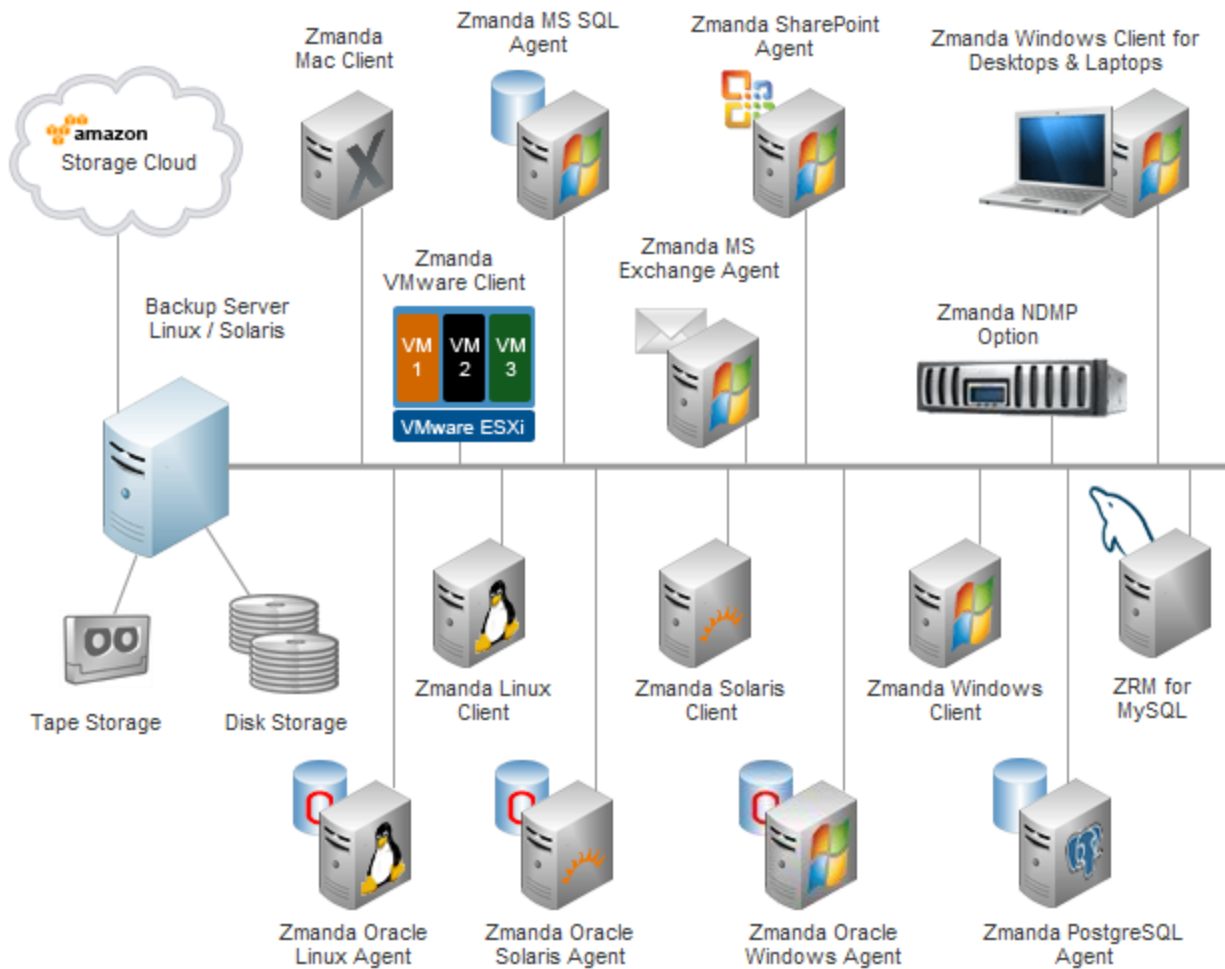




ZMC User's Manual for Amanda Enterprise Edition 3.4

Amanda Community and Amanda Enterprise Editions: Open Source Backup for Everyone

The Amanda open source backup server protects more data than any other open source backup solution. It allows backup administration through a single backup server that can back up one or more client systems. Amanda Enterprise combines the benefits of open source backup with enterprise-level support, a graphical user interface for the server, and enhanced backup client software such as the ability to perform hot backups on Windows server applications (separate licensing required).



