



ESG Lab Report™



Avamar Axion® - Now EMC Avamar Data Protection Software Platform

A validation study
by
ESG Lab
April 2006 - Updated January 2007

Authors:
Tony Asaro Heidi Biggar Brian Garrett Tony Palmer

Table of Contents

Introduction	3
ESG Lab Validation	6
<i>Avamar Data Reduction.....</i>	<i>7</i>
<i>Remote Replication</i>	<i>10</i>
<i>Ease of Use</i>	<i>13</i>
<i>Manager of Managers.....</i>	<i>16</i>
<i>Upgradeability.....</i>	<i>17</i>
<i>Avamar File System (AvFS)</i>	<i>19</i>
<i>Tape Integration.....</i>	<i>20</i>
<i>Integrated Search</i>	<i>23</i>
<i>Reliability and Security</i>	<i>24</i>
ESG Lab Validation Highlights	25
Roadmap Recommendations	25
ESG Lab's View	26
Appendix.....	27

ESG Validation Reports

The goal of the ESG Lab Validation reports is to educate customers about specific storage-related products including storage systems, backup-to-disk solutions, storage management applications, backup/recovery software, storage virtualization platforms, etc. The ESG Lab reports are not meant to replace the necessary evaluation process that end user customers should conduct. The ESG Lab reports are designed to provide insight to what is compelling about various products and how they can solve customer problems. ESG Lab reports also recommend areas we feel the vendor should improve upon. ESG Lab provides third-party expert perspective based on our own hands-on testing in a lab and interviews with customers using these products in production environments.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

EMC Acquires Avamar

Avamar was acquired by EMC in November 2006 and has re-branded the Axion product as *EMC Avamar* software. This report was updated in January 2007 to reflect the new name and product enhancements since ESG Lab first tested the product in April 2006.

Introduction


The backup and recovery process is ripe for reinvention. The majority of customers are struggling with one or more inefficiencies with their backup and recovery eco-systems. There has been a great deal of attention paid to backup and recovery performance. Many organizations are struggling with tape reliability and the challenges that tape creates with media management. The process of handling and managing tapes is a cumbersome process that can lead to human error and potential data loss. Using tapes for offsite archival can also lead to the exposure of confidential information if encryption is not being used.

Data residing outside the data center at remote offices accounts for a significant portion of an enterprise's information store, yet it either isn't protected at all or it's protected with inefficient traditional backup processes, leaving companies at risk for data loss. The backup software and tape-based devices typically deployed at remote sites give rise to numerous challenges for enterprises with remote offices:

- ☑ Untrained remote office staff, responsible for tape management and troubleshooting
- ☑ Tape failures, which must be dealt with by unskilled staff
- ☑ Undocumented processes
- ☑ WAN bandwidth limitations force shipment of tapes (usually unencrypted)

Not only are backup and recovery related issues (such as media management, backup/recovery performance, reliability, scalability, and cost) an ongoing problem for most organizations, but the scale of the problem is expected to worsen over time. Data growth is on the rise, and traditional data protection solutions magnify the problem by multiplying that growth. Backup, snapshots, mirroring, staging and archiving all significantly increase the amount of storage under management, adding to both operational and capital costs. In a distributed environment, centralizing backup operations with traditional backup architectures and existing network capacity is impractical at best. The network bandwidth required to move daily backup data, as well as the time it would take, is not cost-effective or efficient.

Figure One: Adverse Effects of Current Backup Infrastructure



Source: ESG Research Tape Replacement Realities 2005

As Figure One illustrates, backup and recovery time consistently come up as the top issues for customers. However, companies are not willing to break the bank to improve backup and recovery performance. ESG Research has also found that customers will implement solutions, but do need a way to do so cost effectively. Additionally, tape is not going away. Our research shows that companies still want to use tape for off-site long term archival due to its portability, low cost, and longevity.

One of the technologies that ESG expects to be front-and-center in end-users' discussions is EMC Avamar backup and recovery software. Avamar is a unique software-only backup and recovery solution which allows the end-user to select from a variety of popular server and disk hardware platforms. Avamar software can be deployed on industry standard servers from Dell, HP, IBM, and Sun with backup data stored on internal or SAN disk storage. At smaller remote offices, only Avamar software agents are deployed on the systems to be protected. Larger remote offices and datacenters typically deploy Avamar software on industry standard servers to improve recovery performance.

Avamar solves the performance problems found throughout the backup process. Perhaps one of the most compelling aspects of the Avamar solution is its data de-duplication and global single instance storage technology. The Avamar solution efficiently moves data over existing LAN/WAN bandwidth and reduces the cost of retaining backup data on disk for extended periods of time. Avamar's data reduction technology has the following capabilities:

- Only new and changed sub-file data segments are backed up
- Duplicate data is eliminated at the source and globally across sites and servers - before it is transferred across the network
- Data is compressed after duplicates has been eliminated
- Data is encrypted for secure transfer over the network.

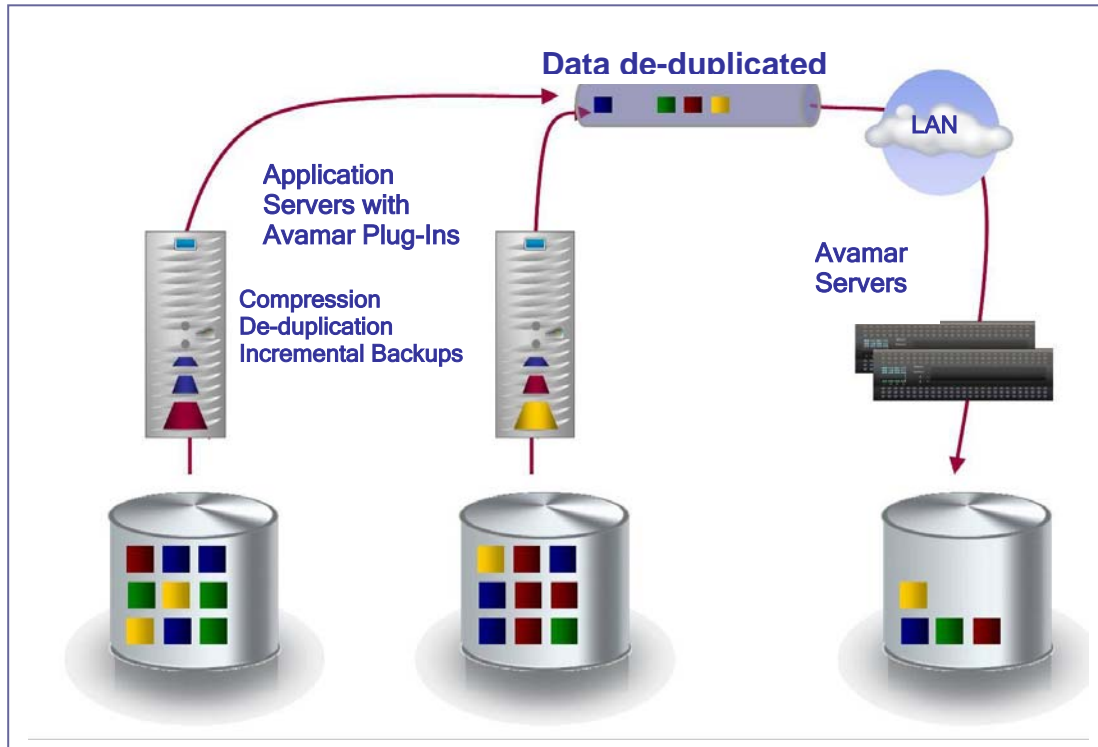
Avamar backs up only new, unique sub-file variable length data segments, which is different than traditional incremental backups. Avamar customers perform an initial full backup just once and subsequent daily backups consist of only new, unique sub-file data segments that are continuously applied forever. In other words, a common legacy backup schedule of *full weekly, incremental daily* changes to *full once, incremental forever* with Avamar. This technique eliminates the issues and delays associated with using one or more incremental backup images during a restore because Avamar provides the equivalent of daily full backups that can be immediately recovered in a single step. Incremental forever backups also eliminate the need to periodically perform full backups, which makes backups run much faster.

Avamar eliminates the need to back up duplicate data by using a powerful set of algorithms that identify duplicate sub-file data segments at the source, before any data is transferred across the network and only backs up what is unique. Avamar also supports data compression algorithms, which eliminate space and redundant byte patterns. EMC refers to this combination of incremental forever backups, redundant data elimination and compression as global data de-duplication and single instance storage (SIS). This patented technology results in significantly improved backup performance, optimized network bandwidth, and efficient use of disk capacity.

To validate these capabilities, ESG performed its own tests of the Avamar software and audited results from ongoing daily backups. For example, an audit of the daily backup of a laptop user indicated that only 4.96 megabytes of capacity was required to protect 1.8 gigabytes of folder data. This represents a data reduction factor of 363x. EMC defines the data reduction factor as the amount of data that is logically backed up daily divided by amount of disk capacity required to store that data¹. While this example is limited in scope, it is illustrative of the powerful data de-duplication capabilities of the EMC Avamar solution.

¹ In this example, a data reduction of 363x = 1,800 MB divided by 4.96 MB.

Figure Two: Data de-duplication and Compression



As Figure Two shows, Avamar software filters primary data at the source - before it is backed up across the network - to ensure that only unique sub-file data segments (represented by the colored boxes) are transferred and stored to disk. The unique data blocks are compressed (represented by the compacted color boxes) and then backed up over an Ethernet LAN or WAN. This dual-process minimizes the amount of data that is ultimately backed up over the LAN to a back-end data repository: There are fewer colored boxes on the right side of the diagram than on the left, illustrating Avamar's claim that it "reduces the required network bandwidth and storage capacity by up to 300x daily".

Other highlights of the Avamar product include an intuitive interface, making it the easiest Enterprise backup solution we have ever tested, the ability to simultaneously perform backup and/or restore processes at any given time, remote replication for disaster recovery and remote site backups, search capability, and integration with legacy tape libraries and leading backup applications.

Avamar is a software-only solution that runs on industry standard servers. The Avamar solution consists of backup-and-recovery software backing up to a centralized SAN-attached or internal disk pool that communicates with software "plug-ins" installed on servers, workstations, or laptops. "Plug-ins" are available for Windows (2003, 2000, NT4, and XP), Linux (Red Hat, SUSE), Unix (Solaris, HP-UX, AIX), and specific applications, including IBM DB2, Microsoft Exchange, Microsoft SQL Server, VMware, Oracle, and NDMP.

