



zmanda
Open Source Backup

Amanda Enterprise

An Introduction

Abstract

This paper provides a technical overview of Amanda Enterprise. It describes design and operations of Amanda Enterprise and how it is unique in its ease of use, flexibility and scalability.

Table of Contents

<u>Overview</u>	3
<u>Introduction to Amanda Enterprise</u>	4
Amanda Enterprise Concepts	6
Backup Server	7
Zmanda Clients.....	8
Zmanda Windows Client.....	8
Zmanda Agents.....	9
Zmanda Management Console (ZMC).....	11
Backup Sets	12
<u>Close Look at Amanda Enterprise Features</u>	12
Ease of Installation and Configuration	12
Security.....	13
Holding Disk (Amanda Staging Area)	14
Media Management	14
Data Formats and Media Layout	15
Backup to Disk	15
Backup to Amazon S3	16
Intelligent Scheduling.....	17
Amanda Enterprise Monitoring and Reporting	19
Recovery Using the Zmanda Management Console	20
<u>Conclusion</u>	21

Overview

As organizations of all sizes become increasingly dependent upon their data, centralized protection and recovery of distributed data across all platforms, databases and applications is essential. Yet, faced with small IT staffs and limited budgets, many businesses cannot afford to purchase and manage expensive and complex backup solutions. Amanda Enterprise software provides enterprise-level backup and restore capabilities, coupled with ease of implementation and management, at less than a fifth of the cost of other well-known backup products.

With Amanda Enterprise, systems administrators can set up a master backup server to back up multiple Linux, Solaris, Windows, and Mac OS X clients to tape, disk, Internet cloud storage, or optical storage devices. Amanda Enterprise also protects databases such as Oracle and MS SQL Server, and MS Exchange. As an open source product, Amanda Enterprise uses open data formats—effectively freeing you from forced lock-in by a backup vendor.

Amanda Enterprise provides you with the following benefits:

- **Enterprise-class protection with simplified administration.** Offers enterprise-level functionality without unnecessary complexity. With an intuitive user interface, Amanda Enterprise is easy to install, configure, and manage.
- **Scalable & heterogeneous.** A single backup server can scale to backup and recover hundreds of desktops, workstations, and servers running Linux, Solaris, Windows, and Mac OS X operating systems.
- **Centralized and secure.** A centralized web-based management console enables protection and recovery of data distributed throughout your enterprise. Industry standard encryption and compression tools are used to secure and store data.
- **Intelligent scheduler.** A unique approach to scheduling automatically optimizes network and storage usage to achieve consistent backup windows and enhance administrator productivity.
- **Open formats.** Only industry standard data formats and tools are used. Unlike other backup products which backup data in proprietary formats—locking you into the vendor.
- **Low cost of ownership.** Subscription to Amanda Enterprise is up to 80 percent lower than licensing cost of proprietary backup products. Simplified and centralized operations improve administrator productivity, further reducing the cost of ownership.
- **Responsive technical support.** Users may choose from three levels of support, complete with knowledge base access, email and phone support, case management, software enhancements, and security updates.

Introduction to Amanda Enterprise

Amanda Enterprise software delivers high performance backup and recovery for Linux, Windows, Solaris and Mac OS X environments, popular databases and applications. All backup and recovery operations are managed from an easy to use web-based console. Amanda Enterprise is based on the most popular open source backup software Amanda.

Amanda started at the University of Maryland in the early 90s with the goal to protect a large number of computers with a single backup server. Today Amanda protects over half a million servers, workstations and desktops worldwide, running every major OS platform, both open and proprietary.

Here are the key benefits of Amanda:

Amanda Features	Benefits to System Administrators
Scalable client-server architecture	Use a single server to back up multiple networked clients to a variety of media types. Centralized backup simplifies administration.
Intelligent scheduling	A unique scheduler optimizes backup level for different clients in a way that equalizes the total backup time from one backup run to the next. Network and hardware resources are used as efficiently as possible.
Designed for backup to tape, disk and optical devices	You can choose the most appropriate medium as a target for your backups.
Unique, multi-tiered backups to tape and disk simultaneously	Data can be backed up redundantly to different media. For example, the same backup can be stored on disk for easy and fast recovery, and to tape for meeting the requirements of disaster recovery and long-term retention policies.
No proprietary device drivers	If the operating system supports it, the device will work well with Amanda. There is no chance that upgrading Amanda will break support for a device you depend on.
No proprietary archive or media formats	Amanda uses native UNIX/Linux dump and/or open archiving utilities such as gnutar and zip . You can always recover Amanda backups with these standard utilities, regardless of whether Amanda is present.
Security	Encryption on a client ensures security of data in transit and encryption on a backup server ensures security of data at rest, for example, on a tape.

Amanda Enterprise is the certified version of Amanda with enhanced usability, security, and expanded platform, database and applications coverage. It is the only open source backup software with enterprise support: Amanda Enterprise is available through subscription from Zmanda, Inc.

Depending on your configuration, you might need multiple subscriptions for different components. For details, please visit <http://www.zmanda.com/shop> or contact zsales@zmanda.com.

Each Amanda Enterprise subscription provides the following:

- Certified software to protect your specific platform, e.g. Windows server or a database such as Oracle on Windows
- Support services with a defined Service Level Agreement (SLA)

- Access to maintenance updates, security patches and new releases of Amanda Enterprise software including all new functionality
- Access to the Amanda Enterprise knowledgebase

Here is the brief overview of functionality of Amanda Enterprise that is beyond what is present in the free Community version of Amanda.

Additional Functionality	Benefits to Amanda Enterprise customers
Zmanda Management Console (ZMC)	<p>ZMC is a web-based Graphical User Interface (GUI) for all day-to-day activities of a backup administrator:</p> <ul style="list-style-type: none"> • Configuring backup policies, schedules and adding new platform clients and application agents • Verifying configuration and client-server communications • Starting immediate backup • Monitoring all backup activities • Reporting about backups and media • Creating, editing and deleting backup users and backup sets • Restoring files, databases and application data
Rapid installer	Provides wizard driven installation and allows you to install Amanda Enterprise in a few minutes.
Backup to Amazon Simple Storage Service (S3)	Amanda Enterprise can take advantage of Amazon S3 to backup, archive and retrieve any amount of data, at any time, from anywhere on the Internet. As an Amanda Enterprise customer, you get access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites.
Backup of open files on Windows	Support for Volume Shadow Copy Service (VSS) allows backup of open files on Windows with minimal impact on applications.
Backup of Windows System State	Simplifies disaster recovery of Windows hosts and enables roll back of any system changes by providing ability to restore a collection of several key operating system elements and their files.
Hot Backup of Active Directory	Ensures recovery of Windows Active Directory in case of hardware failure, software glitch, user error, or any other event that compromises the Active Directory.
Hot backup of Oracle	Backup and recovery of Oracle on Windows, Linux and Solaris from a centralized management console with minimal impact on users.
Hot Backup of Microsoft Exchange	Backup and recovery of MS Exchange from a centralized management console with minimal impact on users.

Additional Functionality	Benefits to Amanda Enterprise customers
Hot Backup of Microsoft SQL Server	Backup and recovery of MS SQL Server from a centralized management console with minimal impact on users.
Enhanced data security	ZMC provides role-based access for backup and recovery operations. Support for backup of Access Control Lists (ACLs) and extended attributes ensures that recovered files have the same permissions as the originals.
Comprehensive reporting	Provides pre-defined and custom reports about success and failure of backups, amount of data protected, performance, compression, media utilization and other backup parameters.

