number of hours available to move data from a nearline staging solution (such as NearStore) to tape. If there is no particular window, simply adjust this to the maximum value of “720”.
- **Number of days between Full Backups** – The number of days between a full backup for a typical system.
- **Full backups run during incremental backups** – Set this to “True” by checking the box if full backups run during the week in parallel with incremental backups. In this case, full backups are typically staggered throughout the week with some systems receiving a full backup on Monday, some on Tuesday, etc. If full backups run separately from incremental backups (e.g., on the weekends), set the value to “False”.
- **Relative size of incrementals vs. full** – In general, considers how large an incremental backup is compared to a full. This is a good measure of the volatility of data in a given environment.
- **Block level change in data (per day)** – This input only applies to block level backups. In order to highlight the benefits of a block level backup solution such as SnapVault/SnapMirror, the number of blocks that change (as a percentage) on a daily basis should be estimated and input here.
- **Average backup failure rate** – Note that different companies define backup failures differently. Some only consider hard failures. Increasingly however, companies are defining any backup that does not complete successfully within its window to be a failure.
- **Average number of restores/week** – How many restores are performed in a typical week?
- **Cost of Data Unavailability (per hour)** – While not all data is valued equally; this is a rough estimate of the cost of being unable to access a given data item for one hour. This calculates the cost of restore delays.
- **Average data retention for incrs (in days)** – How many days are typical incremental backups retained? Most companies seem to retain incremental backups for about 30 days.
- **Average data retention for fulls (in days)** – How many days are typical full backups retained? Note that full backups are considered separately from archived backups. One year is a common value for full backup retention.

- **Average data retention for archive copies (in days)** – How long are archived backups retained? The average value for archive data retention will vary by industry and data type. Some companies will do no archive copies at all (relying instead on fulls). Financial institutions will typically keep at least *some* of their data for 7-8 years. The legal, medical and pharmaceutical industries also tend to have long archive retention periods.
- **Frequency of archive copies (times per year)** – How often are archived copies of data created? For example, many companies consider their month end full backups to be archive copies.
- **SnapVault Snapshot Retention (days)** – Enter the number of days that SnapVault snapshots are retained.
- **Offsite copies of backups** – Choose the frequency (or none) of offsite backups.
- **Monthly cost to store cartridge offsite** – Enter the cost to store a tape cartridge offsite.

## Backup Environment Details Worksheet (2)

Backup Environment Details	
Tape Drive Type	LTO 2 <input type="button" value="Modify Associated Values"/>
Automatically calculate Tape Drive/Media requirements	<input checked="" type="checkbox"/>
GB of disk for backup and media servers	<input type="button" value="←"/> <input type="button" value="0"/> <input type="button" value="→"/> 400
Capital Cost per backup server	<input type="button" value="←"/> <input type="button" value="10,000"/> <input type="button" value="→"/>
Per-server Additional Backup SW cost	Windows <input type="button" value="Modify Associated Values"/>
Server Backplane Throughput (MB/s)	<input type="button" value="←"/> <input type="button" value="1,024"/> <input type="button" value="→"/> <input type="button" value="Sample Figures"/>
Avg Software cost per Backup Client	<input type="button" value="←"/> <input type="button" value="820"/> <input type="button" value="→"/>
Annual Backup software maintenance %	25% <input type="button" value="↓"/>
Average Networking cost per client	<input type="button" value="←"/> <input type="button" value="100"/> <input type="button" value="→"/>
Number of Tape Libraries	<input type="button" value="←"/> <input type="button" value="1"/> <input type="button" value="→"/>
Total capacity of tape libraries (in cartridges)	<input type="button" value="←"/> <input type="button" value="3000"/> <input type="button" value="→"/>
Cost of a Tape Library	\$50,000 <input type="button" value="↓"/>
Backup SW cost per tape drive	\$3,000 <input type="button" value="↓"/>
Number of Master Servers	<input type="button" value="←"/> <input type="button" value="1"/> <input type="button" value="→"/>
Average Cost of a Failed Restore	\$1,000 <input type="button" value="↓"/>
<input type="button" value="BACK"/> <input type="button" value="HOME"/> <input type="button" value="NEXT"/> <input type="button" value="Print Sheet"/>	

Figure 4) Backup Environment Worksheet (2)

- **Tape Drive Type** – Choose the type of tape drive to be used in the backup solution. Clicking the “Modify Associated Values” button reveals the characteristics of each drive type. Customers can modify these values if they disagree. The drive type determines cartridge capacity and the speed of the tape drive.

Backup Environment Details	
Tape Drive Type	LTO 2 <input type="button" value="Modify Associated Values"/>
Number of Tape Drives	<input type="text" value="14"/>
Number of Tape Drives for Duplication	<input type="text" value="3"/>
Tape Cartridges per Year	<input type="text" value="1110"/>
Tape Cartridges per Year for Duplication	<input type="text" value="170"/>
Tape Drive Type #2 (optional)	None
Number of Tape Drives	<input type="text" value="0"/>
Tape Cartridges per Year	<input type="text" value="0"/>
Automatically calculate Tape Drive/Media requirements <input type="checkbox"/>	