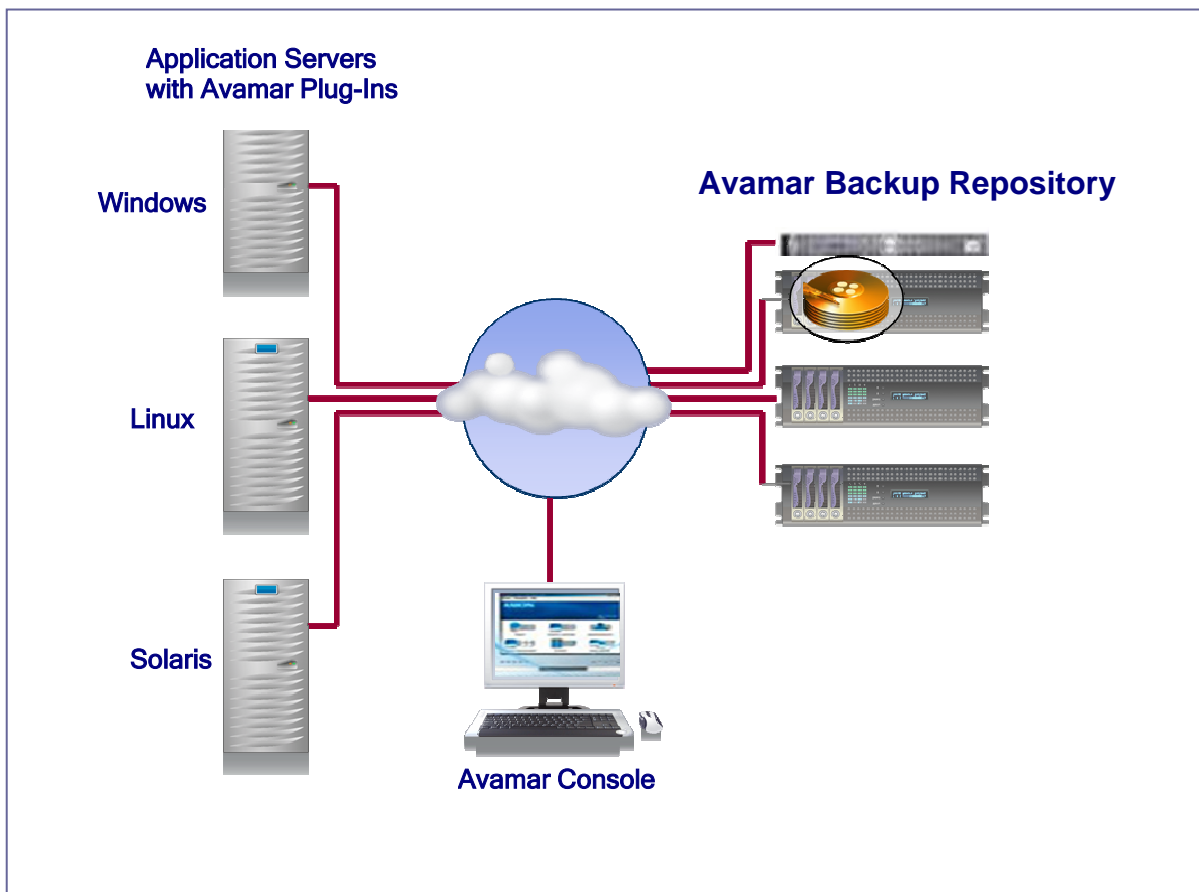
ESG Lab performed hands-on analysis of the Avamar architecture and features in order to validate its value proposition to customers. This ESG Lab report focuses on the fundamentals, including ease of use, rapid backups, and restores from disk and resource conservation. Topics also explored in this report include the use of legacy backup software for archival to tape, replication to a disaster recovery site, and integrated search.

ESG Lab Validation

ESG Lab validated the Avamar data protection solution during two days of hands-on testing. The configuration used during testing is shown in the diagram below. Data residing on Windows, Linux and Solaris servers was backed up using industry standard Intel servers running Avamar software version 3.0.1².

Three servers with internal SCSI hard drives worked together using clustering technology to create a disk-based backup repository. A fourth server running utility node software acted as a traffic cop and was used to configure and control the cluster. Avamar uses the term “3+1” to describe the resulting configuration (3 data nodes + 1 utility node). As shown in Figure Three, backup and restore traffic flowed over an Ethernet-based local area network. Configuration details for the switches, servers, and operating systems used during testing can be found in the Appendix.

Figure Three: The ESG Lab Test Bed

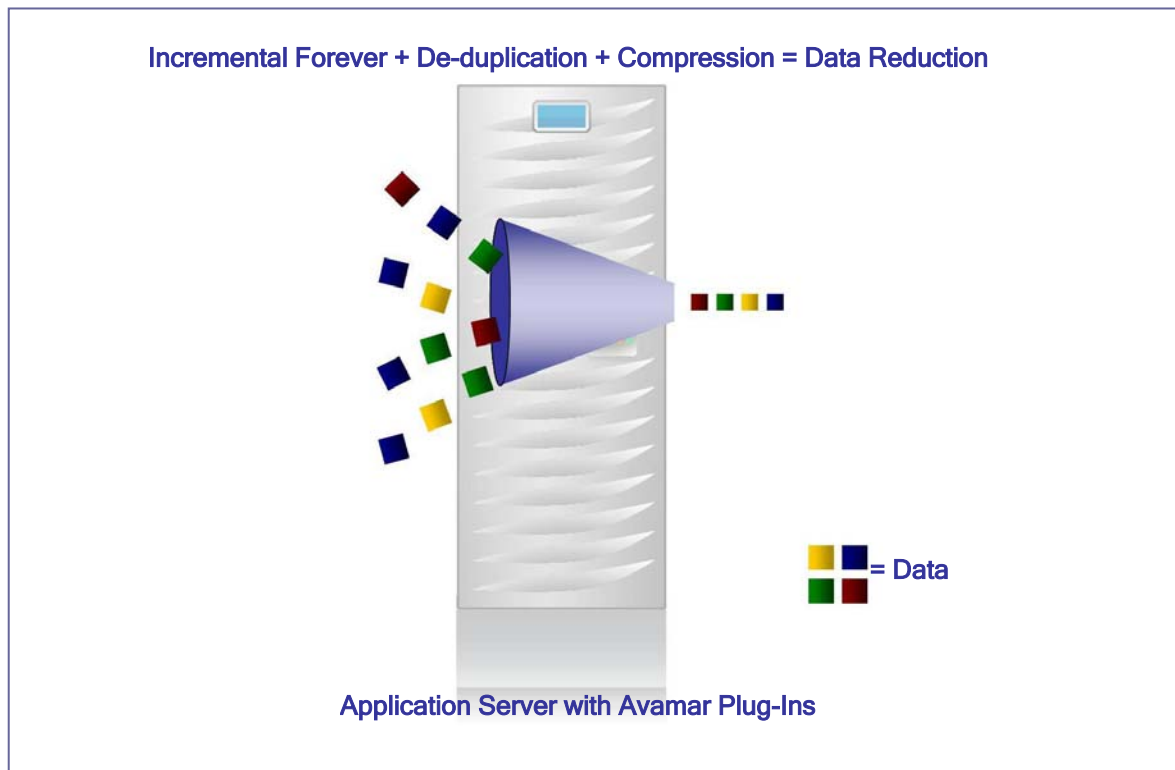


² New features available in a pre-release version of Avamar code (3.5) were also validated during ESG Lab testing. Those results are presented in the sections entitled *Tape Archiving* and *Manager of Managers*.

Avamar Data Reduction

Avamar performs data de-duplication at the source (e.g. server, desktop, laptop) using lightweight software agents and application specific plug-ins. This improves performance from the application server to the backup storage, a capability that is unique in the market.

Figure Four: Data Reduction



ESG Lab Testing

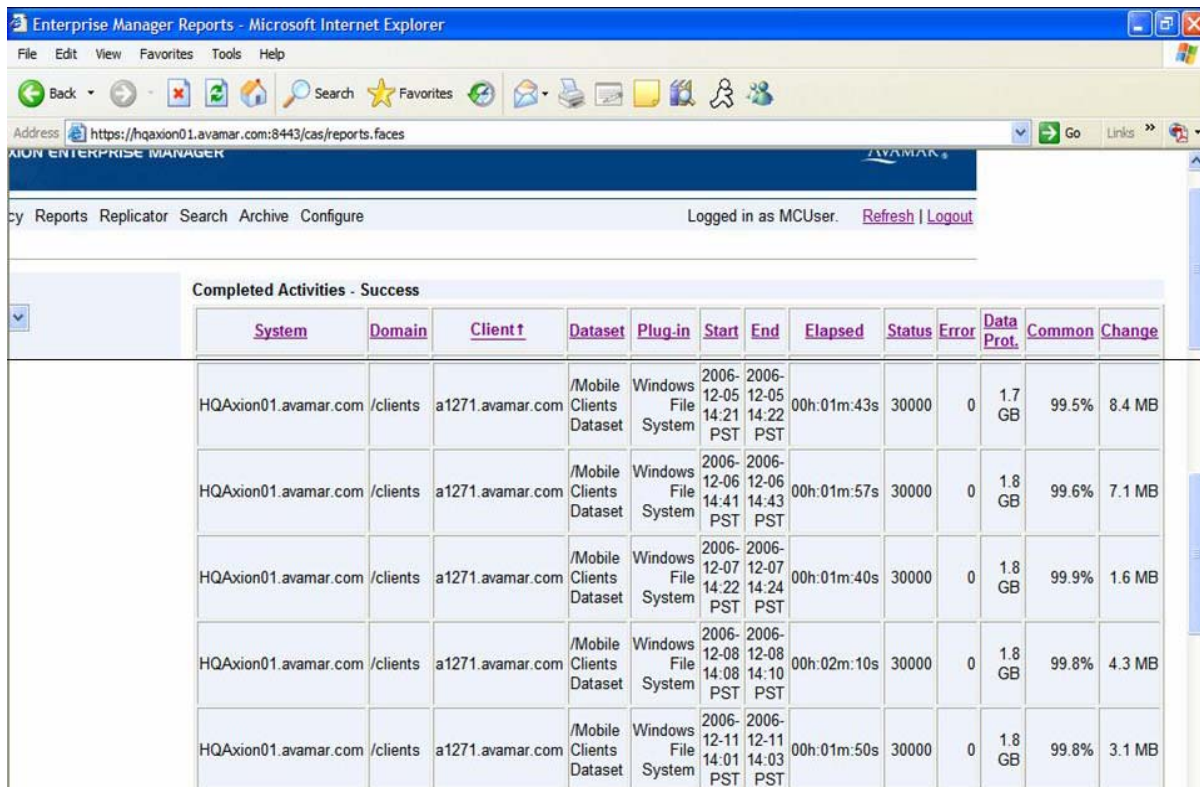
ESG Lab performed hands-on tests and audited historical production backup data to validate capacity and bandwidth efficiency of the EMC Avamar solution. According to EMC, during an initial backup of file data, the amount of data stored to disk is typically reduced by 70% in real-world deployments. Even more impressive is what happens when backups are retained over time. As new daily backups are taken, the amount of common data across backups increases, ultimately reducing the required network bandwidth and backup storage required by up to 300x.

Avamar data reduction technology was investigated initially by examining what happens when a PowerPoint presentation is backed up. An initial backup of a 10 MB presentation was followed by a second backup after the graphics and text in a single slide was modified. The second backup completed in a matter of seconds and the Avamar management interface was used to learn that only 3.7% of the data had changed. In other words, the second backup provided an effective full backup of 10 MB using only 370 KB of data on disk.

Having witnessed the benefits of data reduction over the course of two backups, ESG Lab audited a series of Avamar backup logs to determine the cumulative data reduction benefits that can be achieved over longer periods of time based on real-life application data. Avamar software is used internally by the Avamar team at EMC to protect all of their data. The data stored on servers used by the engineering, finance, and operations groups, as well as on desktops and even laptops, is protected using Avamar software.

For example, Figure Five shows a series of daily backup logs for laptop users. ESG Lab was impressed not only by the fact that backups run in a matter of minutes, but even more so by the reduction in the amount of backup capacity required for each daily backup.