Amanda Enterprise Concepts

Amanda Enterprise allows you to set up a single backup server to protect multiple Linux, Windows, Solaris, and Mac OS X hosts to a large selection of media.

Typically, Amanda Enterprise protects networks with 5 to 100 backup clients and several databases. However, there are multiple Amanda installations with hundreds of clients protected by a single backup server. There are also installations where Amanda protects just a single client.

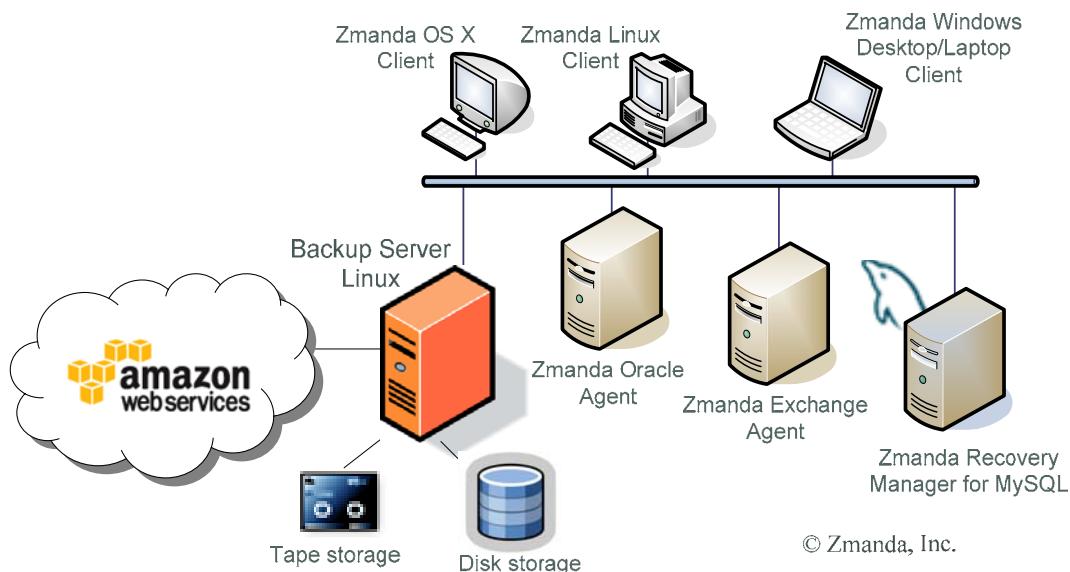


Figure 1. An example of a typical Amanda deployment.

The complete Amanda Enterprise deployment consists of the following:

- Backup server and ZMC centralize all backup operations and control backup media. You need at least one backup server for your environment.
- Zmanda clients that are installed on Windows, Solaris, Linux or Mac OS-X servers, workstations, desktops and laptops.
- Zmanda agents that are installed in addition to Zmanda clients for hot backup of Oracle, MS Exchange and MS SQL Server.

- Zmanda Internet Backup is an option for backup over the Internet to Amazon S3 Service.
- Zmanda Recovery Manager (ZRM) for MySQL is an option for comprehensive backup and recovery of MySQL servers. For more information about ZRM please download “Zmanda Recovery Manager: Guide to MySQL Backup & Recovery” white paper.

Note: Amanda Enterprise also supports backup over NFS and CIFS. If you can remotely mount your file system to Amanda backup server, you can protect that system, without installing Zmanda client on that host. In fact, this is the recommended method for backup and recovery of NetApp filers and other appliances that require agentless backup.

Backup Server

The backup server manages all backup and recovery operations and keeps track of backup data through the backup catalog that contains an index of all files, database data and log files, backup levels, backup times and locations of backup files on media.

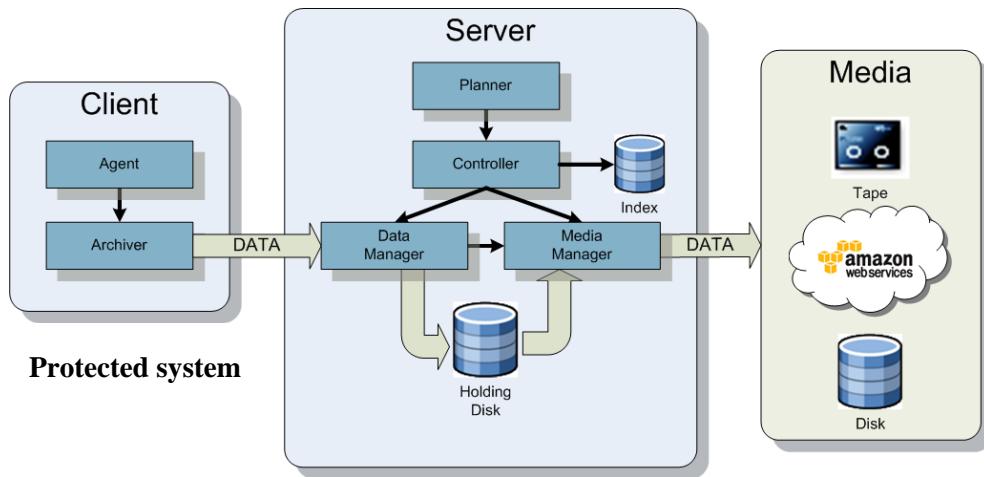


Figure 2. Amanda Enterprise architecture.

The client server architecture provides the following benefits for backup operations:

- It ensures scalability of Amanda Enterprise environments. Many configurations are possible. You can use Amanda to backup a single client. However, you can also deploy it for thousands of clients with hundreds of terabytes of data that you can backup to large tape libraries with multiple drives.
- It allows all configurations to be done and stored in one place – on the backup server. Once the initial configuration is completed, adding additional clients requires no changes to tested backup procedures.
- It allows processing load to be distributed. Some CPU-intensive operations such as compression or encryption can be done on the client before sending backup images to the backup server.

The Amanda backup server also manages storage devices by using non-proprietary device drivers. Any storage device, for example, a tape drive, supported by the operating system of the backup server, will work. System Administrators do not have to worry about breaking support for a device when upgrading Amanda Enterprise software to the latest release, or when upgrading an operating system.

Currently we provide support for the following types of media and storage devices:

- LTO, DLT, SDLT, AIT, SAIT, DAT and many other tape drives
- Tape libraries, jukeboxes and auto changers

- JBOD and RAID arrays, NAS devices
- Virtual Tape Libraries (VTLs) and deduplication appliances such as Data Domain Storage Systems
- Optical devices
- Amazon Simple Storage Service (S3)

Zmanda Clients

Zmanda clients are available for Windows, Linux, Mac OS X and Solaris. For the up-to-date list of supported operating systems, please review [Zmanda Compatibility Guide](http://www.zmanda.com/supported-platforms.html) (<http://www.zmanda.com/supported-platforms.html>). All clients interface seamlessly with the ZMC to provide one intuitive and consistent interface to backup and recover multiple Linux, Solaris, Windows, and Mac OS X systems.