- **Automatically calculate Tape Drive/Media requirements** – If this box is checked the following inputs (as described immediately below) are automatically calculated. Otherwise, they can be adjusted manually. The inputs are as follows:
  - **Number of tape drives** –The number of tape drives used for backups.
  - **Number of tape drives for duplication** –The number of tape drives necessary to duplicate tapes.
  - **Tape cartridges per year** – The number of cartridges used for primary copies of backup data per year.
  - **Tape cartridges per year for duplication** –The number of cartridges used for duplication of backup data per year.
  - **Tape Drive Type #2 (optional)** – Allows an optional second type of tape drive to be entered. Some customers have hybrid tape backup environments.
  - **Number of tape drives** –The number of “Type #2” tape drives that are being used.
  - **Number of tape cartridges per year** –The number of “Type #2” tape cartridges that are used per year.
- **GB of disk for backup and media servers** – Backup software requires considerable disk capacity to store file catalogs and indexes.
- **Capital Cost per Backup Server** – Enter the purchase cost of the server being used to run the backup software.
- **Per Server Additional Backup SW Cost** – Choose Windows, UNIX or Linux. Typically, backup software packages charge a fixed price based on operating system for a backup server or media server. There is a “Modify Associated Values” button here which will allow customers to modify these costs. It is difficult to generalize these costs as different companies reach different pricing agreements with their backup software vendors.
- **Server Backplane Throughput (MB/s)** – The backplane speed of a server and the read/write speed of a tape drive determine the number of media servers

necessary to build an effective backup solution. Putting too many tape drives on a server results in a saturated server backplane and suboptimal tape drive performance. This value will drive the auto-calculation of the number of necessary media servers. Note that sample figures for various server types are provided here. Many customers will find these examples helpful.

- **Avg Software Cost per Backup Client** – Most backup vendors charge a license fee that is, in part, based on the number of backup clients.
- **Annual Backup Software Maintenance %** – Most backup vendors charge an annual maintenance percentage based on the purchase price of the backup software. This fee covers support as well as software upgrades.
- **Avg Networking Cost per client** – Data backups are typically one of the most network intensive applications in an environment. This cost estimates the average cost of the network infrastructure necessary to run backups. For companies that build a dedicated network for backups, this number will be considerably higher.
- **Number of Tape Libraries** – This is the number of tape libraries required for a traditional tape backup solution, exclusive of tape duplication.
- **Total capacity of tape libraries (in cartridges)** – How many cartridges do these libraries hold in aggregate?
- **Cost of a tape library** – Enter the average cost of each library.
- **Backup software cost per tape drive** – Most backup vendors charge a license fee that is, in part, based on the number of tape drives in the environment. *Note that– the number of tape drives required for each type of solution is calculated automatically based on tape drive type, backup window and amount of data.*
- **Number of master servers** – Each backup master server controls scheduling, file indexes and restores for each server whose backups it manages.
- **Average cost of a failed restore** – On average, what is the cost to the business of being unable to perform a restore.

## Disaster Recovery Details Worksheet

The calculator examines the cost of both site failure and individual server data unavailability. The calculator does not consider the failure of servers themselves. Rather it is concerned with whether or not the data itself is available. Thus, it assumed that if data is available, there will be systems available to access the data. The costs of these servers are not considered in the calculator's results.

Disaster Recovery Details			
Number of servers to protect	<input type="text" value="100"/>		
Percent of Systems that suffer outages per year	<input type="text" value="5%"/>		
Avg Duration of System outage using tape restore (hours)	<input type="text" value="12"/>		
Avg Freq Site outage per year	<input type="text" value="0.5"/>		
Avg Duration of Site outage (in hours)	<input type="text" value="24"/>		
Hard Costs			
Cost of NetApp Primary Storage	<input type="text" value="\$1,100,000"/>		
Cost of non-NetApp Primary Storage	<input type="text" value="\$1,300,000"/>		
Hardware maintenance costs (for Networks and Storage)	<input type="text" value="25%"/>		
Software maintenance costs (for Replication Software)	<input type="text" value="13%"/>		
Cost of FC Switched ports (HBA + switch port)	<input type="text" value="\$3,000"/>		
Cost of Ethernet port for replication (NIC + switch port)	<input type="text" value="\$2,000"/>		
Cost of non-NetApp Synchronous Replication Software	<input type="text" value="\$180,000"/>		
Cost of non-NetApp Asynchronous Replication Software	<input type="text" value="\$170,000"/>		
Avg Cost of FC remote Link (per month)	<input type="text" value="\$15,000"/>	Number of FC Links	<input type="text" value="0"/>
Avg Cost of IP remote Link (per month)	<input type="text" value="\$10,000"/>	Number of IP Links	<input type="text" value="1"/>
Soft Costs			
Cost (per hour) of data unavailability	<input type="text" value="\$10,000"/>		<input type="button" value="Sample Figures"/>
Cost (per hour) of degraded availability	<input type="text" value="\$3,000"/>		
<input type="button" value="BACK"/> <input type="button" value="HOME"/> <input type="button" value="NEXT"/> <input type="button" value="Print Sheet"/>			

Figure 6) Disaster Recovery Worksheet

- **Number of servers to protect** – Enter the number of servers to protect in the environment.
- **Percent of servers that suffer system outages per year** – On average, how often does a given system become unavailable?
- **Average Duration of System outage using tape restores (hours)** – In general, how long does a server outage last? It is assumed that tape restore is the only method available to recover data.
- **Avg Frequency Site outage per year** – On average, how often is an entire site unavailable.
- **Avg Duration of Site outage (in hours)** – When a site outage does occur, how long does it last?
- **Cost of NetApp Primary Storage** – How much is FAS storage per TB (this is additional storage associated with DR)?
- **Cost of non-NetApp Primary Storage** – What is the cost of non-NetApp primary storage per TB (this is additional storage associated with DR)?
- **Hardware Maintenance Costs (for Networks and Storage)** – The annual cost of hardware maintenance expressed as a percentage of the purchase price. Note that the network costs are represented here as an effective disaster recovery solution will typically require network infrastructure for replication.