Figure Five: Avamar GUI



System	Domain	Client t	Dataset	Plug-in	Start	End	Elapsed	Status	Error	Data Prot.	Common	Change
HQAxion01.avamar.com	/clients	a1271.avamar.com	/Mobile Clients Dataset	Windows File System	2006-12-05 14:21 PST	2006-12-05 14:22 PST	00h:01m:43s	30000	0	1.7 GB	99.5%	8.4 MB
HQAxion01.avamar.com	/clients	a1271.avamar.com	/Mobile Clients Dataset	Windows File System	2006-12-06 14:41 PST	2006-12-06 14:43 PST	00h:01m:57s	30000	0	1.8 GB	99.6%	7.1 MB
HQAxion01.avamar.com	/clients	a1271.avamar.com	/Mobile Clients Dataset	Windows File System	2006-12-07 14:22 PST	2006-12-07 14:24 PST	00h:01m:40s	30000	0	1.8 GB	99.9%	1.6 MB
HQAxion01.avamar.com	/clients	a1271.avamar.com	/Mobile Clients Dataset	Windows File System	2006-12-08 14:08 PST	2006-12-08 14:10 PST	00h:02m:10s	30000	0	1.8 GB	99.8%	4.3 MB
HQAxion01.avamar.com	/clients	a1271.avamar.com	/Mobile Clients Dataset	Windows File System	2006-12-11 14:01 PST	2006-12-11 14:03 PST	00h:01m:50s	30000	0	1.8 GB	99.8%	3.1 MB

Backup logs as shown above were audited for a variety of server and data types. The last three values shown on the right for each completed backup (Data Protected, Common, Changed) were used to calculate the data reduction factors shown in the following table.

Table One: Data Reduction Benefits

Data Type	Changed Data Daily	Data Reduction Factor
Laptop	0.28%	363x
Engineering Server	0.2%	611x
Exchange Server	0.5% to 4%	25 to 200x

Avamar End-User Interview

Armed with hands-on experience and an audit of the Avamar backup logs, ESG Lab interviewed an Avamar customer to validate the data reduction benefits of Avamar software over time. The customer worked for a construction firm with about 12 TB of Oracle, email and file data. Veritas NetBackup and tape libraries were being used before the introduction of Avamar software. Full backups run once a week were beginning to take all weekend long. Backups frequently failed and portions had to be re-run on Monday. Restore requests that came in over the weekend or on Monday often had to be deferred until

Tuesday. Handling over 60 TB of on and off-site DLT and LTO tape cartridges was becoming untenable - a fact that wasn't fully realized until tape was phased out after the installation of Avamar software.

Avamar software was installed and initially deployed to protect project data living on a file server. It took some time to convince the database administrators, but now all of the company's data is backed up using Avamar. A month's worth of backups are kept on disk for quick and reliable restores. The user interface made a great first impression. As a matter of fact, ESG Lab heard something we thought we'd never hear about backup software: "getting it going was quite a bit of fun." It's easy and it works reliably. Running full backups is as simple as a right mouse click. Backups that run in parallel every night are done in a couple hours. Restores of lost or corrupt file are fast and simple. Now that backups aren't running all weekend, Saturdays can be used for maintenance.

Avamar logs as shown in Figure Five are checked each morning to verify that last night's backup completed successfully. The percentage of common data across all data types is typically 99.7%. That implies that a daily effective full backup of 12 terabytes of company data takes only 36 gigabytes of capacity on disk for an effective data reduction factor of 333x.

Based on the results of ESG Lab testing, interviews with EMC Avamar customers, and an audit of backup logs, ESG is confident that customers can run daily full backups and achieve a data reduction factor in the 300x range while retaining backups on disk for six months.

Why This Matters

EMC Avamar's patented data reduction technology significantly changes the economics and efficiency of performing backups. EMC Avamar software claims to reduce the amount of data transferred and stored for backups by up to 300x daily. Tests demonstrated the ability to perform daily full backups of 1.8GB on a corporate user's laptop using an average of only 4.9MB of actual disk capacity to store the entire 1.8GB of backup data; that's a data reduction factor of 363:1. Avamar Data Reduction technology puts disk and tape on equal footing in terms of cost while increasing data availability by eliminating the need for traditional incremental backups.

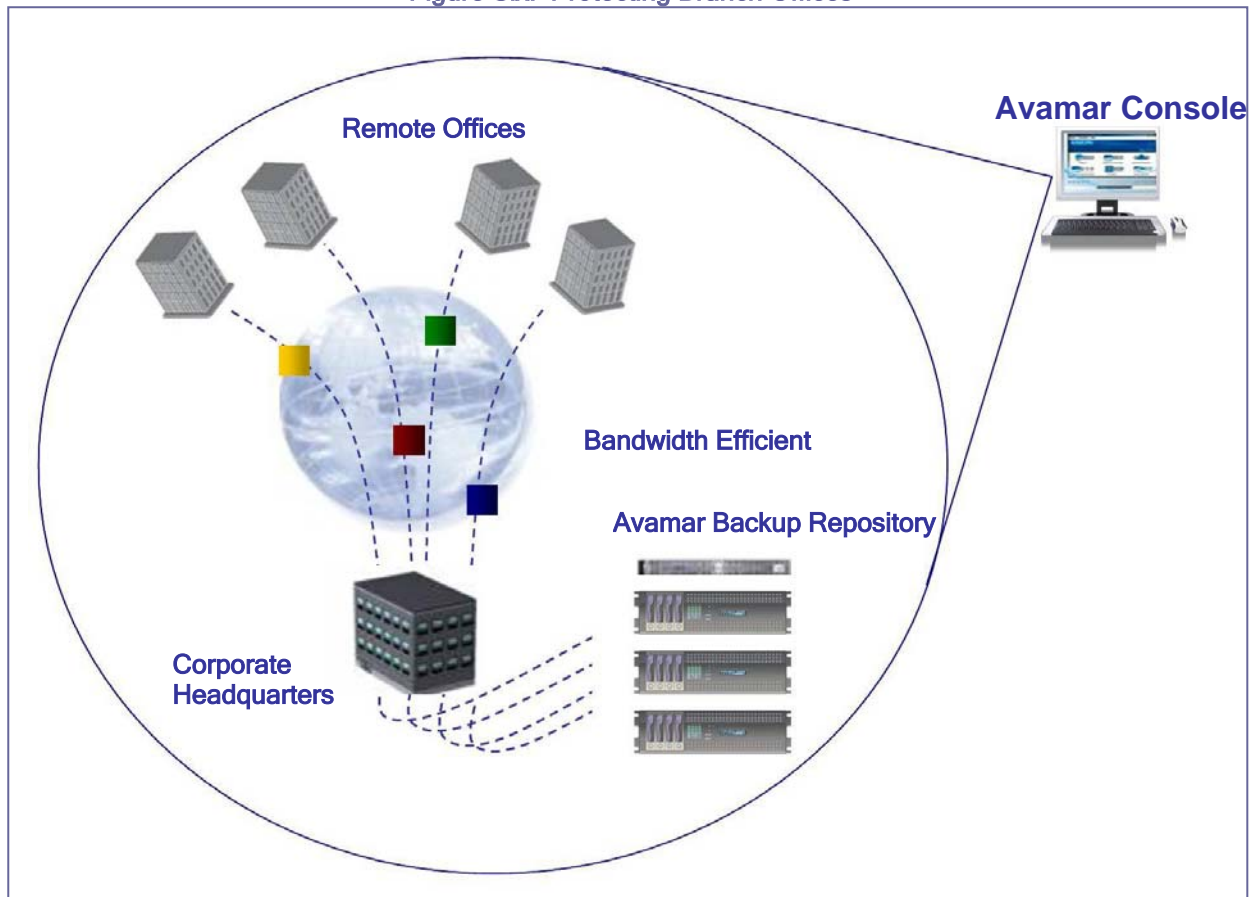
Another important benefit is that Avamar is one of the only backup solutions that perform data de-duplication at the client source, which adds efficiency throughout the entire backup process. As a result, Avamar is the only solution that solves the backup performance problem entirely. ESG has spoken to customers that have reduced their backups from eight hours down to literally 15 minutes using Avamar. ESG also spoke with one customer who said that they could not afford the time or the money required to perform remote backups using traditional solutions. However, using Avamar for this purpose was fast and cost-effective.

Remote Replication

In addition to providing disk-based backup protection within a local data center and at remote offices, Avamar supports disaster recovery and off-site archival by replicating backup data over a wide area network to a remote data center using the Avamar Replicator software. Data reduction and Replicator capabilities ensures that only changes since the last backup, comprised of unique and compressed sub-file variable length data segments, are replicated. This is extremely valuable, reducing the cost of bandwidth and improving the performance of backups and replication over the WAN.

Backups of a single data center can be replicated to a remote site using Avamar Replicator for disaster recovery in a one-to-one configuration (one local archive, one remote replica). Avamar Replicator can also be configured in a many-to-one configuration, as shown in Figure Six. Note that many branch offices are replicated to a centrally managed Avamar system. The Avamar systems at the branch offices can be managed by a single administrator at the corporate office.

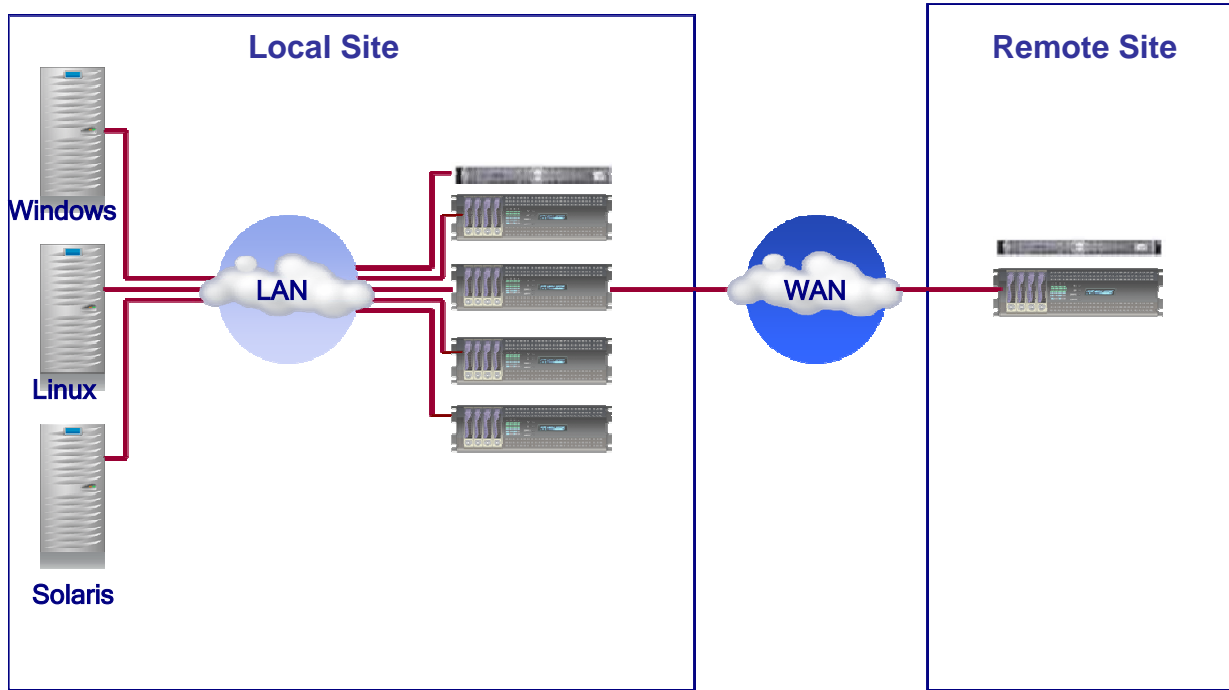
Figure Six: Protecting Branch Offices



ESG Lab Testing

A single-node Avamar system was configured as shown in Figure Seven to test remote replication of the system used earlier in ESG Lab testing. An internal GigE network was used instead of a WAN during this phase of ESG Lab testing.