Zmanda clients provide the following functionality common for all platforms:

- **Simple, flexible backups.** Supports full and multiple levels (0-9) of incremental backups.
- **Open source formats.** Uses only industry standard data formats such as ***tar*** for Linux, Solaris and Mac OS X) and ***zip*** for Windows data formats.
- **Cross platform restores.** You can restore data from a backup to a host with a different operating system. That simplifies the disaster recovery process.
- **Support for hard links.** Maintains the integrity of hard links during backup
- **Support for sparse files.** Supports the backup of sparse files to optimize the use of storage media for backups.
- **Timestamps unchanged.** File timestamps remain intact during backup
- **Exclude lists.** Administrators can choose to include or exclude specific files and directories during backup
- **Encryption.** Administrators can choose which data to encrypt, which standard encryption algorithm to use, and whether to encrypt data on the client or on the server.
- **Compression.** Uses only industry standard compression utilities, and allows administrator to choose the optimal compression algorithm for the data and the resources available, as well as whether to compress the data on the client or the server.
- **Wide platform support.** Extensive and up-to-date support for Windows, 32-bit and 64-bit Linux x86 distributions, Solaris on SPARC and AMD, and Mac OS X on POWER and Intel.

Note: Zmanda Mac OS X client supports resource forks used by many Mac applications and files.

Zmanda Windows Client

Zmanda Windows client provides the same functionality as other clients and more. The backup data is stored in a Windows ***zip*** format used by Compressed Folders in Windows XP, Vista, Windows Server 2003 and Windows Server 2008.

The Windows client is able to backup and restore all metadata information about the file: ACLs, extended attributes, reparse point information, hardlink information, object identifiers, alternate data streams (ADS) and property data. The client is also able to backup and restore encrypted files, sparse files, compressed files, junction points and hardlink files on NTFS.

The Windows client provides functionality that is not applicable to Unix and Linux but is important to the Windows platform:

- **Protects System State.** Simplifies disaster recovery of Windows hosts and enables roll back of any system changes by providing the ability to restore the key operating system elements and files.

- **Backup of Active Directory.** Ensures recovery of Windows Active Directory (AD) in case of hardware failure, software glitch, user error, or any other event that compromises the Active Directory.
- **VSS support.** Takes advantage of Microsoft Volume Shadow Copy Service (VSS) snapshot technology to backup open files and databases on Windows without impacting users or applications.

Zmanda Windows client is a Requestor for VSS Writers such as Active Directory service or MS SQL Server. VSS creates a storage volume with shadow copies (a.k.a. snapshots) while a source volume continues to operate with minimal impact on users. For details on how Zmanda Windows client creates shadow copies, please refer to **Figure 3** in the next section.

Depending on what computing role is assigned to a backup host (e.g. workstation, server, domain controller), the Zmanda Windows client will backup different System State components. The table below provides an overview of System State components required for a successful recovery:

Component	When is this component included in System State?
Registry	Always
COM+ Class Registration database	Always
Boot files, including the system files	Always
Certificate Services database	If backup client is a Certificate Services server
Active Directory directory service	Only if backup client is a domain controller
SYSVOL directory	Only if backup client is a domain controller
IIS Metadirectory	If it is installed
System files that are under Windows File Protection	Always

Specifically, the System State backup for Windows XP includes only the registry, COM+ Class Registration database, files under Windows File Protection, and boot files.

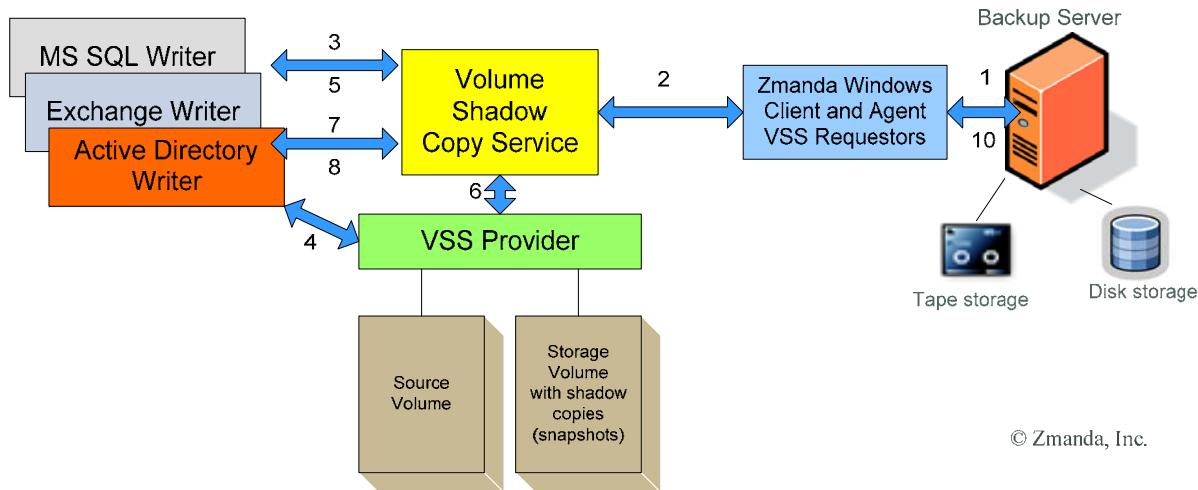
For Windows 2003, backup of the System State includes the registry, COM+ Class Registration database, files under Windows File Protection, and system boot files. Depending on the configuration of the server, other data may be included in the System State data. For example, if the Zmanda Windows client is installed on a server with a domain controller, the Active Directory and the SYSVOL directory are also contained in the System State data.

When you choose to back up the Windows System State, all of the System State files are protected. Because of dependencies among the System State components, you cannot select for backup only the individual components of the System State. However, you can restore the individual System State files to an alternate location. For example, you might choose to restore only the registry.

Zmanda Agents

Zmanda Enterprise Agents provide reliable online backup and recovery for Microsoft SQL Server and Oracle databases, as well as Microsoft Exchange. Database administrators benefit from fast implementation, enterprise-class functions, centralized management, standard formats, and responsive technical support. Users benefit by being able to continue to access their data during backup operations.

The following diagram illustrates how Zmanda Windows client and Database Agent work with VSS to create backup copy of a volume with Active Directory (AD) and MS Exchange Server (also applicable for MS SQL, Oracle and any other VSS-aware application).



© Zmanda, Inc.

Figure 3. Zmanda Windows client integration with VSS for backup of VSS-aware applications.

1. Backup server initiates scheduled backup of Zmanda Windows client and Exchange agent.
2. Zmanda Windows client is a Requestor and asks the VSS to prepare for shadow copy creation.
3. AD and Exchange Writers create an XML description of the required components for backup. VSS notifies AD and Exchange Writers to prepare data for making shadow copies.
4. Both Writers prepare the data as appropriate. For example, the Exchange Writer completes all open transactions, rolls transaction logs and flushes cache. When data is ready, AD and Exchange Writers notify VSS.
5. VSS tell the Writers to quiesce AD and Exchange and freeze the Zmanda Windows client I/O write requests for a few seconds to create the shadow copy of Source Volume. VSS flushes the file system buffer and then freezes the file system. This ensures that file system metadata is written and that data is written in a consistent order.
6. VSS tells the Provider to create the shadow copy. Usually it takes a few milliseconds, but according to Microsoft specifications for VSS it should take 10 seconds maximum.
7. VSS thaws the file system. After the shadow copy is done, VSS releases the Exchange and AD Writers and all queued I/Os are completed.
8. VSS verifies that all writes were successfully held. If not, the shadow copy is deleted (because it might be inconsistent) and Zmanda Windows client is notified so it can re-try the process.
9. If the copy is successful, VSS gives the information about location of Storage Volume with shadow copy backup to Zmanda Windows client.
10. Zmanda Windows client streams backup data to backup server in **zip** format.