The advantages of Amanda/Amanda Enterprise include:

- Backs up clients of all the relevant server and desktop OSs: Linux/Unix, Solaris, Windows, and Mac OSX.
- Uses non-proprietary native archive utilities (such as dump, tar, and zip).
- Backs up to any device accessible to the Amanda backup server (including Amazon S3). Use of virtual tapes allows backup to disk media.
- Parallel operation and intelligent scheduling ensure scalability both in number of systems backed up and in size of backup sets.
- Separately-licensed application modules allow intelligent backup of selected application servers such as Oracle and MS Exchange (Enterprise only).

Amanda Community Edition Documentation

- Amanda documentation is maintained on the Amanda Wiki at <http://wiki.zmanda.com/>.

Supported Platforms

Please see the latest platform support information on the Zmanda Network at <http://www.zmanda.com/supported-platforms.html>.

Pre-Installation Requirements

This page describes the environmental requirements you should verify *before* running the install program. The Amanda Enterprise components required on each are described in [Downloading and Installing Amanda Enterprise Components](#).

Throughout this document, The *Amanda Client* refers the system being backed up by the *Amanda Enterprise Server*, which also called the *Amanda Server* or *backup server*.

- *Identify the hosts and filesystems to be protected using Amanda software.* The list of hosts to be protected are the Amanda backup clients. The Amanda backup clients and the filesystems or directories to be backed up on each client will need to be specified during Amanda configuration.
- *Identify the Amanda backup server.* The Amanda backup server should have access to the media (tape, disk, changers) where the Amanda client backups will be stored on tapes or the backup process is managed (in case of NDMP). Any Amanda client can also act as an Amanda server, and it is recommended that every Amanda server also be a backup client. In selecting your Amanda backup server machine, keep in mind that the server might use a significant amount of memory during backup runs, and that it must have good network connectivity to all Amanda clients.

Amanda Enterprise Server Requirements

Backup server performs various CPU, Memory, Network, and Disk intensive operations. While hardware requirements will vary based on your backup environment, we recommend a server with at least 4GB of memory and a modern quad-core server-class CPU. The bandwidth of the network link between backup server and your network switch is also very critical. If network bonding is supported by your switch, we recommend providing a bonded connection to the backup server.

- You should have at least 2 GB of disk space on the disk where Amanda Enterprise software is being installed.
- The Amanda server and clients must have the latest service packs or updates to the distribution installed.
- Amanda server and client cannot be installed on an NFS-mounted file systems. For example: You should not install the Amanda Enterprise ZMC (`/opt/zmanda/amanda`) on an NFS partition.
- The Amanda server must have the locale set to `en_*` (for example, `en_US.UTF-8`), `C`, or `POSIX`.
- In addition to the space required for the software itself, you should ensure that there is always 10% free space in the Zmanda installation and temporary directories, defined [here](#). If Amanda Enterprise runs out of space on either of these directories while a backup is in progress, the backup will fail with MySQL errors. Also ZMC will print warnings messages when there is insufficient disk space.