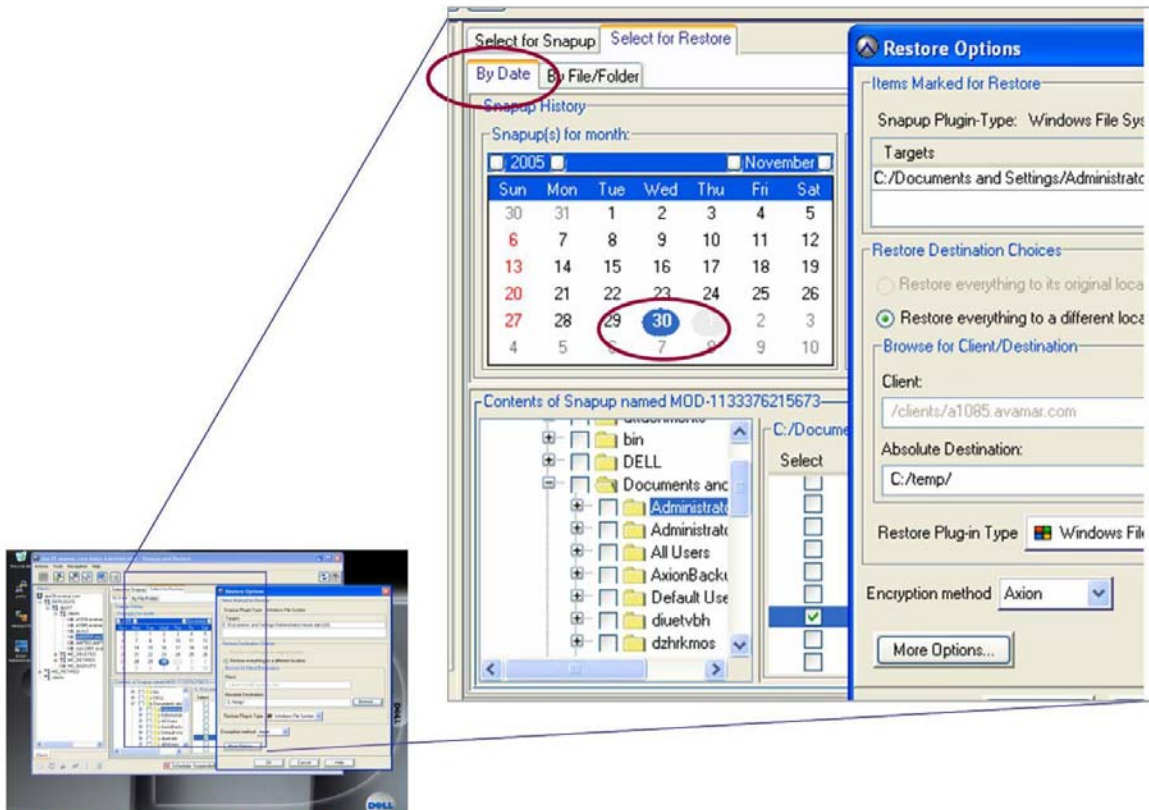
Figure Seven: Testing Avamar Replicator



Although installation of Avamar Replicator at a customer site is performed by trained professional services staff, ESG Lab performed the installation with the guidance of an EMC expert. The installation of the remote system with approximately 1 TB of internal SCSI drive capacity began with an Avamar specific Linux distribution and was completed in about 30 minutes using a series of command line options. Within an hour ESG had configured a replication policy to run at 6:00 PM nightly. A restore from a remote replica was then performed using another cluster in the lab.

As shown on the left side of Figure Eight, the remote restore began with the selection of a Dell server. Using the panel shown to the right, a single file (ntuser.dat.LOG) was used to restore a file from the remote site to a different location (c:/temp on a local Windows server). The restore completed without error in less than three minutes.

Figure Eight: Restoring a Remote File



Why This Matters

Avamar Replicator software is a plug-in software module for EMC Avamar software. It easily and cost effectively enables companies to make copies of data to remote locations for disaster recovery. Typically, companies use expensive technologies to perform remote mirroring to protect their most mission-critical data. However, the majority of their data is not protected. Using Avamar Replicator software, companies can perform asynchronous replication of all their data including internal, DAS, SAN, and NAS storage, leveraging the advantages of the Avamar global data de-duplication and single instance storage technology.

Ease of Use

Configuring and managing legacy backup software can be difficult and often requires specialized training and skill. The routine tasks that need to be managed with a backup and recovery solution include:

- Installing software plug-ins, often referred to as agents, on the servers to be backed up.
- Configuring backup software to use tape and/or disk for storing recovery images.
- Choosing which applications and files to back-up.
- Creating backup schedules.
- Monitoring backup completion status.
- Performing file, application, system, and data center level restores.

ESG Lab Testing

EMC Avamar software is installed and configured by an EMC customer service representative. A morning or afternoon is typically set aside for installation, initial configuration, and a quick orientation lesson at the customer site. ESG Lab testing simulated that customer experience using a pre-wired Avamar system.

Installing Plug-Ins

The first step in configuring Avamar is the download and installation of a thin software plug-in onto each system to be backed up. Installation of software plug-ins is typically performed later to protect a new or upgraded server. Avamar offers plug-ins for a variety of operating systems including Windows (2003, 2000, NT4, and XP), Linux (Red Hat, SUSE), Unix (Solaris, HP-UX, AIX) and specific applications including IBM DB2, Microsoft Exchange, Microsoft SQL Server, VMware Oracle and NDMP³.

During ESG Lab testing, the Avamar software plug-ins were downloaded using a browser onto the servers to be protected. Directing the browser to the host name of the Avamar utility node brought up a menu of installation, configuration, and maintenance options. Installing the plug-in was quick and easy and conformed to the installation standards of each server (setup.exe for Windows, rpm for Linux, pkgadd for Solaris). Ten minutes after testing had begun, the Windows, Linux, and Solaris servers were equipped with Avamar plug-ins and ready for the first backup.

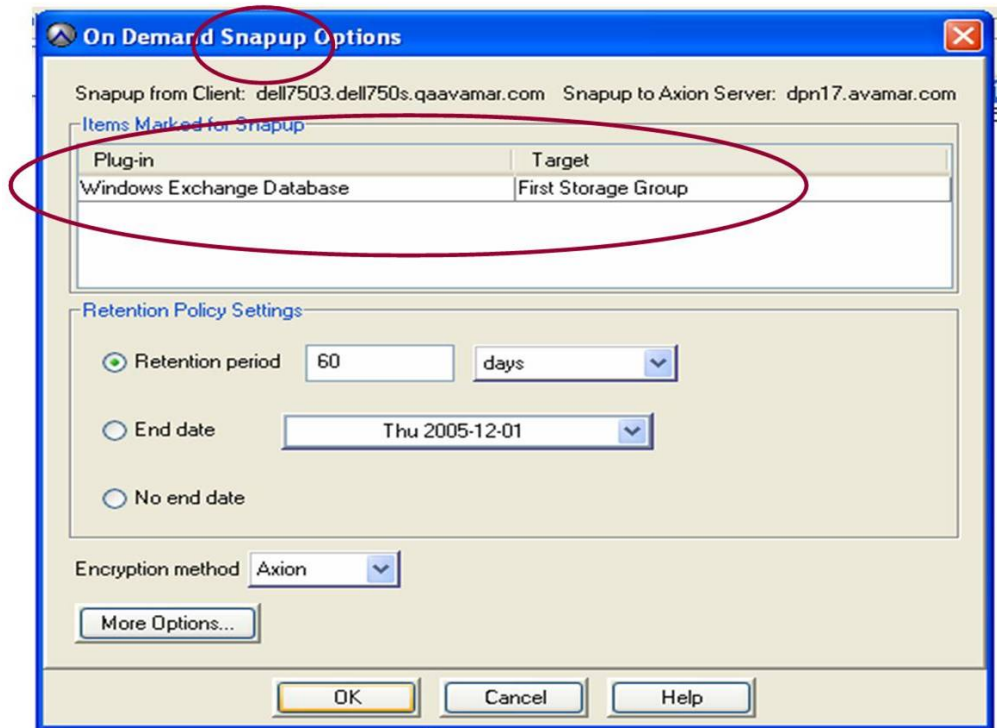
Backup

With literally no guidance from Avamar, ESG Lab used the user interface running in a browser to select the server to be backed up (Dell7503), the type of data to be backed up (Windows Exchange Database), and the storage group to be backed up (First Storage Group). As shown in Figure Nine, finding and selecting the server and data to be backed up felt very intuitive. The tree structure felt much like the Windows Network Neighborhood and Explorer interfaces. Once selection was complete, a backup on demand was easily executed using an intuitive right mouse click.

Note that Avamar uses the term *Snapup* instead of *Backup* in the user interface. Although the term *Snapup* may not be familiar to users, ESG Lab appreciates the fact that the name is an interesting play on words that says exactly what is happening. Avamar stores every full backup as a network mountable view of the system, like a snapshot, which provides fast and space efficient protection for individual files, directories, or entire systems - without the need to apply incremental backups to a full baseline image.

³ NDMP is industry standard protocol for backing up network attached file systems.