The recovery process also takes advantage of VSS-awareness of Windows applications and NTFS in a similar way.

Zmanda Oracle Agent allows administrators to perform backup at any time, while users and applications continue to access Oracle databases. For Oracle backup, the Zmanda Agent uses Oracle's preferred mode to back up a given operating system:

- VSS for Oracle on Windows

- Oracle Recovery Manager (RMAN) for Oracle on Linux and Solaris.

For Oracle on Windows, the Zmanda Oracle Agent takes advantage of VSS in a way that is described for MS Exchange on **Figure 3** above.

For Oracle on Linux and Solaris, the Zmanda Oracle Agent, through integration with RMAN, works intimately with the Oracle server, providing block-level corruption detection during backup and restore. RMAN takes care of all underlying database procedures before and after backup or restore, freeing dependency on operating system and SQL*Plus scripts. It provides a common interface for backup tasks across different versions of Linux and Solaris, and offers features such as parallelization of backup/recovery data streams, backup files retention policy, and detailed history of all backups.

The table below provides an overview of Zmanda Database Agents' features:

Features and benefits	Zmanda MS Exchange Agent	Zmanda MS SQL Server Agent	Zmanda Oracle Windows Agent	Zmanda Oracle Solaris/Linux Agent
Hot backup with minimal impact on users	Yes	Yes	Yes	Yes
Seamless integration with Zmanda Management Console	Yes	Yes	Yes	Yes
Supported backup levels	Full and incremental	Full and incremental	Full and incremental	Full and incremental
Open format for backup data	zip	zip	zip	tar
Database vendor preferred way to backup	VSS	VSS	VSS	RMAN
Encryption	Yes	Yes	Yes	Yes
Compression	Yes	Yes	Yes	Yes
Supported releases	2007 2003	2005 2000	Oracle 11g	Oracle 10 and 11

Zmanda Management Console (ZMC)

In addition to a powerful Command Line Interface (CLI), Amanda Enterprise also provides the Zmanda Management Console (ZMC).

ZMC is a browser-based user interface for setting up and managing all backup and recovery activities. It is integrated with the Zmanda Network, which provides certified Amanda Enterprise binaries, white papers, demos, technical support and a knowledgebase to help you to deploy your backup solution. ZMC provides context-sensitive help and the most current information about Amanda Enterprise from the Zmanda Network. After registering with the Zmanda

Network you will be notified of new software releases via RSS feed. You can also use ZMC for communicating with Zmanda. For example, right from ZMC you can provide feedback about the product.



Figure 4. Zmanda Management Console top-level navigation.

ZMC provides you with an easy to use tool for managing all backup and recovery operations for multiple clients and various databases. Using ZMC, a backup administrator can accomplish all day-to-day backup activities:

- Configure backup policies, schedules and clients
- Verify configuration and client-server communications
- Start immediate backups
- Monitor Amanda Enterprise activities and access reporting about backups and media
- Create, edit and delete backup users and backup sets
- Restore files

Backup Sets

Backup set is a key policy concept of Amanda Enterprise that simplifies and optimizes the enterprise-wide backup system with multiple clients protected in a similar way. The backup set defines the following:

- List of hosts, applications, directories, excluded files, and other details of what to backup
- Type of media (disk, tape, VTL, Amazon S3) used to backup a given backup set
- Backup time and parameters for backup level scheduling.

Each backup client belongs to at least one backup set. Backup sets are easy to copy, add and modify. There are several templates for common backup policies to help you get started with Amanda Enterprise:

- Backup to disk
- Backup to a single tape drive
- Backup to a tape changer
- Backup to Amazon S3

All backup, reporting, monitoring and restore actions are performed on the backup sets. This enables the backup administrator to abstract the backup policies and apply the same policy for multiple clients.

Close Look at Amanda Enterprise Features

Ease of Installation and Configuration

The Rapid Installer enables easy installation of the Amanda Enterprise backup server, ZMC and all of its components that are built on certified Apache, MySQL, and Perl/PHP stacks for versions of Linux and Solaris supported for the

backup server. Through wizards, the rapid installer also takes care of important administrative tasks such as creating the backup user **amandabackup** that has proper permissions for backup and recovery.

For any backup software, the initial configuration of a server and the backup clients is key for error-free operations. Amanda Enterprise provides sophisticated verification tools to ensure that both the backup server and the backup client configurations are done properly. You should run the “Verify Clients” test every time you install a new client.

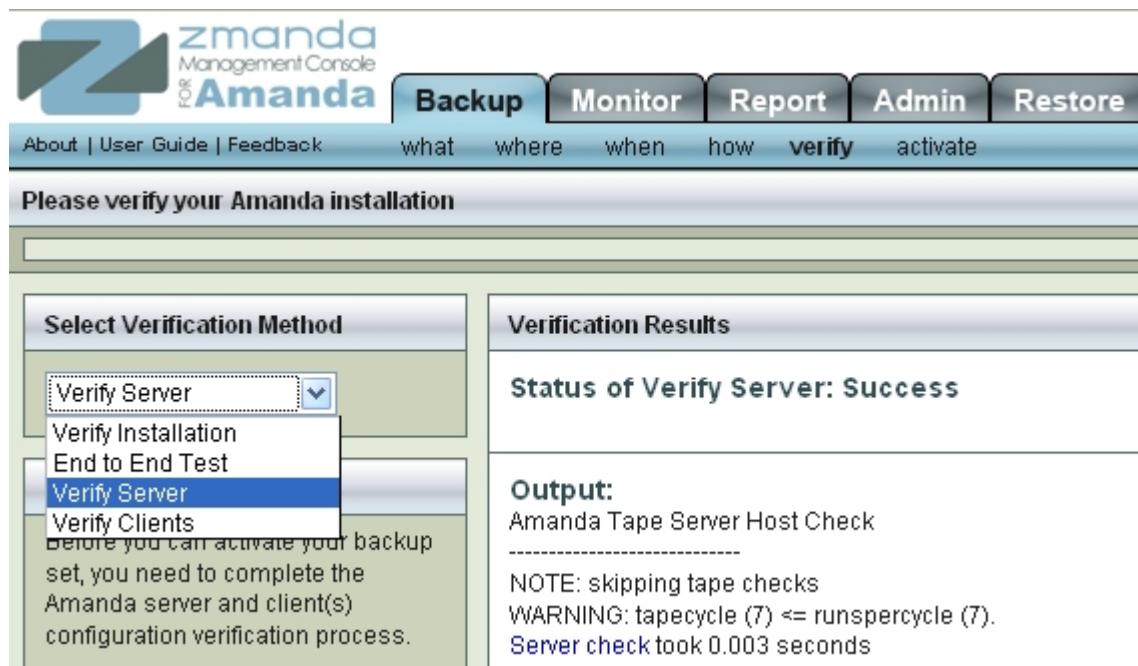


Figure 5. The verify page of ZMC to test the initial setup.