- **Software Maintenance Costs (for replication software)** – On average, what is the annual cost for support and upgrades for replication software (e.g., SnapMirror, SnapVault or other non-NetApp replication software) expressed as a percentage of the purchase price.
- **Cost of FC Switched ports (HBA + switch port)** – The cost of a Fibre Channel (FC) connection. This cost includes the Host Bus Adapter (HBA) and the cost of a port on a FC switch.
- **Cost of Ethernet port for replication (NIC + switch port)** – The cost of an IP network connection being used for replication. Note that this is typically a high bandwidth connection (i.e., gigabit Ethernet). The cost includes the Network Interface Card (NIC) and the cost of a port on an IP switch.
- **Cost of non-NetApp Synchronous Replication software** – The purchase price of synchronous (typically FC-based) replication software. This input reflects the aggregate price necessary to protect the entire environment.
- **Cost of non-NetApp Asynchronous Replication software** – The purchase price of asynchronous (typically IP-based) replication software. This input reflects the aggregate price necessary to protect the entire environment.
- **Avg Cost of FC Remote Link (per month)** – The cost of a FC connection between data centers. This is generally not a one time cost, but rather a recurring monthly charge.
- **Avg Cost of IP Remote Link (per month)** – The cost of an IP connection between data centers. Again, this is generally a recurring monthly charge.
- **Cost (per hour) of data unavailability** – On average, what is the business cost of data being unavailable for one hour. Note that sample figures by industry are provided by choosing the “Sample Figures” button.
- **Cost (per hour) of degraded availability** – On average, what is the business cost of data only being available in a degraded mode. Note that degraded mode is defined as data that normally resides on primary storage residing on nearline storage instead, reducing the level of availability compared to primary storage. For some companies, this is a relatively small number as the effect may not be that severe. For other companies in a highly transactional environment, the cost may be very high as data availability performance will be severely reduced.

## How Results Are Derived and Varied

### The Upper Left Corner

The upper left corner of the TCO worksheets contains several check boxes. Selecting them will toggle columns on or off. Most of these are straightforward (e.g., Virtual Tape Library, Traditional Tape Backup) and should be selected or de-selected depending upon what the customer is comparing against NetApp. The less intuitive boxes are:

- **Show soft costs** – Many customers will discount (sometime totally) soft costs such as operations time and such. This box allows these costs to be excluded or included.
- **Show formula details** – Checking this box will show the intermediate values that drive the final results. Some customers will want to see these details.
- **Leverage DR site for backups (Backup TCO only)** – Setting this value to “True” will reduce the amount of NearStore storage required for backups. It will assume that since a full copy of the environment is retained on NearStore for disaster recovery, this copy will be leveraged as a full backup. Thus, only enough storage for incremental backups on NearStore will have to be purchased.
- **NearStore used for backups (DR TCO only)** – This removes the hardware cost from the NearStore backup column.

### Backups

Many customers will undoubtedly question how the calculator arrives at its results. This section documents the equations used. It is assumed that the “Show Formula Details” option has been selected on the “TCO” tabs (see above). Note that all formula details reflect the automatic calculation of tape drives and tape media quantities. Only the final results reflect the values that are manually entered when “Automatically calculate Tape Drive/Media requirements” is unchecked.

- **Number of Tape Libraries** – This number is input directly on an earlier worksheet.
- **Number of Tape Drives** – Unlike libraries, this number is calculated based on backup window; data volatility and tape drive characteristics. The number is sum of “Num Tape Drives” and “Tape Drives for Duplication.” Both are described below. Notice that in the case of VTL, it is assumed that only a single drive is needed in order to occasionally offload a virtual tape to a physical image. This calculation is perhaps one of the less intuitive computations in the calculator. The number of drives for full and incremental backups considers the amount of data to be backed up, the backup window, the speed of the tape drives, and whether or not full backups run simultaneously with incremental backups. Drives for duplication are calculated as described below. Note that if the “Automatically calculate Tape Drive/Media requirements” box is not checked, then the number of tape drives is to be input directly by the user. No automatic calculation is performed in this case.