Figure Nine: Backing Up an Exchange e-mail Storage Group

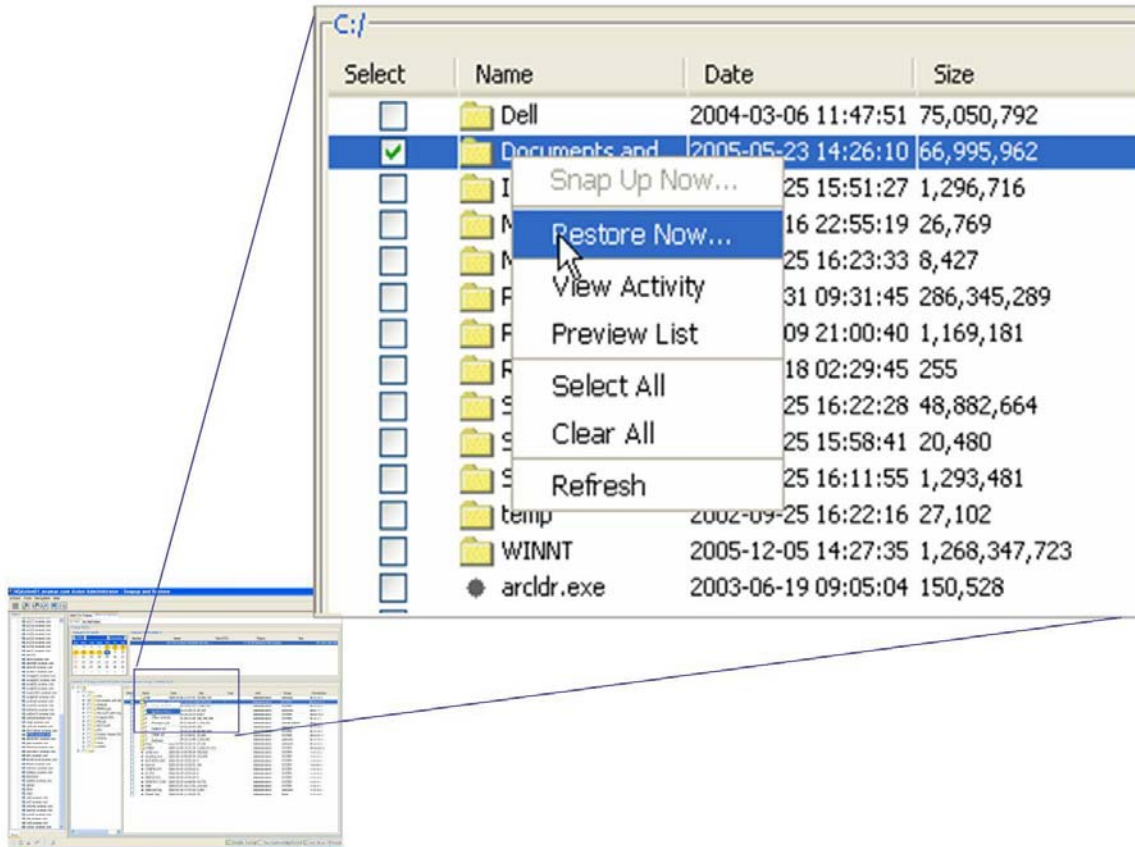


Restore

After file systems and application data on the Windows, Linux, and Solaris servers had been backed up, ESG Lab next used Avamar software to perform restores. Once again, the user interface felt very intuitive and with no guidance from EMC, ESG Lab restored a file on a Windows server, as shown in Figure Ten. Note that in addition to the server and data navigation trees used in the *Snapup* user interface, the restore interface includes a calendar which is used to pick the backup to restore from. The yellow circles clearly indicate on which days backups had been performed. The backup(s) performed on the date selected appear on the right. The file system navigation window is then used to select directories or files to restore using an intuitive right click and the *Restore Now* option. When restoring Exchange mailboxes and Oracle tables, the same intuitive look and feel was observed.

Single file and single directory restores completed nearly instantaneously due to the fact that Avamar restores from disk don't incur the performance latency associated with waiting for a tape to load. Files restored were compared to files backed up and were observed to be error-free.

Figure Ten: Restoring a File



Why This Matters

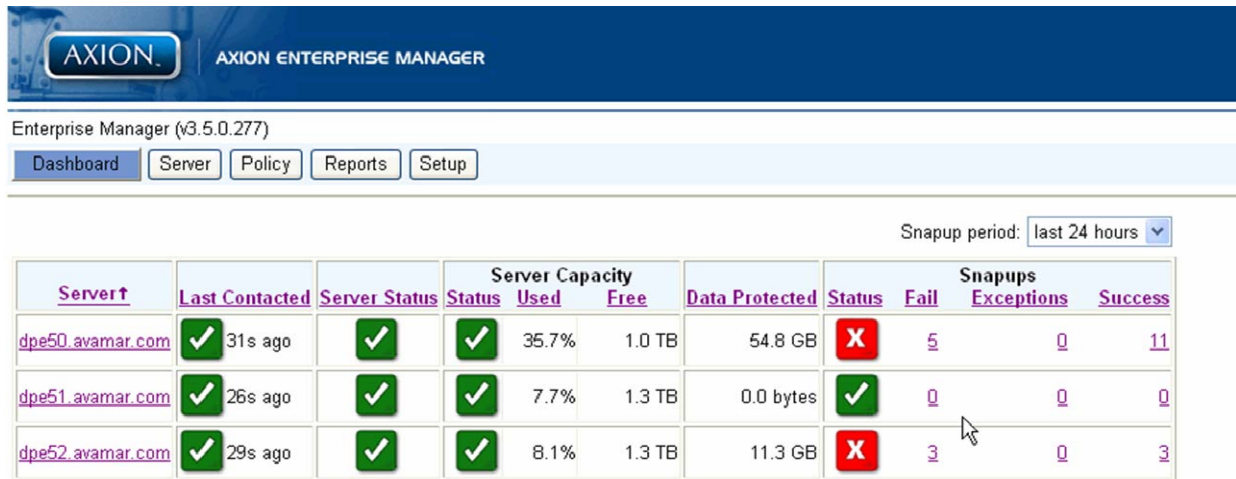
Legacy backup and recovery software architectures can be decades old and often suffer from “feature-creep.” With too many ways to get the same things done, and a long history to maintain, the user interface can appear daunting at first. In many IT shops, ESG has found that only experts who have worked with the backup and recovery software for years know how to make changes.

ESG Lab believes that it is obvious that Avamar was built from a clean slate with simplicity and ease of use in mind. Avamar is easy to use and intuitive. In our view, it is the easiest Enterprise backup software we have ever used.

Manager of Managers

ESG Lab verified the ability to manage multiple repositories from a single user interface, as shown in Figure Eleven. Using a simple web-based interface, it is possible to determine the health of multiple EMC Avamar repositories at a single glance. From this single interface, ESG Lab was able to point and click on a specific repository and further diagnose a potential problem. This functionality, which ESG refers to as a manager of managers, makes it easier for companies to standardize on Avamar for backup and recovery services across the enterprise.

Figure Eleven: The Avamar Manager of Managers



Enterprise Manager (v3.5.0.277)

Dashboard | Server | Policy | Reports | Setup

Snapup period: last 24 hours

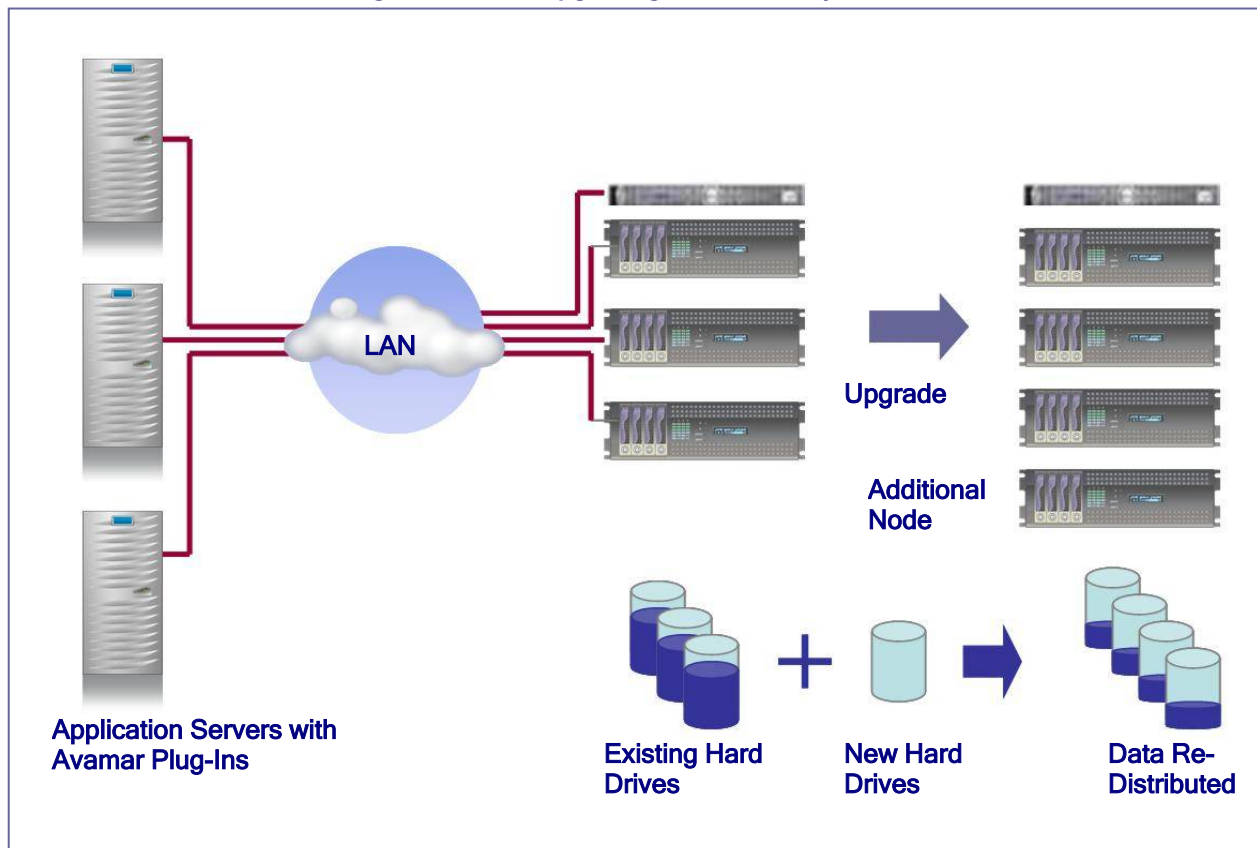
Server†	Last Contacted	Server Status	Server Capacity			Data Protected	Status	Snapups		
			Status	Used	Free			Fail	Exceptions	Success
dpe50.avamar.com	✓ 31s ago	✓	✓	35.7%	1.0 TB	54.8 GB	✗	5	0	11
dpe51.avamar.com	✓ 26s ago	✓	✓	7.7%	1.3 TB	0.0 bytes	✓	0	0	0
dpe52.avamar.com	✓ 29s ago	✓	✓	8.1%	1.3 TB	11.3 GB	✗	3	0	3

Upgradeability

Avamar implements a grid architecture that can be upgraded simply by adding servers to the grid. Hard drives within each additional server increase the capacity of the total system. Adding servers also adds processors, bandwidth, and memory which increases total system performance. As servers are added, existing backup data is automatically rebalanced and spread evenly across the Avamar servers. As a result, backup workloads are always spread evenly across all nodes, which also increases system performance.

Avamar supports a concept that ESG refers to as a single level of management. Regardless of the size of an Avamar grid, it is manageable as a single system. Whether the system is built from three nodes or 20 nodes, it can be managed from a single user interface.