Performing "End to End" test will quickly backup and restore a test file on the server to check the operation of the Amanda Enterprise server. This test does a disk backup and recovery of a file and is independent of the backup set configuration.

Security

Security is a key concern in enterprise backup. Amanda Enterprise provides many layers of security to ensure that backup data, communications, and access to the backup process itself are secured.

The Zmanda Management Console (ZMC) provides role-based access controls for backup administrators. An organization can have both **Administrators**, who have universal access, and **Operators**, who have more limited access defined. This allows for configuration and management of more important data to be performed under tight control, and less sensitive data to be manageable by more people.

Separation of roles also allows sharing Amanda Enterprise for multiple backup users. For example, you can provide backup services to hundreds of backup customers. Each customer with Operator's rights might have multiple backup clients and applications, but under no circumstances will a customer have access to somebody else's data.

Amanda gives you the flexibility to encrypt data either on the backup client or on the backup server. Having that choice, you can encrypt data where it is most appropriate. For example, if you were concerned about the backup traffic being intercepted as it flows over TCP/IP, encrypting the data on the client would be most appropriate. In other cases, network security is provided by other means, for example via VPN. Since encryption can take significant CPU resources, you may choose to encrypt the backup data on the backup server so data is secure on backup media, e.g. on a tape or Amazon S3.

Amanda Enterprise comes with interfaces to standard open-source encryption tools such as *openssl* and *GPG*. Through ZMC you can easily turn on server-side encryption. Simply click a check box labeled “Encryption” when defining the directory to be backed up (shown below).



Figure 6. Enabling encryption with ZMC

Holding Disk (Amanda Staging Area)

One of the features that enable Amanda to scale up and provide high performance is the “holding disk”. As the name suggests, a holding disk is a cache to store backup data from Zmanda clients and application agents. Modern tape drives are very fast. An LTO-3 has throughput of 80 MB/s. Even a Gigabit connection can not push backup data fast enough to keep the tape drive streaming.

Using holding disks as a staging area for backups has several benefits:

- Holding disks can accept data streams from multiple clients in parallel to overcome the serial nature of a tape. Instead of writing one backup to tape after another, you can configure multiple backups running in parallel, making full use of your available network bandwidth and reducing total backup time.
- If network becomes the bottleneck for performance, you can reduce total backup time by adding another NIC to your backup server or dedicating a separate network for backups.
- Using holding disks provides additional safety in case there is a bad tape, the wrong tape is inserted in the drive, or there are no available tapes. Backups will still be complete and safe on the holding disks.

Amanda supports multiple holding disks so that backup images from different clients can be sent to different holding disks. This means better scalability by providing an easy mechanism for load balancing: to increase throughput, simply add more controllers and holding disks.

On the other hand, the holding disks are not required. Amanda can write backups directly to media. In fact, it is a good practice not to use a holding disk when disk is the backup medium. However, for backup to tape and to Amazon S3 you will benefit from using the holding disk.

Media Management

Amanda Enterprise manages all media through the use of a “label”. Tape labeling allows the backup server to keep track of all tapes that were labeled.

A new tape is used for each backup run. Based on the backup retention policy, Amanda keeps track of the expiration date for each labeled tape. Once an old backup image expires, the backup server will re-use that tape for new backups. However, Amanda Enterprise can be configured not to re-use specific tapes. This allows certain tapes to be used as archives.

For backups of large amounts of data Amanda supports tape changers and tape libraries for writing to multiple tapes in a single backup run. Backups for multiple clients can fit on a single tape and large clients can span multiple tapes.

Data Formats and Media Layout

Amanda Enterprise is the only backup and recovery product that does not use proprietary tools and data layouts when backing up your data. Both **tar** and **zip** formats are industry standard and common for Linux/Unix/Mac OS X and Windows. Their specifications and data layouts have been stable for many years, with the promise that this will continue for years to come.

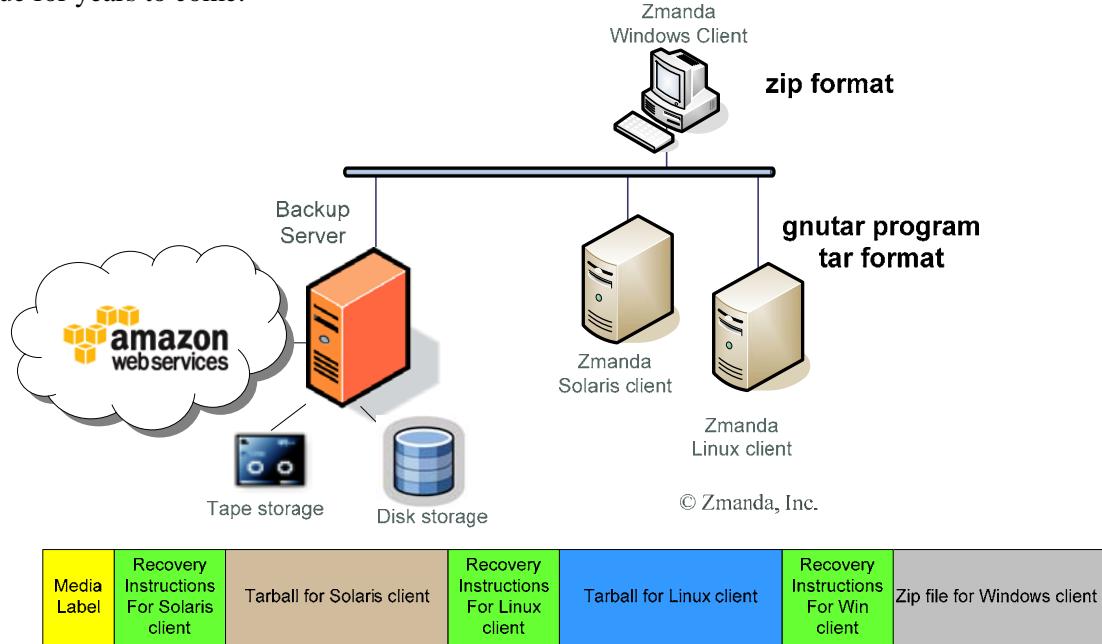


Figure 7. Media layout for backup images.

Using holding disk for backup of multiple clients at the same time, Amanda allows for backup images from several sources to reside sequentially on the same piece of media. Amanda separates each backup on tape using a tape file mark.

When Amanda writes the output of a backup to tape (or other medium) it does not alter the output in any way. It does precede each **tar** or **zip** file with a header. This header contains the exact command sequence needed to restore the data with standard operating system utilities. The instructions in this header are optional – you only need to use these instructions if you are trying to recover data from the dump without the use of ZMC or CLI recovery commands - for example, in a disaster recovery situation with even Amanda not running.