- **Number of Tape Drives for Full Backups** – The aggregate amount of data in the environment, full backup window and tape-drive speed are examined to derive the number of tape drives required to complete full backups within the window. Notice that with backup strategies that incorporate NearStore, the backup window is significantly longer since data is first written to NearStore and then written to tape at a later point. Also notice that if full backups run simultaneously with incrementals, then it is assumed that the full backup job schedule is divided equally over the number of days between full backups. This can significantly reduce the necessary number of tape drives to run fulls.
- **Number of Tape Drives for Incremental Backups** – For NearStore-based incrementals, no tape drives are required. For tape-based incrementals, the aggregate amount of data in the environment is multiplied by a volatility factor (which is directly input on an earlier worksheet). The tape-drive speed and incremental backup window (which is typically much shorter than the full backup window) are considered to arrive at the number of tape drives necessary to successfully complete incremental backups.
- **Num Tape Drives (intermediate result only displayed under formula details)** – On an earlier worksheet, the user chose whether or not full backups ran simultaneously with incrementals. If this is the case, it is assumed that the correct number of tape drives will be the sum of drives necessary for both fulls and incrementals. If they do not run simultaneously, it is assumed that whichever value is greater (the number of drives for fulls or the number of drives for incrementals) is the correct value. This is true because it is pointless to have a significant number of drives sitting idle. Many customers run incrementals during the week and fulls on the weekends. In cases such as this, incremental drive requirements never conflict with full backup drive requirements.
- **Duplication Window** – This only applies to customers who create duplicate copies of their backup data. This is input on an earlier worksheet. The duplication window determines how many additional tape drives are needed in order to duplicate tapes. This value represents the number of hours between the end of a full backup and the start of the next full backup. It is assumed that this is the time available to duplicate the backup media.
- **Tape Drives for Duplication** – The aggregate amount of data in the environment is divided by the amount of data that can be backed up by a single tape drive in an hour. The result is then further divided by the number of hours in the duplication window. This yields the number of drives necessary to duplicate all data in the environment within the duplication window. Notice that this assumes that drives operate at full capacity for the entire window.
- **Tape Cartridges/yr** – Tape cartridges are a significant operational spend for most companies. From a sales point of view, it is worth noting that operational costs such as tape media are often in a different budget than capital costs such as

hardware. This often affects sales strategy. Tape cartridge cost is a function of retention cycles, tape capacity, number of copies and cartridge cost. The details are explained below:

- **Tapes used per full bkp** – This is the amount of data in the environment divided by tape cartridge capacity.
- **Tapes used all full bkps per year** – Based on the retention and frequency of full backups, the number of cartridges required in a year is calculated. It is assumed that full backup cartridges are reused when their retention period has expired.
- **Tapes used per incr backup** – The average incremental size (derived from the “relative size of an incremental vs full” input on “Bkp1” worksheet) is divided by the cartridge capacity.
- **Tapes used all incr bkps per year** – Based on the retention and frequency of incremental backups, the number of cartridges required in a year is calculated. It is assumed that incremental backup cartridges are reused when their retention period has expired.
- **Tapes used all archive bkps per year** – Archive backups are presumed to be full backups that never expire. This means that these cartridges are written exactly once. Archive tapes represent a significant expense because of this.
- **Backup (master) servers** – Master servers perform all control and management functions of the backup. They also store all file catalogs for restores. This value comes from an input on the “Bkp-2” worksheet.
- **Backup (media) servers** – Media servers are servers whose primary purpose is just to power tape drives. Typically, only a limited number of tape drives can be attached to a server before saturating its backplane. The total number of servers is calculated by examining tape drive speed and the backplane of a typical backplane. These values are controlled by the tape drive type and server type entries on the “Bkp2” worksheet.
- **Backup Server Capital Cost per year** – Average cost of a backup server is a input on the “Bkp2” worksheet. Notice that the “per year” cost is arrived at by using straight line depreciation based on the number of years used for asset amortization. This value is entered on the general “Inputs” worksheet.
  - **Tape Drives per media server** – This is based on backplane speeds. Higher backplane capacity allows more drives per media server. This is controlled by the choice of backup server type on the “Bkp2” worksheet. The “Modify Associated Values” button on the “Bkp2” worksheet can be used to alter these values.
  - **Num Servers** – Derived by taking the number of tape drives necessary (from result above) to successfully perform backups and calculating the number of servers necessary to power the tape drives. Again, this is largely dependent on the backplane capacity of the servers.
- **Backup software per year** – Backup software vendors typically price their software according to the number of backup servers and the number of tape drives. Customers can input the price per backup server and tape drive on the “Bkp2” worksheet. Likewise, customers can input software cost for NearStore

Software costs are assumed to be capital costs and are amortized over a number of years (this value is input on the “Inputs” worksheet). Software maintenance is also factored in. The components in this cost are explained below:

- **Software cost/client + software cost/tape drive** – The number of clients and tape drives is simply multiplied by the appropriate software costs entered on the “Bkp2” worksheet.
- **Software cost/bkup server + Software cost/media server** – The number of backup and media servers are multiplied by the software licensing costs entered on “Bkp2” worksheet.
- **Software cost NearStore** – This is the cost for any software such as SnapMirror, SnapVault, VTL, etc.
- **Software costs \* maintenance** – Backup software costs are multiplied by maintenance percentages entered on the “Bkp2” worksheet.
- **SW costs above/number of years** – Aggregate software costs are divided by the number of years over which capital costs are amortized.
- **Backup disk capital cost per year** – Regardless of backup methodology, disk can represent a significant cost. This is because when backups are performed by traditional backup software, file catalogs are created. File catalogs are essentially indices of every file that has been backed up. File catalogs typically make up 1% of the total amount of data being backed up. Creating file catalogs can become a significant bottleneck to the backup process if the underlying disk infrastructure does not perform significantly well. Thus, the calculator assumes that primary storage costs apply for disk that is attached to the backup server.
  - **NearStore Disk Requirement (TB)** – The data capacity required of the NearStore for the particular environment. Based on retention cycle and amount of data being protected.
  - **Backup Server Disk Costs** – Primary disk attached to the backup servers and the cost of primary storage are both entered on the “Bkp2” worksheet.
  - **NearStore costs** – The cost of the NearStore hardware entered in the “Inputs” worksheet. This cost should be for a NearStore solution that can accommodate the “NearStore Disk Requirement” above.
  - **Number of Years to Amortize** – From the “Inputs” worksheet
- **Tape Drive Capital Costs per year** – This is straightforward multiplication based on the number of tape drives (explained above) times the cost of a drive (from the “Bkp2” worksheet) divided by the number of years over which capital purchases are amortized. Notice that the default cost of a drive can be modified by choosing the “Modify Associated Values” option on the “Bkp2” worksheet.
- **Library Cost per year** – Library cost is input on the “Bkp2” worksheet. The “per year” value is simply this cost multiplied by the number of libraries (determined above) divided by the number of years over which capital costs are amortized (from the “Inputs” worksheet).
- **Media** – Calculated by multiplying the number of cartridges used in a year (determined above) times the cost per cartridge. Note that the cost per cartridge is determined based on default values for each type of tape (e.g., DLT, LTO, etc.) The tape type is chosen on the “Bkp2” worksheet. To modify these default values, choose the “Modify Associated Values” button at the top of the “Bkp2”