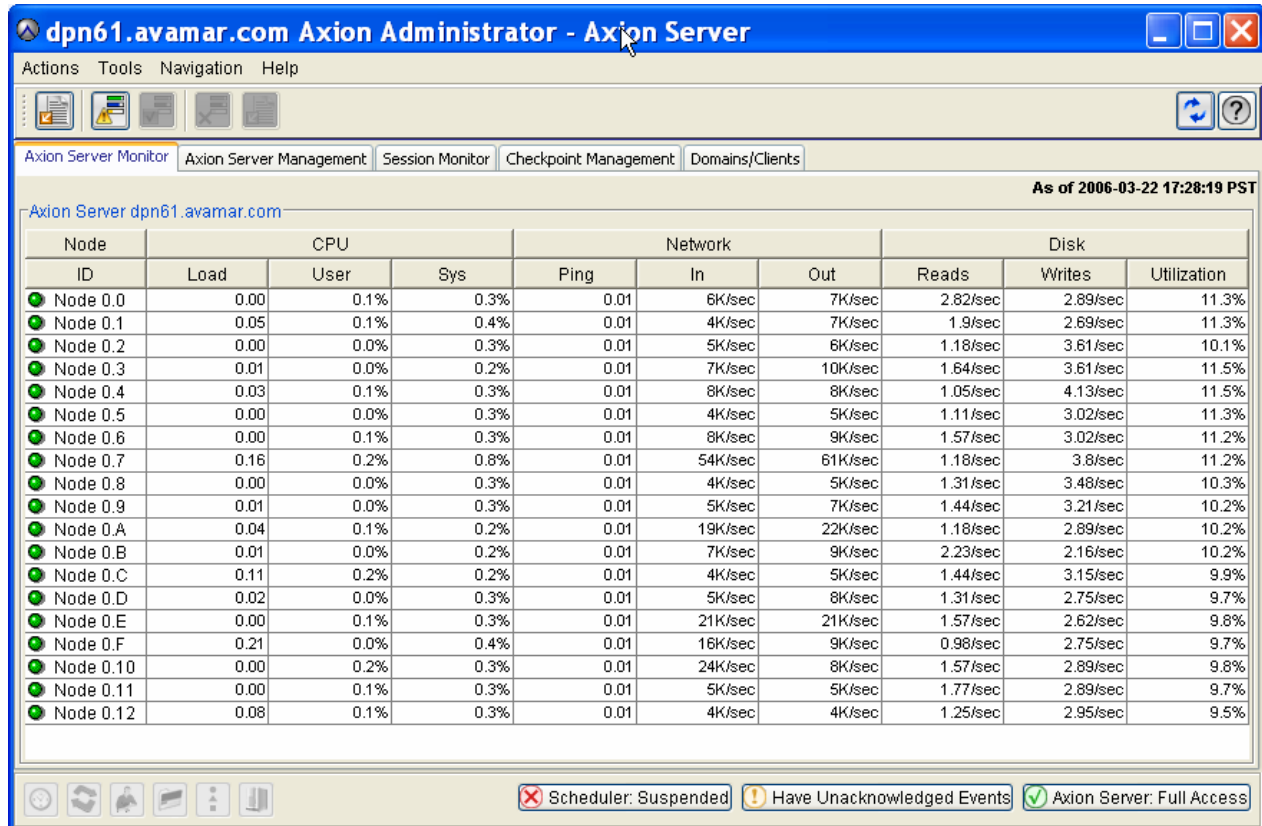
Figure Twelve: Upgrading an Avamar System



ESG Lab Testing

A fourth server with approximately 1 TB of internal SCSI drive capacity was added to the three node system during ESG Lab validation testing, as shown in Figure Twelve. Using Avamar’s terminology, this was an upgrade from a 3+1 to a 4+1 configuration. Although Avamar upgrades at a customer site are performed by a trained customer service representative, ESG Lab performed the upgrade with Avamar guidance. Once the installation had completed and the new server had joined the cluster, existing used capacity was automatically re-balanced over the cluster. Re-balancing progress was observed on the user interface. ESG Lab observed a 19 node system running in the lab at Avamar. The maintenance interface screen in Figure Thirteen shows the status of that 19 node cluster.

Figure Thirteen: A 19 Node EMC Avamar System



As of 2006-03-22 17:28:19 PST

Node ID	CPU			Network			Disk		
	Load	User	Sys	Ping	In	Out	Reads	Writes	Utilization
Node 0.0	0.00	0.1%	0.3%	0.01	6K/sec	7K/sec	2.82/sec	2.89/sec	11.3%
Node 0.1	0.05	0.1%	0.4%	0.01	4K/sec	7K/sec	1.9/sec	2.69/sec	11.3%
Node 0.2	0.00	0.0%	0.3%	0.01	5K/sec	6K/sec	1.18/sec	3.61/sec	10.1%
Node 0.3	0.01	0.0%	0.2%	0.01	7K/sec	10K/sec	1.64/sec	3.61/sec	11.5%
Node 0.4	0.03	0.1%	0.3%	0.01	8K/sec	8K/sec	1.05/sec	4.13/sec	11.5%
Node 0.5	0.00	0.0%	0.3%	0.01	4K/sec	5K/sec	1.11/sec	3.02/sec	11.3%
Node 0.6	0.00	0.1%	0.3%	0.01	8K/sec	9K/sec	1.57/sec	3.02/sec	11.2%
Node 0.7	0.16	0.2%	0.8%	0.01	54K/sec	61K/sec	1.18/sec	3.8/sec	11.2%
Node 0.8	0.00	0.0%	0.3%	0.01	4K/sec	5K/sec	1.31/sec	3.48/sec	10.3%
Node 0.9	0.01	0.0%	0.3%	0.01	5K/sec	7K/sec	1.44/sec	3.21/sec	10.2%
Node 0.A	0.04	0.1%	0.2%	0.01	19K/sec	22K/sec	1.18/sec	2.89/sec	10.2%
Node 0.B	0.01	0.0%	0.2%	0.01	7K/sec	9K/sec	2.23/sec	2.16/sec	10.2%
Node 0.C	0.11	0.2%	0.2%	0.01	4K/sec	5K/sec	1.44/sec	3.15/sec	9.9%
Node 0.D	0.02	0.0%	0.3%	0.01	5K/sec	8K/sec	1.31/sec	2.75/sec	9.7%
Node 0.E	0.00	0.1%	0.3%	0.01	21K/sec	21K/sec	1.57/sec	2.62/sec	9.8%
Node 0.F	0.21	0.0%	0.4%	0.01	16K/sec	9K/sec	0.98/sec	2.75/sec	9.7%
Node 0.10	0.00	0.2%	0.3%	0.01	24K/sec	8K/sec	1.57/sec	2.89/sec	9.8%
Node 0.11	0.00	0.1%	0.3%	0.01	5K/sec	5K/sec	1.77/sec	2.89/sec	9.7%
Node 0.12	0.08	0.1%	0.3%	0.01	4K/sec	4K/sec	1.25/sec	2.95/sec	9.5%

Scheduler: Suspended | Have Unacknowledged Events | Axion Server: Full Access

Why This Matters

Managing many systems is harder than managing a single system. This fact is known all too well by backup administrators who are struggling with too many backup media servers. A centrally managed Avamar system with scalable processing power, bandwidth, and capacity can be used to replace many media servers. Consolidation of multiple legacy media servers into a centrally managed Avamar repository reduces complexity, saves time, and reduces cost.

Avamar File System (AvFS)

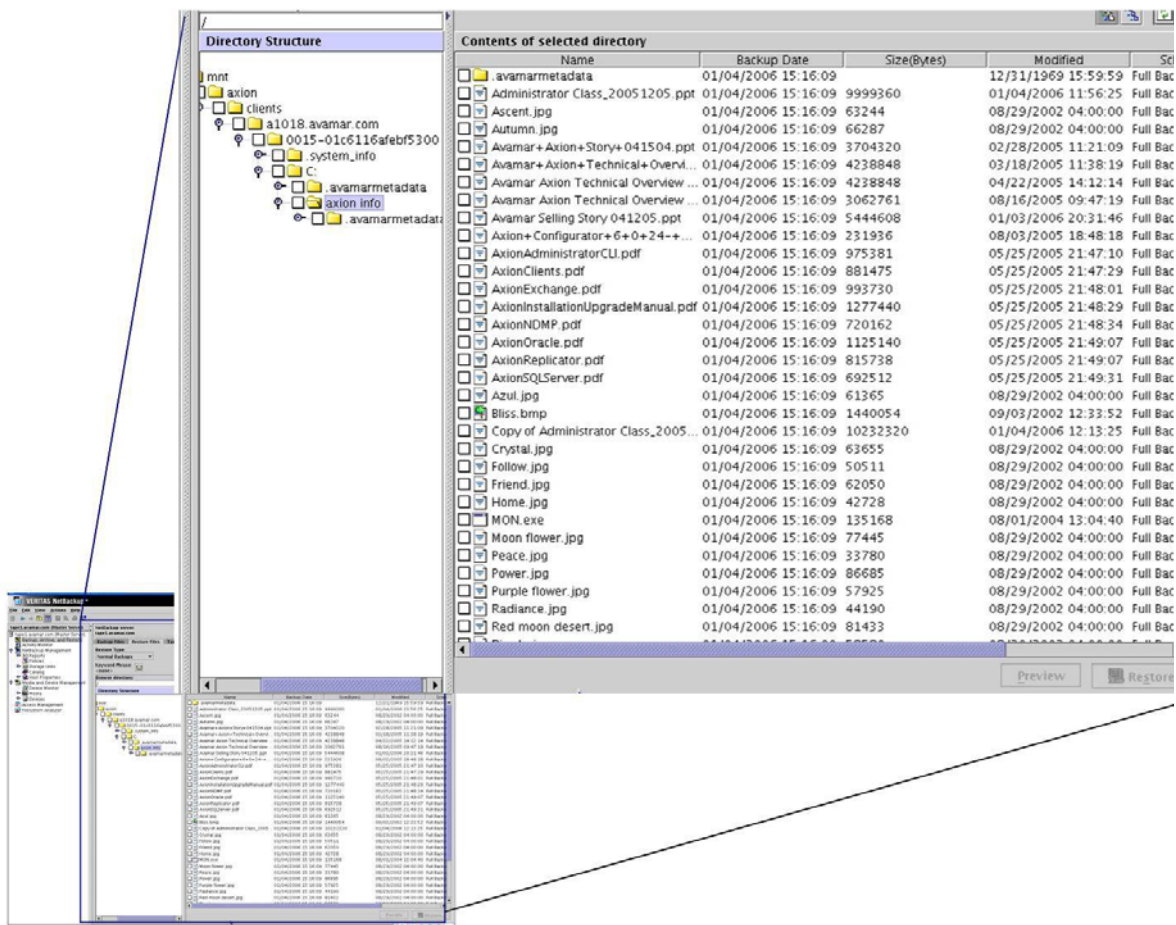
Avamar stores backup data on an internally managed file system. As such, Avamar has file-level awareness. Avamar has wisely exposed this file system to users and the value that this brings to customers is compelling. For example, the file system interface can be used to:

- Output to Tape - A Windows or Linux file system view of Avamar backups can be accessed by any standard 3rd party backup software solution to export disk-based copies of backup data to tape in native format.
- Search - The file system view of Avamar data can be indexed and scanned by enterprise-class search engines, such as Google and FAST.
- Direct restore - Administrators can utilize the file system view of backups to perform direct restores of data using familiar operating system level copy commands and utilities.

User Restores

Since Avamar stores backup data on a file system, users can now restore their own files. This is exactly the same as users restoring their own files from snapshot directories within NAS storage systems. By enabling users to restore their own files, companies can reduce the number of help desk calls and interventions by the IT department to recover files. Avamar provides security tools, permissions, and privilege levels so users are limited in what they can restore.