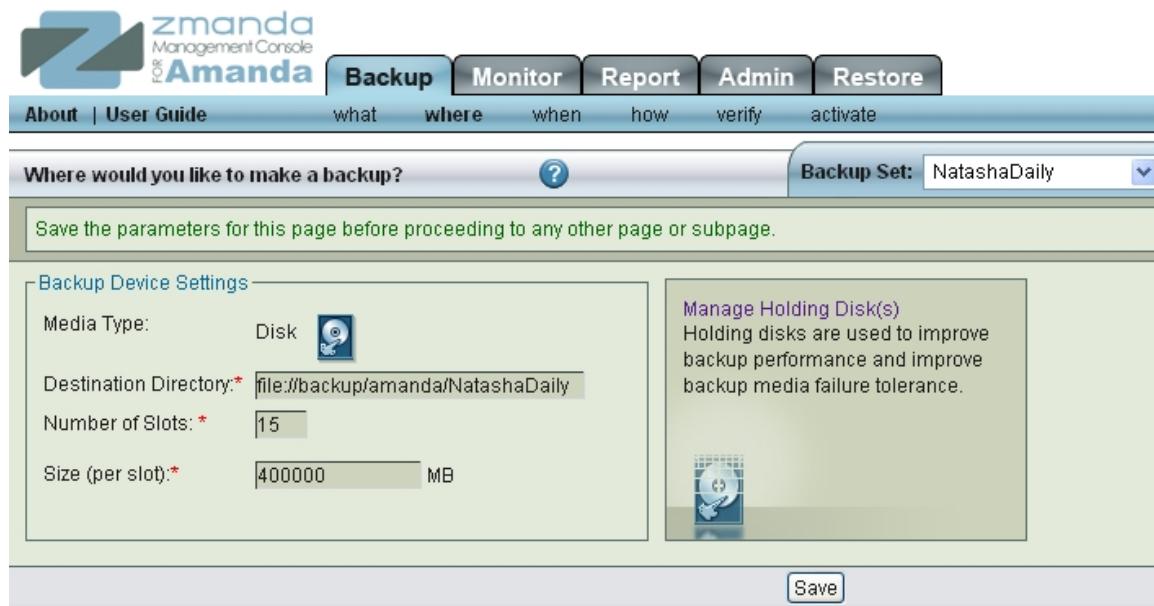
Amanda can span multiple media volumes to accommodate very large backup jobs, and it uses the very same data layout for any supported media. Using open formats and standard media layouts means that unlike proprietary solutions Amanda does not hold your data hostage. If Amanda is unavailable for any reason you can still recover your data.

Backup to Disk

More and more system administrators are using backup to disk. Backup to disk has several advantages over backup to tape, one being the speed of recovery. Most commercial backup packages include some level of support for backup-to-disk, and Amanda does as well.

Some commercial products can only emulate a single tape drive with a single piece of media. Amanda on the other hand emulates an entire tape library. This allows you as the system administrator to configure both backup to disk and backup to tape in the same way – there is no need to learn two different ways.

In ZMC you start your virtual tape configuration by basing your configuration on the “hard-disk template”. You then define a few parameters on the **Backup -> where** page:



The screenshot shows the ZMC interface with the 'Backup' tab selected. The 'what', 'where', 'when', 'how', 'verify', and 'activate' tabs are also present. A 'Backup Set' dropdown is set to 'NatashaDaily'. The 'Backup Device Settings' section contains fields for Media Type (Disk), Destination Directory (file:///backup/amanda/NatashaDaily), Number of Slots (15), and Size (400000 MB). A 'Manage Holding Disk(s)' section is visible on the right, with a note about holding disks improving backup performance and media failure tolerance. A 'Save' button is at the bottom.

Figure 8. Configuring vtapes in Backup Where page

If you selected backup to disk template while creating a new backup set, Disk will be the only choice for the *Media Type*. The *Size* is the maximum size of each virtual tape. This space is not pre-allocated, so Amanda will only consume the amount of space required for each virtual tape.

The *Destination Directory* specifies which directory will hold the virtual tape library. ZMC provides a default location (under the *amandabackup* user's home directory), but you can change it to a more suitable location. A best practice is to use a separate disk for your virtual tape system.

Backup to Amazon S3

In addition to traditional backup to disk, tape and optical devices, Amanda Enterprise can use Amazon Simple Secure Storage (S3) Service to backup, archive and retrieve any amount of data, at any time, from anywhere on the Internet. You are getting access to the same scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites. Storing a copy of data on S3 gives you many benefits:

- Streamlined offsite backup storage & recovery process
- Ability to retrieve data anytime, anywhere.
- Restoration of data from S3 is much faster than from offsite tape storage
- Cost effective “pay as you go” (\$0.20 per GB per month)
- Online storage service backed by Amazon's 99.9% SLA

When you configure a new backup set you simply select Amazon S3 as a backup destination as shown below. You will benefit from configuring a holding disk so you can continue backups even if Amazon S3 is temporarily not available. After initial configuration, Amanda Enterprise securely sends an encrypted copy of your data for storage on Amazon S3.



Figure 9. Configuring backup to Amazon S3

In comparison to traditional storage of tapes off-site for disaster recovery, the Amazon S3 provides streamlined recovery process with minimal downtime for your business.

Recovery from off-site tape	Recovery from Amazon S3
Identify location of the required tapes	
Request your off-site tape storage vendor to deliver the tapes on truck (that step takes the most time)	Find files or directories using intuitive ZMC and click Restore
Mount and read the tape until required data is found and recovered	
Unload the tape and repeat with another tape if required	
Total downtime for best case scenario is hours, but usually takes several days	Total downtime is a few minutes

Intelligent Scheduling

For all other backup products you must specify which backup levels to run and on which specific days to perform these backup levels. If you have a large number of backup clients to configure, this can be complicated and time consuming. Amanda intelligently determines the backup levels for you, thus relieving you of this burden.

Also, using any other backup product you tend to do all of your full backups on one day, for example, on Sunday and perform incremental backups the rest of the days. All backup clients that use this schedule will perform their full backup on Sunday. This dramatically increases the amount of time for Sunday's backup relative to the rest of the week. You no longer have a consistent amount of time for your backups each day. This is of course because very different amounts of data are backed up each day. In addition, much more backup media will be consumed on Sunday meaning you have to verify if enough media is available to finish the full backup. Also, if one of the clients had

significant change of data on Tuesday, running just incremental backup for that day (as opposed to a full backup) might be very inefficient for recovery.

Amanda, on the other hand, has the following design goals related to backup levels:

- Backup a consistent amount of data each backup run making consistent backup window
- Thus consume a consistent amount of backup media each run.
- Recalculate and re-balance the schedule when you add a new client to the network or an existing client is not available or a client has significant change of data.

To accomplish these goals Amanda has an estimation or planning phase at the beginning of a backup run. A backup run is a backup of multiple file systems, directories and databases, often from multiple backup clients. The goal of the planning phase is to determine the backup levels for the individual backups in the backup run. Each backup run results in a different mixture of full and incremental backups.

One determining factor for the individual backup levels to be run is how much data that level will produce. Amanda attempts to keep the total amount of data for each backup run the same, which results in a consistent backup window.