worksheet. It should be expected that larger customers may have more favorable media prices.

- **Additional Networking Costs** – In general, it is difficult to determine a standard cost for networked backups. Some companies build dedicated backup networks (including switches, wide area links, network interface cards, etc.) Others attempt to run over their existing network. We approximate this by allowing users to enter a “Cost Per Backup Client of Gigabit Ethernet connection” on the “Bkp2” worksheet. This cost is multiplied by the number of clients (and divided by the number of years over which capital costs are amortized). Companies with dedicated networks will have significant costs here. Companies using existing networks will have significantly lower costs. It is a rare (and possibly non-existent) case that the additional network costs are zero. Even with an existing network, any type of backups will consume a significant amount of bandwidth. Eventually, this load will have to be considered in networking purchases and upgrades.
- **Offsite storage** – The per cartridge cost of offsite storage is input on the “Bkp1” worksheet.
- **Operations** – Personnel costs are usually a significant portion of backup costs. This number includes not only backup administrators, but also the operators who are responsible for monitoring backups on weekends or nights. Costs considered by this calculator include:
  - **Cost of Staff** – this is a straightforward calculation of the number of people working on backups times their average salary. Keep in mind that often times these people are not fully dedicated to backups. In this case, it is important to make sure that partial headcount (e.g., overnight operators who have many responsibilities) are being included. These values are input on the “Inputs” worksheet.
  - **Time to change tape (minutes)** – This is the average tape handling time to move a tape for backup. Typically, tapes are rotated in batches in and out of libraries and to offsite storage. This number represents an average. At present, this is hard-coded at five minutes as most customers will not have a good handle on the average.
  - **Number of tapes changed per year** – It is assumed that each cartridge is handled once. Given that most non-archive tapes reside in a library (and do not have to be handled after being entered into the library), this is a reasonable assumption.
  - **Tape Change Labor (for backups)** – Based on cost of staff, the number of tape changes and the time necessary to change a tape, the calculator arrives at a cost.
  - **Time to find and retrieve tape (minutes)** – This is currently hard-coded at sixty minutes. Again, most customers won’t have a good idea of a time here. An hour seems reasonable given that some percentage of tapes will be archived offsite or not in a library. Further, operators will typically watch the progress of a restore as it occurs.

- **Number of tapes retrieved per year** – This is calculated by taking the number of restores/week (from the “Bkp1” worksheet) and multiplying by 52 weeks.
- **Tape Retrieve/Load Labor (for restores)** – Based on cost of staff, the number of tape retrieve/loads and the time necessary to load a tape, the calculator arrives at a cost.
- **NearStore Labor Savings** – Since NearStore can handle offsite replication and can eliminate tape handling, it is assumed that NearStore can remove the tape handling labor charges described above.
- **Avg Cost of Data Unavailability During Restores** – This is often one of the most stark differences between NearStore based solutions and traditional backup. NearStore offers remarkably rapid restores. This is particularly true for cases where multiple files have to be restored simultaneously. The numbers for data unavailability cost were derived by multiplying the number of restores in a year by the average time to perform a restore by the average cost of data unavailability per minute. The details follow:
  - **Data Unavailability cost per minute** – This is based on the data unavailability cost per hour (divided by 60) that is entered on the “Bkp1” worksheet.
  - **Time to restore (minutes)** – This number is hard-coded based on typical averages. It is assumed that tape based restores (including mount time, seek time and transfer time) require 30 minutes while disk based restore times average 3 minutes.
  - **Number of restores per year** – This is based on user input from the “Bkp1” worksheet.
- **Avg Cost of Failed Restores** – Backup failures are typically the single most expensive aspect of a suboptimal backup environment. They introduce risk and exposure into a company. Backup failures are defined as any backup that does not complete successfully within its assigned backup window. This is an important distinction. It is not enough to consider a backup successful simply because it completes without errors. Backups that do not complete within their assigned window create a number of issues surrounding data integrity. Introducing NearStore into a backup environment can dramatically increase backup success rates.
  - **% Backup Failures** – This is taken from industry averages. For example, Enterprise Strategy Group estimates that 40% of networked backups fail (using the definition above). The failure rate falls dramatically, for disk based backups.
  - **Number of Restores per year** – This is based on user input from the “Bkp1” worksheet.
  - **Number of Failed Restores per year** – The “% Backup Failures” is multiplied by the “Number of Restores per year” to arrive at this number.
  - **Avg Cost of Failed Restore** – Taken from user input on the “Bkp2” worksheet.
  - **Cost of Failed Restores** – “Avg Cost of Failed Restore” multiplied by the “Number of Failed Restores per year”.