Figure Fourteen: File System View Enabled Tape-out

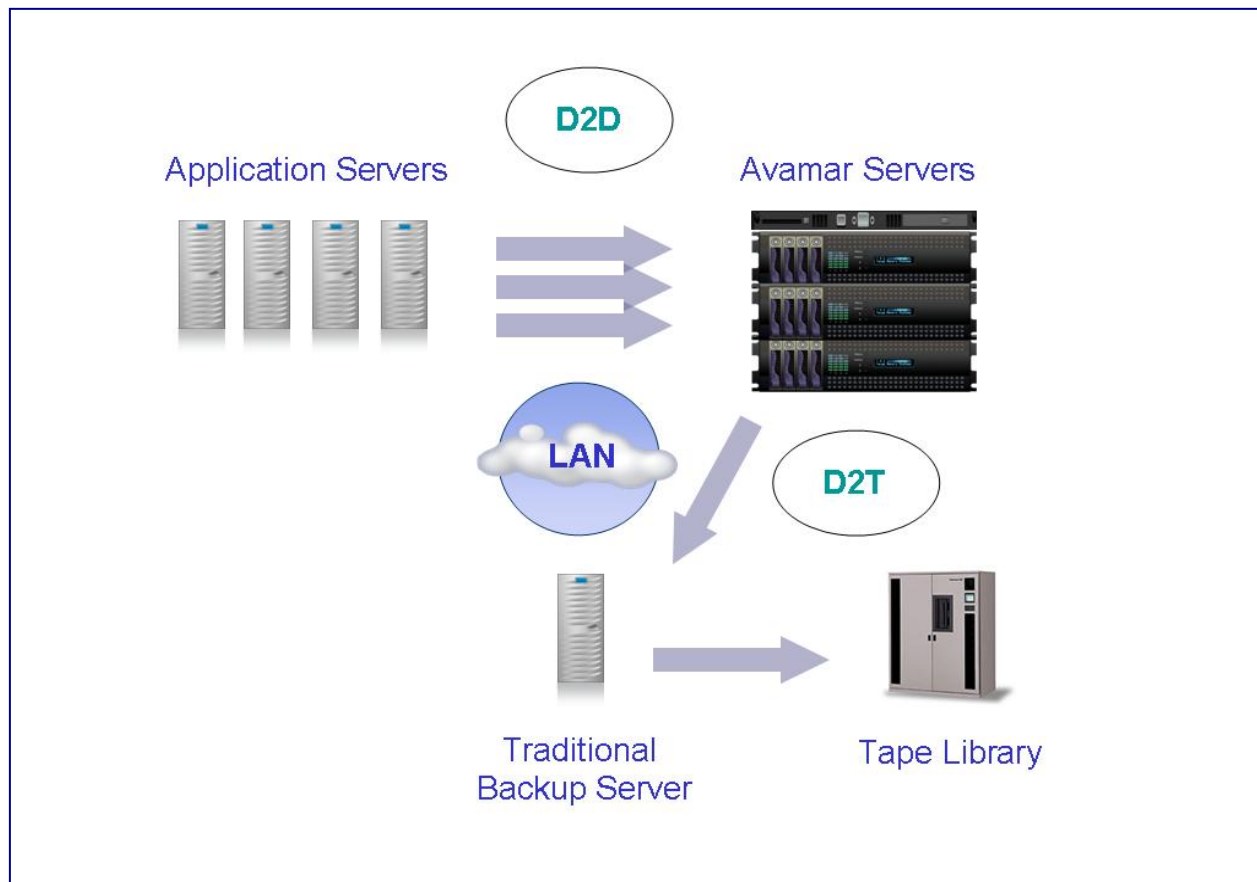


Tape Integration

ESG has found that customers want to back up data to both disk and tape. Avamar has recently implemented integration with tape systems. Since Avamar stores its backup data as a file system, any traditional backup software (e.g., Legato, TSM, and Veritas) can archive to a tape or disk archive system (e.g. EMC Centera) in native format so that it can be recovered directly from tape via backup software if desired. This is an elegant and simple approach for using the Avamar backup to disk solution with traditional backup software and tape.

This enables customers to implement a strategy that is commonly referred to as Disk to Disk to Tape (D2D2T). The way it works is simple: applications with data stored on disk (first **D**) are backed up to the Avamar system on a regular basis (second **D**). Since the Avamar data is stored as a file system, traditional backup software (e.g., Legato, TSM, Veritas) can be configured to periodically perform full archives to tape (**T**).

Figure Fifteen: Avamar D2D2T



To the traditional backup software, the Avamar file system looks like any other server or NAS storage system and treats it as such. The process for backup and restore from tape is easy and transparent.

Another thing customers may want to consider is reducing the backup or media servers required by their traditional backup software. Instead, they can use Avamar to back up their application servers and then use traditional backup software to store Avamar backups onto tape. This would allow customers to leverage the performance, capacity and bandwidth optimization, and ease of use of Avamar, and at the same time continue to use their traditional backup solution for their tape processes. Customers could potentially remove dozens or even hundreds of backup or media servers.

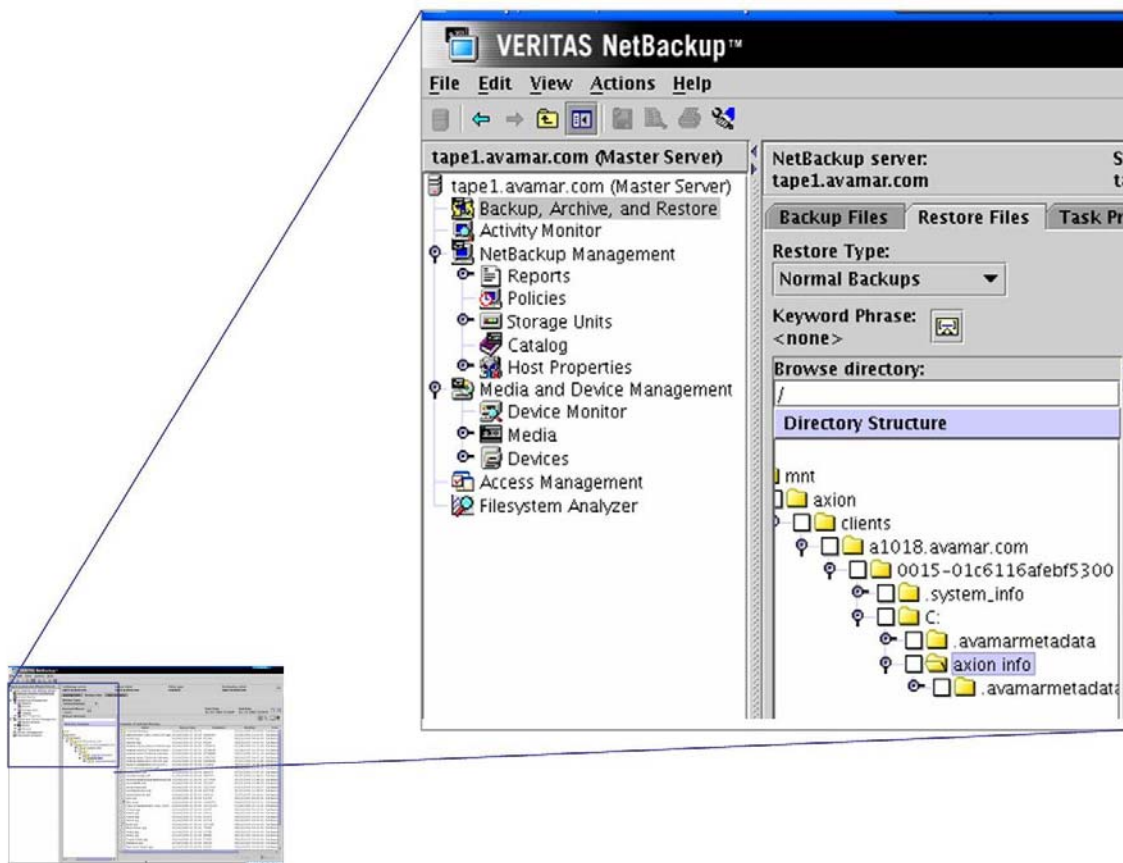
ESG Lab Testing

Backups to tape using legacy backup software are run against a network attached file system view of an Avamar repository. ESG Lab tested the Avamar file system view and observed Veritas NetBackup as it ran tape-based backups and restores from an Avamar repository.

The file system view presented by Avamar software was accessed as a network mounted drive on a Windows and a Linux server during ESG Lab testing. The directory structure made sense - under each server name there were a series of date- and time-stamped directories that represented each logical full backup performed by Avamar software. A directory shortcut aptly named "latest" was used to find the latest backup performed on a Windows server during ESG Lab testing. ESG Lab manually navigated directory trees with Windows Explorer and Linux cd commands to find, copy, and verify a few files that had been copied to a temporary directory.

Manual verification of the file system view was followed by a backup to tape using Veritas NetBackup. As shown in Figure Sixteen, configuring Veritas to back up and restore logically complete full backup images is the same as backing up an application server. Instead of pointing the backup software to a specific server, a directory on the Avamar server for each protected application is used.

Figure Sixteen: Veritas NetBackup using the Avamar File View



Why This Matters

The Avamar File System (AvFS) allows users to restore their own data without requiring the assistance of IT staff. This enables users to be more efficient and frees up the IT staff from dealing with recovering files. AvFS is extremely compelling because it extends the use of the data already stored by Avamar beyond backup and recovery. Users can now search and access Avamar data as a digital archive. This allows for multiple users of the same data.

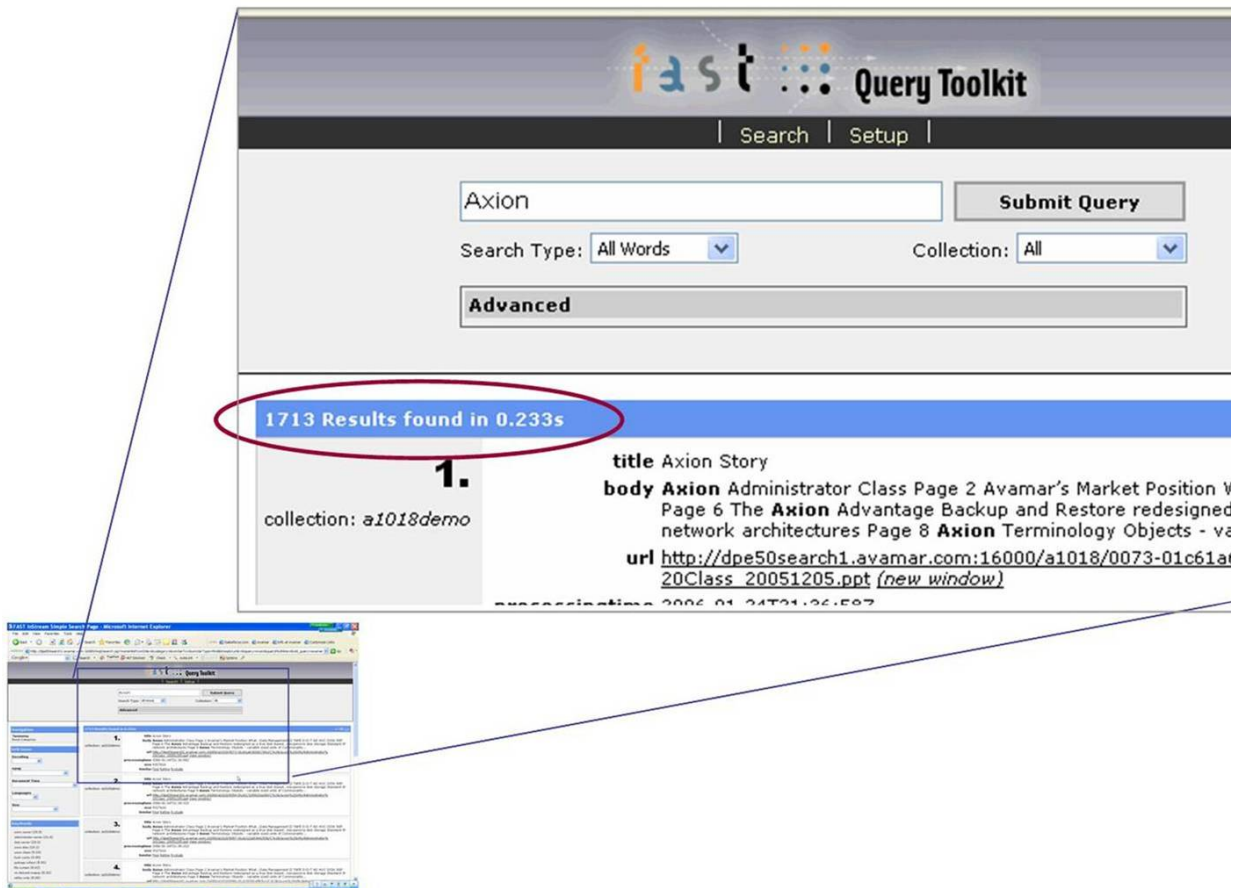
Additionally, since Avamar stores its data as a file system, traditional backup software can back it up like any other server. Therefore, customers can back up Avamar data onto tape systems using Legato, TSM, Veritas and other popular backup applications. Support for tape is important, since customers may need to follow corporate policies that mandate that they back up or archive data to tape and vault it off-site.