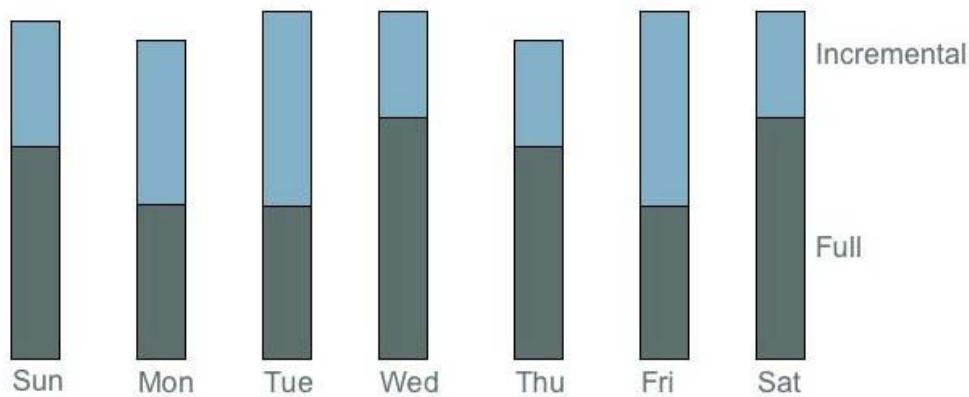


Figure 10. Amanda equalizing backup jobs every day

Amanda's backup level algorithm also manages spikes in data changes. For example, suppose that on Tuesday an unusually large amount of data is changed. Amanda's planning algorithm can balance the large amount of incremental data by promoting some backup jobs from incremental to full backups.

Backup Fault Tolerance

Since backup touches most of your IT infrastructure, any of the moving parts can have a potential glitch. The architecture of Amanda Enterprise allows it to keep functioning even in the face of multiple faults.

Amanda Enterprise gracefully skips systems (e.g. laptops) which can not be backed up in a particular run. When this skipped system is available in a future backup run, Amanda's intelligent scheduler promotes its backup to an appropriate backup level.

In case of backup media errors such as tape drive errors or temporary unavailability of remote cloud storage, Amanda Enterprise caches the backup data on the holding disk. You can migrate this data to backup media after the fault has been resolved.

Amanda Enterprise Monitoring and Reporting

The complexity of backup operations for various operating systems and databases makes it difficult to estimate how well you meet the business objectives of data protection. Comprehensive backup reporting provided by Amanda Enterprise helps you to ensure your data is protected well. It also allows you to fine-tune the backup operations for optimal performance because your initial configuration might be based on guesses or even incorrect assumptions about your clients and data. Reports will provide you with information to adjust your assumptions.

You can monitor each backup run and see where you are in a backup process as shown below. However, you don't have to be in front of your computer to monitor details of backup runs. All relevant information is saved, and you can review the progress through different backup phases later.

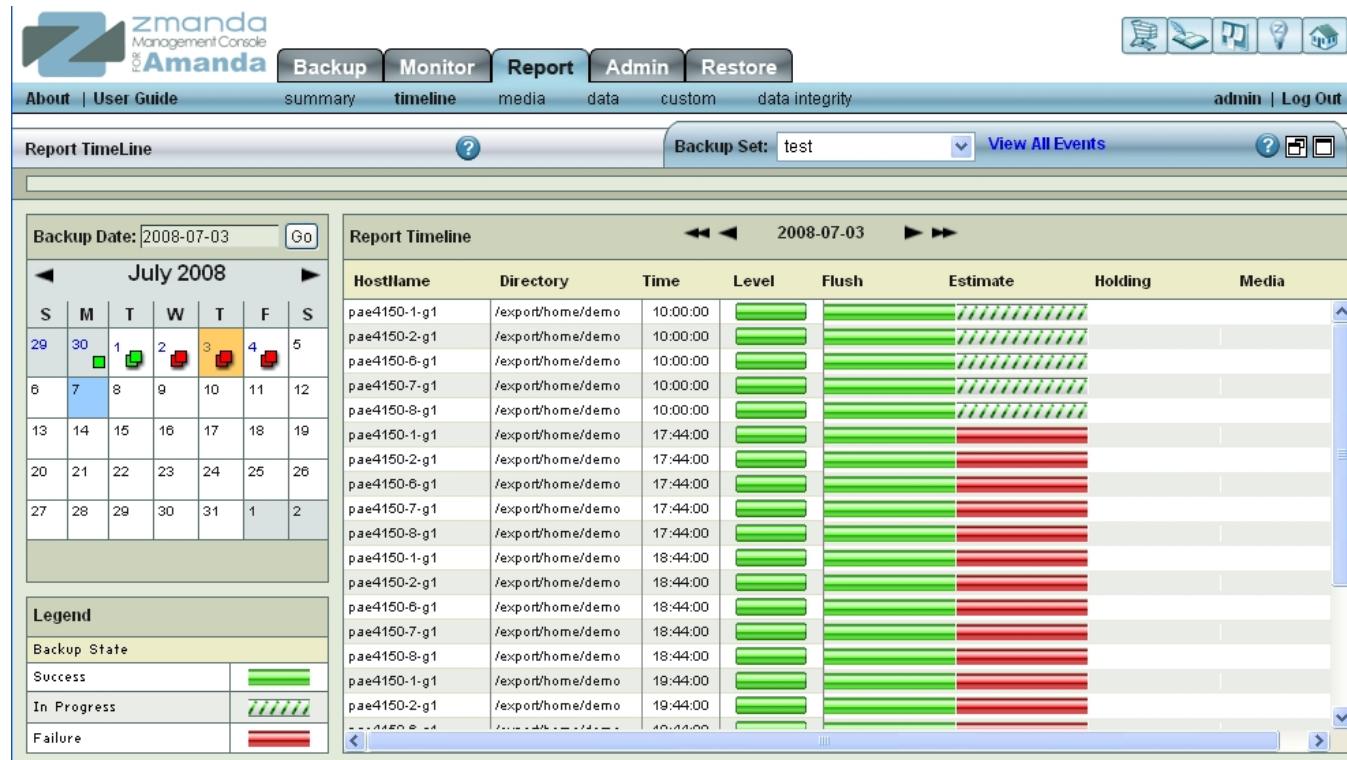


Figure 11. Timeline report

Timeline reports provide you with a historical record of monitoring information. It allows you to see all phases of the backup run and troubleshoot any errors because you can easily see where exactly in a backup process an failure happened. For example, on **Figure 11** you can see that none of the backups on July 3 were successful because there was a failure during estimate phase of backup.

Calendar view provides an easy to understand visual summary of your backups for a month with color-coded indicators of success, failure or warning. Multiple squares indicate several backup runs that day. In case of a problem, you can quickly drill down to a report with details about any backup run.