## **Assumptions**

Building calculators such as this is an inexact science. Assumptions always have to be made. It is hoped that these assumptions are generally correct. There are, however, always exceptions to the general case. Here, the less intuitive assumptions for each type of backup are explained. If a situation occurs where an assumption is clearly incorrect for a given customer, it should be relatively straightforward to manually adjust the results to accurately reflect the customer's environment.

## **Typical Tape Solution**

The way in which the number of master and media servers is calculated is not very intuitive. Once the appropriate number of tape drives is calculated (based on the amount of data to be backed up and the length of the backup window), the calculator calculates the number of servers based on the backplane speed of the servers. The backplane calculation is often overlooked (even by customers). Today's tape drives are actually quite fast. They are still relatively poor at seeking and positioning. However, sequential reads and writes are quite fast. Putting more tape drives on a server than the backplane can handle will cause the tape drives to perform at suboptimal speeds. This can lead to somewhat unexpected results on occasion. For example, putting a large amount of data, with a very short backup window will often cause the calculator to report a large number of tape drives and a large number of media servers. Often times, a customer will find that their requirements simply make a tape solution unrealistic. It is possible to toggle automatic calculation of tape drives and media servers on or off. If a customer strongly disagrees with the automatically calculated values, the values should be input manually.

Generally, backup software costs are directly correlated with the number of tape drives. Thus, inputs that generate a large number of tape drives will also likely yield very significant backup software costs.

Networking costs tend to be higher for a traditional tape backup solution than for any other type of solution. Because tape backup solutions involve a large number of servers writing to a smaller number of tape drives, it is very important that no one server take more time than necessary to finish its backup. With disk based solutions, there is far less (in some cases none) contention for access to the nearline disks being used for backups. This is because, as a random rather than sequential access device, disk can easily support multiple servers writing simultaneously. In this scenario, if one server writes particularly slowly, it won't delay another server's backup. With tape, it is often necessary to have dedicated backup networks (or at least to make sure that there are no bottlenecks in the network) in order to avoid significant queuing delay to write to a tape drive.

## **NearStore Staging**

Another disk based backup approach is staging. This serves as an alternative to synthetic backups as described above. With staging, both incrementals and fulls are written to tape. However, all data is written first to NearStore. The data is written to tape from the NearStore. This reduces (or eliminates) the problem of contention for tape drives. It also

allows data to be written to tape over a much longer period of time (since the backup is already on disk) greatly extending the backup window. The calculator thus assumes that a single tape library (possibly a large one) will be adequate for this approach. Of course, this can be argued with depending on a customer's preferences. However, in general, this should be an accurate, cost efficient assumption.

This method also assumes that tapes are not being duplicated. This is because two copies of data will exist for some period of time (i.e., one on NearStore, one on tape). For most customers, this is adequate. For others, the media costs will have to be increased to reflect the cost of tape duplication.

### **VTL (Virtual Tape Library)**

Virtual tape libraries are becoming increasingly popular. To backup software, VTLs look like tape libraries. However, in reality, they are disk. This allows a backup administrator to create a large number of "virtual drives" so that contention for tape resources (and the queuing associated with this contention) largely disappears. VTLs can be used in a number of ways. The calculator assumes that in general, data written to a VTL will stay on a VTL and will only be written out to tape as an exception. This eliminates the need for a tape library or more than a single tape drive. It also eliminates the need for tape media and the media servers associated with tape drives. Some customers may feel the need to duplicate their VTLs. In this case, SnapMirror is the most straightforward way to accomplish this. The calculator does not take this duplication into account.

### **Open Systems SnapVault or SnapMirror/SnapVault**

OSSV and SnapMirror are combined in that it is assumed that a complete backup solution will use OSSV for data not stored on NetApp systems and SnapMirror/SnapVault for data stored on NetApp systems. Both are assumed to be necessary. It is also assumed that full and archive backups will be copied to tape. This assumption makes the cost profile of OSSV very similar to that of VTL with full and archive tape duplication. The key difference from the VTL case is the backup disk cost. This is largely dependent on the SnapVault Retention Days variable set on the Bkp1 worksheet. Setting this value to a large number of days can result in a large cost (particularly if data volatility is high).

### ***Disaster Recovery***

- **Amount of Data to Protect** – This is entered on the "Inputs" worksheet.
- **# of systems to protect** – This is entered on the "DR" worksheet.
- **Annual Data Unavailability Cost** – This reflects the cost of data unavailability after implementing a specific disaster recovery strategy. The end result is the product of multiplying the cost per minute of data unavailability by the average number of minutes of data unavailability in a year after implementing a particular disaster recovery strategy. Notice that the calculator does not measure server unavailability. Server unavailability is a different (but related) problem. It is assumed that a successful disaster recovery strategy will take servers into account, but that is beyond the scope of this calculator. The calculator assumes that if the

data is available and accessible, servers and the surrounding infrastructure will also be available. The following data is used to calculate the cost of annual data unavailability:

- **Avg Number of System outages per year (excl site outages)**– Calculated by multiplying the number of servers to protect times the percentage of systems that suffer outages per year. Both of these values come from the “DR” worksheet.
- **Avg Number of Site outages per year** – Input from the “DR” worksheet.
- **Reduction in outages per year due to DR solution** – This is a hardcoded value based upon our expectations of various disaster recovery solutions. The values can be tweaked by overwriting the cells if customers disagree with our estimates.
- **Duration of Avg Site outage (hours)** – Input from the “DR” worksheet.
- **Duration of Avg System outage (hours)** – Input from the “DR” worksheet.
- **Reduction in duration of outage due to DR solution** - This is a hardcoded value based upon our expectations of various disaster recovery solutions. The values can be tweaked by overwriting the cells if customers disagree with our estimates.
- **Site outages per year after DR solution** – This value is calculated by multiplying the “Reduction in outages per year due to DR solution” by “Avg Number of Site Outages per year.”
- **System outages per year after DR solution** – This value is calculated by multiplying the “Reduction in outages per year due to DR solution” by “Avg Number of System outages per year (excl site outages).”
- **Duration of Avg Site Outage (Hrs) after DR solution** – This value is calculated by multiplying “Reduction in duration of outage due to DR solution” by “Duration of Avg Site Outage.”
- **Duration of Avg System Outage (Hrs) after DR solution** – This value is calculated by multiplying “Reduction in duration of outage due to DR solution” by “Duration of Avg System Outage”.
- **Avg Site Unavailability Per Year (Hours)** – This represents the number of hours (on average) that a site is unavailable per year after implementing a DR solution. This is calculated by multiplying the “Duration of Avg Site Outage(Hrs) after DR solution” by the “Site Outages per year after DR solution.”
- **Avg System Unavailability Per Year (Hours)** – This represents the number of hours (on average) that a system is unavailable per year after implementing a DR solution. This is calculated by multiplying the “Duration of Avg System Outage (Hrs) after DR solution” by the “System Outages per Year after DR solution.”
- **Avg Total Data Unavailability Per Year (Hours) due to Site Outage** – This value is calculated by multiplying “# of systems to protect” by “Avg Site Unavailability Per Year (Hours).”

- **Avg Total Data Unavailability Per Year (Hours) due to System Outage** – This value is calculated by multiplying “# of systems to protect” by “Avg System Unavailability Per Year (Hours).”
  - **Avg Total Data Unavailability Per Year (Hours)** – This result is calculated by adding the data unavailability numbers for both site and system outages.
  - **Cost (per minute) of Data Unavailability** – The customer input for cost per hour of data unavailability (entered on the “DR” worksheet) is multiplied by 60 to get a per minute cost.
- **Annual Degraded Mode Cost** – Disaster recovery and backups are often confused. In general, backup solutions and disaster recovery require totally different solutions. Disaster recovery is optimized for rapid restoration (sometimes instantaneous) of current data. Backup is optimized to economically store historical (often only minutes or hours old) copies of data. Backup solutions are typically engineered to provide quick restoration of a few files, directories or databases. However, backup solutions are not engineered to provide rapid recovery of multiple systems simultaneously. Most companies will need both a backup solution and a disaster recovery solution. It is possible, however, to combine solutions and actually run in degraded mode (i.e., lower availability). An example of this would be to use NearStore for both backup and disaster recovery. Obviously, the impetus behind this would be cost. The major disadvantage of this approach is performance. NearStore is optimized for high capacity at a low cost. Most primary storage is optimized for performance. Most environments that are accustomed to running on FAS will not perform well on NearStore. Thus, we use the term “degraded mode”. This implies that the data will be available, but performance will be below normal production levels. Depending upon the underlying application, this can have serious implications as to the usability of the application. Often times, it is not possible to run production computing loads in degraded mode. Annual degraded mode cost is a straightforward calculation based on the cost per minute of running in degraded mode times annual minutes spent in degraded mode.
    - **Degraded mode time (hours)** – This number is assumed to be the difference between the outage hours before the DR solution is implemented and the outage hours after the DR solution. The logic here is that while a NearStore based DR solution will greatly reduce the number of outage hours, there will still be a process of migrating the data from NearStore back to primary storage. The environment will be in degraded mode during this period.
    - **Cost (per min) of degraded unavailability** – This value is entered on the “DR” worksheet.
  - **Operations Time (hours/year)** – It is assumed that for each hour of data unavailability or of a system being in degraded mode, an administrator will also have to spend an hour of time. This is a very conservative estimate as it does not take into account testing time after the data becomes fully available or other resources such as network administrators, etc.