Companies may also wish to consider using Avamar to back up all of their application servers to disk and using existing traditional backup software to back up the Avamar system to tape. This would improve backup and recovery performance, lower cost, reduce the number of media servers required, and still provide tape backups.

Integrated Search

Since Avamar stores its backup data on a file system, it can be integrated with search engines (e.g. Google, FAST) for fast access of data. ESG Lab observed FAST index and search operations running against a 3+1 Avamar configuration. Searches against an indexed Avamar repository from the web-based search interface, as shown in Figure Seventeen, completed nearly instantaneously. For example, a search for the string “Axion” in a 55 GB repository took 0.233 seconds to find 1,713 matches.

Figure Seventeen: A FAST Avamar Archive Search



Why This Matters

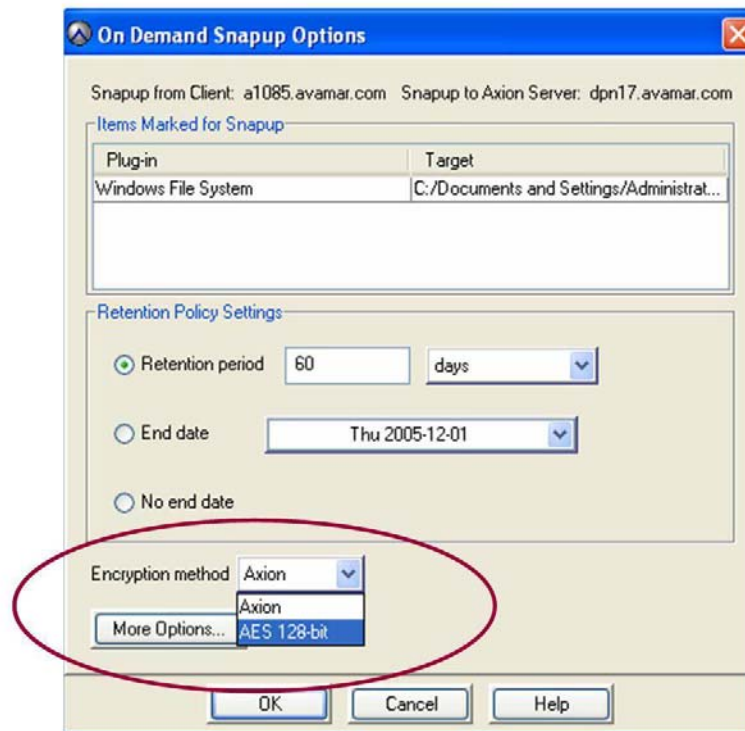
Instead of searching primary storage, where applications and users may be impacted, searches can be run against the Avamar backup repository using its file system view. The number of files being stored is increasing and having quick access to this data can improve business efficiency. Imagine trying to locate thousands, hundreds of thousands, millions, and even billions of files through traditional methods. Performing search on stored data is much more practical and efficient. Additionally, litigation protection and compliance mandates are driving the need for rapid and efficient search of digital records. Fast and efficient discoveries save time and reduce risk.

Reliability and Security

Avamar utilizes patented redundant array of independent nodes (RAIN) technology to deliver high availability and eliminate single points of failure. Avamar data nodes are arranged in a cluster and have been architected to survive a data node or hard drive failure. After a hardware failure, the Avamar system remains available for backups and restores. A data node was powered off and a drive was failed during ESG Lab testing. In each case, the Avamar system remained available for backups and restores.

Avamar backup, restore, and remote replication traffic travels over a local or wide area network. ESG Lab noted that Avamar always encrypts data using a light-weight Avamar-specific protocol, and that industry standard AES 128 bit encryption is supported, as shown in Figure Eighteen.

Figure Eighteen: Encryption Options



Why This Matters

ESG Lab believes that it is important that the solution used to protect your data is extremely reliable and secure. While this may seem self-evident, many customers who invest in fault tolerance for primary application servers tend to ignore fault tolerance when architecting a backup and recovery solution. Avamar delivers built-in fault tolerance using RAIN technology running on commodity servers, commodity RAID technology, and two levels of encryption, including industry standard AES for security.

ESG Lab Validation Highlights

- ☑ Avamar is the easiest enterprise-level backup software that ESG Lab has ever worked with. Configuring and executing backups and restores on Windows, Linux, and Solaris servers was extremely easy and intuitive.
- ☑ ESG Lab verified that the combination of global data de-duplication and single instance storage at the source, compression, and incremental forever algorithms can reduce the amount of backup data transferred across the network and stored to disk by up to 300x daily, compared to a traditional full backup to disk.
- ☑ The power of Avamar's data reduction technology was verified by auditing backups of systems in Avamar's lab and in conversation with an end user using Avamar software in production. In both cases, data reduction of approximately 300x was confirmed.
- ☑ An upgrade of an Avamar cluster was performed (from 3 data nodes to 4). A single web-based interface into the upgraded cluster was used to restore files created before the upgrade had occurred.
- ☑ Remote replication was configured from scratch and a restore from the remote Avamar system onto a client in the simulated primary data center completed without error.
- ☑ The Avamar file system view was used to perform a backup from the Avamar repository to tape using Veritas NetBackup. A restore from tape was performed using Veritas and did not require Avamar software.
- ☑ A search of an Avamar backup repository indexed using a search index created by FAST Search and Transfer completed in a fraction of a second.
- ☑ The manager of managers interface, new in the last version of Avamar software, was used from a web browser to monitor and diagnose the health of three Avamar backup repositories.

Roadmap Recommendations

The following recommendations were made by ESG Lab:

- Initial installation and upgrade is a manual process performed by EMC/Avamar service personnel. Automation is recommended to limit the amount of time on-site and reduce the likelihood of errors.
- Search capabilities could be enhanced with industry-specific plug-ins to help end-users develop policies and procedures for performing litigation and compliance-driven discoveries.

ESG Lab's View

ESG Lab, without hesitation, feels that EMC Avamar was the easiest enterprise-level backup software that we have ever worked with.

Unlike some GUIs (in particular, those that come with legacy backup products), the Avamar user interface is intuitive. ESG Lab was particularly impressed with its “manager of managers” interface, which allows end-users to manage multiple Avamar clusters from a central location. Whether performing an initial backup, executing a restore, replicating data from one Avamar cluster to another for disaster recovery purposes, or adding nodes to a cluster to increase capacity and performance, the process required just a few mouse-clicks. The only perceived negative is that the initial Avamar installation and upgrade process is manual, requiring EMC service personnel on-site.

At the core of the Avamar value proposition is its data reduction technology that provides incremental forever backups, global de-duplication at the source, single instance storage, and compression. Avamar is the only backup solution that solves the entire backup performance problem via its unique and powerful patented data reduction technology. This significantly improves backup performance, lowering the cost of disk storage for backup data, enables remote backups using existing WAN bandwidth, and provides a practical disaster recovery solution. ESG spoke to customers who went from a 16 hour backup window to being able to perform a backup in just 15 minutes. Another customer implemented a backup procedure for its remote offices, which, in the customer's words, was not possible using their traditional backup software.

Avamar added tape support which importantly allows end-users to deploy an Avamar solution in a traditional tape-based backup environment, leveraging existing backup applications and processes. This is necessary for customers mandated to use tape as part of their existing backup policies for off-site vaulting. Integration with tape is a necessary evil that Avamar has finally embraced, a fact that should remove a major objection that customers had with implementing its solution.

Another compelling aspect of the Avamar solution is that users can search and access data using third-party search engines (e.g. FAST, Google). ESG is a big proponent of reusing data to perform different functions versus creating multiple copies. Avamar can now act as a backup and recovery solution and as a digital archive. ESG believes that the use of existing stores of data for multiple purposes is a trend that will continue. And Avamar is one of the most efficient ways to store data, making it an even more attractive central repository for backup and archive. At this time, Avamar is the only solution that provides this dual-purpose functionality and customers should evaluate the potential value it brings to the table.

Avamar customers have consistently told ESG that they are able to reduce backup and restore times to just a fraction of what they were compared to traditional methods. ESG Lab has confirmed that the innovative EMC Avamar software delivers performance gains while drastically reducing the cost of disk capacity compared to traditional disk-based backup methods.

ESG believes that EMC's acquisition of Avamar was an excellent strategic move. With the Avamar solution in its portfolio, EMC can now offer a remote and branch office backup consolidation solution with a powerful and clear value proposition. ESG also believes that Avamar technology could be leveraged in EMC's existing portfolio, including NetWorker and RecoverPoint.

ESG feels that it bears repeating - the backup software market is ripe for reinvention and Avamar is a compelling solution, which leads the way.

Appendix

The equipment used during the ESG Lab validation is listed below along with the types of application data used for backup and restore testing.

Clients and Data Types Tested

Server	Amount of Data	Server Type	O/S	Application
Linux file system	8.5 GB	Dual 3.0 GHz Xeon, 4 GB RAM	RHEL ES3	ext3
Windows file system	8 GB	Xeon 3.0 GHz, 4 GB RAM	W2K3	NTFS
Windows/Oracle	30 GB	Dual 2.4 GHz Xeon, 2 GB RAM	W2K3	Oracle 10 R1
Windows/Exchange	12 GB	2.8 GHz Pentium 4, 1 GB RAM	W2K3	Exchange 2003
Solaris file system	10 GB	Dual UltraSparc-IIIi, 2 GB RAM	Solaris 9	ZFS

Configuration Details

Primary Data Center	
Avamar software revisions	3.0 and 3.5
Utility Node	1U Avamar server, 4x250GB SATA
Data Node 1	2U HP DL380, 6x300 GB SCSI
Data Node 2	2U Dell PE 2850 , 6x300GB SCSI
Data Node 3	2U SuperMico Optima, 6x300 GB SCSI
Data Node 4	2U IBM x346,, 6x300 GB SCSI
GigE switches	Dell PowerConnect 48 port, Cisco Catalyst 4506 6-slot.
Remote Data Center	
Single Node Avamar system	2U Dell PE 2850 , 6x300GB SCSI
GigE switch	ProCurve 2848 48-port