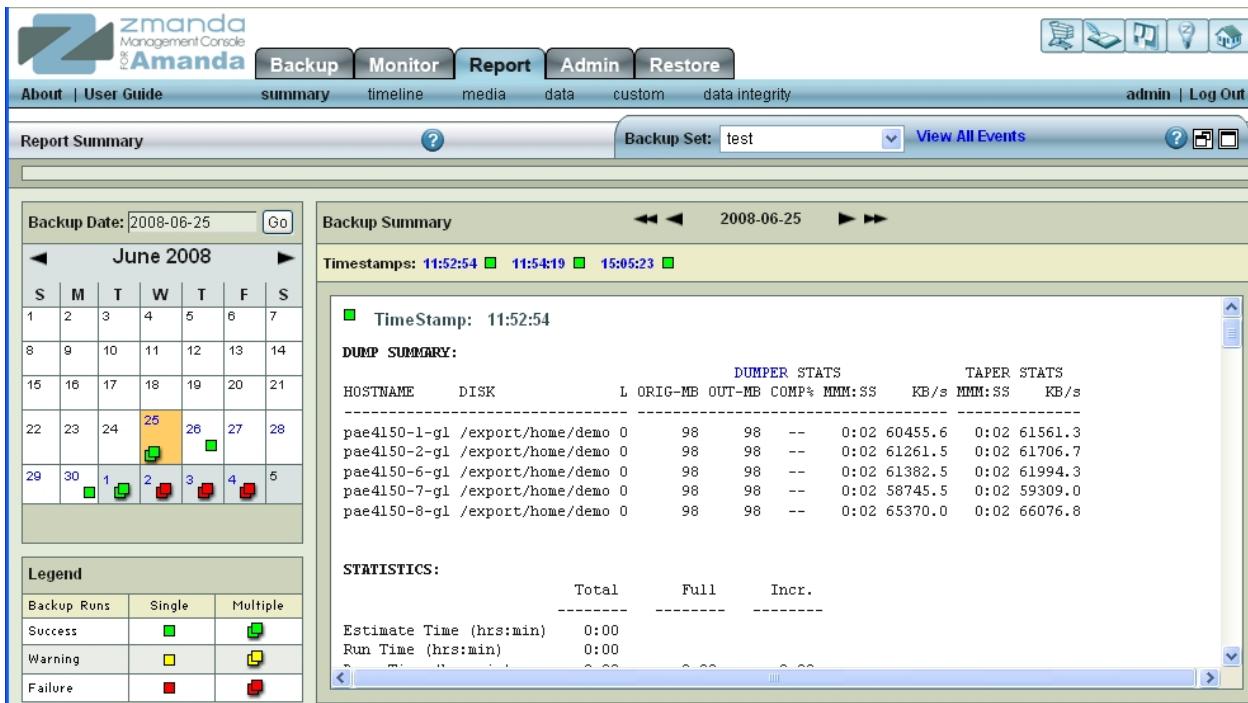


Figure 12. Calendar view of recent backup reports

Each backup run (successful or not) results in a report with details about all clients scheduled for that run:

- List of all clients and applications in that backup run
- Backup levels for each backup set
- How much data for each client was backed up
- Compression ratio achieved
- Backup anomalies such as errors and other unexpected events (missing client)

In case of a warning or an error, the report provides a link to a relevant article in Amanda Enterprise documentation. You can configure Amanda Enterprise to receive backup reports by e-mail.

Media reports provide visual overview of media utilization and performance. This helps you to understand your data growth, identify backup performance bottlenecks and plan ahead when to add more media.

You can customize the reports to extract and present data in a variety of ways. Specifically, you can monitor success and failures, backup performance, compression, amount of data written, and many other parameters. The reports can be saved, e.g. for compliance purposes, or exported into your dashboard monitoring application and analyzed for time trends.

Recovery Using the Zmanda Management Console

As with backups, recovery using the ZMC is based on a simple workflow:

- Determine what data should be restored – **restore what**
- Determine where the data should be restored – **restore where**
- Run the restore process – **restore**

Using the **restore what** screen you select what to restore. This selection can be as simple as a single file, database or directory or more complex such as all directories/files under a single directory. Alternatively, you can start the restore

process by clicking on a timestamp link included in any backup report. Cross-platform recovery allows you to recover for example, backups from a Solaris client on a Linux host.

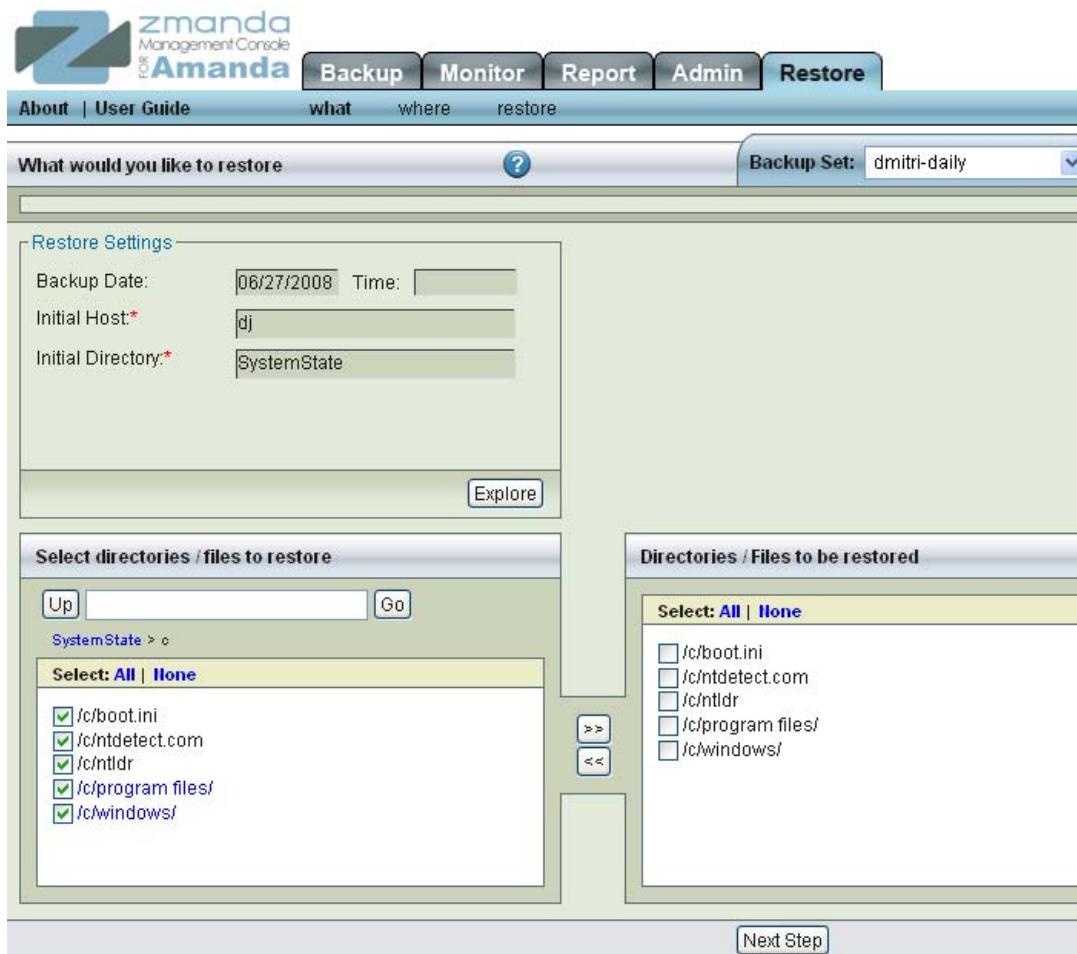


Figure 13. Example of restoring the System State

In the example above the System State of the Windows workstation will be restored as a whole without selecting individual components. The **restore where** page is to select where the data is to be restored. One of the key features of **restore where** is to prevent potential conflicts in restoring data. If the existing files match the file to be restored, **restore where** can be configured for the following options:

- Keep existing files
- Overwrite existing files
- Rename existing files
- Rename restored files

Finally, the **restore** page is for initiating the restore process. This is where the settings can be checked, and where prompting for media is done (if required.) Once ready, you start the restore process and monitor the progress until successful completion.

Conclusion

Based on the proven Amanda open source backup software, Amanda Enterprise provides a powerful yet easy to use software to protect Linux, Solaris, Mac OS X and Windows systems as well as MS Exchange, MS SQL Server and Oracle databases from a web based management console. For more information, please contact zsales@zmanda.com