- **Operations Cost of Data Degradation/Unavail.** – This value is calculated by multiplying the associated operations time as a result of a system or site outage by the hourly cost of the staff (“operator”).
  - **Avg Operator Cost** – The annual cost of an operator is entered on the “General Inputs” worksheet. In order to convert from an annual to a hourly rate, the annual cost is divided by 230 days (it is assumed that there are roughly 230 working days in a year) and further divided by 8 hours (assumed to be a normal working day).
- **Additional Cost of Storage Hardware** –Storage hardware is a significant component of any effective DR strategy. The difference in cost between solutions varies widely. DR solutions that are based upon nearline storage (such as NearStore) will be significantly less expensive than solutions built on primary storage (such as FAS or other vendors’ primary storage devices). The total cost is a straightforward calculation multiplying the number of TB of data to protect by the cost of the chosen storage solution.
  - **Additional Storage Requirement (TB)** – In most cases, this is simply the “amount of data to be protected” that is entered from the “General Inputs” worksheet. In cases where the backup solution is being leveraged for DR, this may not be the case however. It is possible to simply leverage backup tapes or backup images made to nearline storage to address DR. This approach and its direct and indirect costs are reflected in the calculators cost equations.
  - **Cost of Storage (TB)** – Depending on the solution, the appropriate cost for DR storage is entered. These values are gathered from the inputs on the “DR” worksheet. Depending on the case, NetApp FAS, NetApp NearStore or another vendor’s storage may be used in the calculation. Non-NetApp storage is used for non-IP based replication schemes. All other storage choices are clearly labeled. The storage costs entered in the “Inputs” worksheet should reflect prices for solutions that can accommodate the “Additional Storage Requirement (TB)”, which will in most cases simply be the “Amount of data to be protected.”
- **Additional Cost of Storage Software per year** – The cost of licenses for either SnapMirror, SnapVault or non-NetApp replication software divided by the amortization period. These costs are entered on the “DR” worksheet.
- **Associated Costs for Remote Links per year**– The cost of either FC (in the case of non-IP-based replication strategies) or IP links is factored in. Note that some companies will use existing network connections and others will add additional bandwidth. The number of links to be added and the monthly link costs are entered on the “DR” worksheet. Note that network charges are generally calculated monthly. For purposes of this cost, all numbers are multiplied by 12 months to get an annual cost.
- **Cost of FC Switched Ports and HBAs per year**– FC-based replication strategies will likely require a certain number of FC ports (i.e., Fibre Channel switches) and HBAs to be purchased. This value is input on the “DR” worksheet. These costs are divided by the amortization period.

- **Cost of IP Port and NICs per year** – IP-based replication strategies may require a certain number of IP network interface cards and dedicated IP switch ports to be purchased. This value is input on the “DR” worksheet. These costs are divided by the amortization period.
- **Hardware Maintenance Costs** – Most hardware has an associated annual maintenance cost. An average annual hardware maintenance percentage is entered in the “DR” worksheet. To determine this annual cost, the maintenance percentage is multiplied by the sum of total storage hardware and networking hardware costs. Note that the maintenance percentage is based upon total cost, not the annual cost.
- **Software Maintenance Costs** – Much like hardware, most software has an associated annual maintenance cost. This covers support, upgrades and the like. The annual maintenance percentage is entered on the “DR” worksheet. To determine annual maintenance cost, the percentage is multiplied by the total purchase cost (not the annual cost) of the replication software cost.
- **Total Annually** – This is the sum of all annual costs including amortization of hardware and software, labor, cost of downtime, media and maintenance.
- **Total Over Period** – This is the total cost over the life of the solution. This value is attained by multiplying the annual cost by the amortization period. It is assumed that the useful life of the solution is roughly equivalent to the amortization period.

## Assumptions

### Traditional Tape Backup

Tape backup is really not a credible disaster recovery solution. As such, the data unavailability costs of this solution will typically lead to an unacceptably high cost. Aside from this cost, relying on a company’s existing backup solution will undoubtedly yield a cheaper overall cost than any solution built to optimize disaster recovery. Notice that additional hardware and software costs are zero for this solution since it is assumed that a company already has a backup solution in place.

### NearStore Backup

Again, using a NearStore centric backup solution as a DR solution is really not acceptable in most cases. This approach represents an improvement over tape backup in that recovery time will generally be quicker with a NearStore based backup solution than with a tape based backup solution. The solution is not, however, engineered for rapid failover or large scale, multi-system rapid recovery. Again, additional hardware and software costs are assumed to be zero as this approach just reflects leveraging an existing backup solution for DR.

### Synchronous (non-IP) based Replication

It is assumed that this means Fibre Channel based replication. In general, while effective, this will be a more expensive solution than IP based replication. Note that synchronous solutions generally appear to be cheaper than asynchronous solutions. This is because asynchronous solutions almost always entail some degree of downtime for failover.

Synchronous solutions, on the other hand, can provide hot (and in some cases automated) failover.

#### **Synchronous IP based Replication FAS**

NetApp synchronous replication will almost always offer a cost advantage over Fibre Channel based replication.

#### **Synchronous IP based Replication NearStore**

This is the same solution as above but rather than failing over to primary storage (i.e., FAS), the failover is to NearStore. This provides immediate access to data, but the storage will run in degraded mode as NearStore is not designed to run at the same performance levels as FAS.

#### **Asynchronous (non-IP based) Replication**

It is assumed that this means Fibre Channel based replication. This will provide a more expensive solution than IP based solutions. Notice that most of the cost associated with any asynchronous method is in recovery time. While the recovery will be far quicker than with a tape backup solution, it is generally necessary to manually check each system before completing the failover.

#### **Asynchronous IP based Replication FAS**

NetApp asynchronous replication almost always offers a cost savings over a Fibre Channel based solution. All asynchronous solutions require some degree of system verification in order successfully complete a failover.

#### **Asynchronous IP based Replication NearStore**

This is the same as the FAS solution described above, but once the failover is complete, the systems will suffer degraded performance in that the NearStore is not designed to run at the same performance levels as FAS. Depending on the value chosen for the cost of being in degraded mode, this may or may not be a significant drawback to using NearStore for failover.