- The **amandabackup** user should have privileges to have crontab entries to perform scheduled backups. In some distributions, the **amandabackup** user must be added to **cron.allow** file. This file is usually in **/etc** or **/etc/cron.d** directory.
- The **amandabackup** user should have permissions to read and write to disk and tape devices (if they are used for backup). The **amandabackup** user is automatically added to **disk** user group on Linux systems. In case of Ubuntu/Debian/RHEL 6/Oracle Enterprise Linux 6 platforms, the **amandabackup** user must be manually added to **tape** user group also. It is important to restart ZMC server processes (Run **/etc/init.d/zmc_aee restart**) after adding **amandabackup** user to the **tape** user group.
- If you intend to back up an Oracle server, the **amandabackup** user must have permission to write to the **\$ORACLE_HOME** directory.
- The **amandabackup** user should have permission to write to the **/tmp** directory on the Amanda server.
- If a tape changer is the intended backup media, the backup server also requires both the **mt-st** and **mtx** package in addition to the platform-specific dependencies listed below. Please note **mt-st** package is not required for Suse/SLES distribution. If a tape drive is the intended backup media, **mt** is required. Without the appropriate packages, ZMC will return an error message when you try to view these devices.
- The directories **/etc/amanda** and **/etc/zmanda** must be on the same file system.
- Amanda clients must have open inbound TCP ports 10080 and 10081.
- Internal firewalls and selinux policy must not prevent PHP processes from using HTTP over the loopback address 127.0.0.1
- If you are using SE Linux on the server in an enforcing mode, you should run the following command to allow Zmanda Management Console to work as superuser:

```
setsebool -P allow_execstack on
```

Package Dependencies

The following packages are required on the Amanda backup server and clients:

[Linux](#)

[Solaris 10/11 and Open Solaris](#)

[Mac OS X](#)

Linux

Program Dependencies: The following programs are needed for Linux backup server and client

- GNU tar (version 1.15.1 or greater, except Red Hat/CentOS/Oracle Enterprise Linux, which supports version 1.14 or greater)
- xinetd
- ssh server
- perl version 5.8.8
- Schily tar (star 1.5final) - Required only if file system extended attributes have to be backed up.
- Gnu GPG - if data encryption is being performed

These additional programs are required on the Linux backup server

- ssh client
- gnuplot
- gettext
- mailx
- redhat-lsb (for Redhat Enterprise/Cent OS/Oracle Enterprise Linux)
- lsb-release (for Debian/Ubuntu)
- mt-st, mtx and lsscsi commands for tape backups.

Library Dependencies: The server packages are 32bit, so 32bit libraries are required. Native 32bit systems already have the required libraries. 64bit systems require the 32bit compatibility libraries. They are available for RHEL/CentOS/SLES/Oracle Enterprise/Fedora/Open Suse Linux platforms. On Ubuntu/Debian platforms, they may not be available for older platforms. On Ubuntu/Debian platforms, they may not be available for older platforms. See [Linux Debian/Ubuntu](#) for more details.

These packages are installed by default on most Linux distributions. If you need to add them, you can use **yum**, **apt-get**, or **yast**. Packages can be obtained from the distribution media or from a distribution repository (run either as root):

```
#yum install package_name
```

or

```
#apt-get install package_name
```

or

```
#yast -i package_name
```

In the case of yast on SLES, you can also use the YaST Online Update (YOU) to install new package and to keep the SLES distribution updated (which is recommended).

For example:

Amanda server installation on 64bit RHEL 6/CentOS 6/Oracle Enterprise Linux will require the following packages:

```
# yum install compat-libstdc++-33.i686 curl glib2.i686 glibc.i686 gnupg libcurl.i686 libgcc.i686  
libidn.i686 libstdc++.i686 libXp.so.6 mailx ncurses ncurses-libs.i686 readline.i686 tar xinetd zlib.i686  
openldap24-libs.i386
```

Amanda server installation on 64bit Debian 6 will require the following packages:

```
# apt-get install gettext gnuplot ia32-libs lib32gcc1:i386 libcurl3:i386 libcurl3-gnutls:i386  
libglib2.0-0:i386 libidn11:i386 libpcre3:i386 libreadline6:i386 libssh2-1:i386 lsb-release lsscsi  
mailutils mt-st mtx xinetd
```

Amanda server installation on 64bit Debian 7 will require the following packages:

```
# apt-get install lib32stdc++6 bsd-mailx gcc-4.7-base:i386 gettext libc6:i386 libffi5:i386  
libgcc1:i386 libglib2.0-0:i386 libncurses5:i386 libpcre3:i386 libreadline5:i386 libselinux1:i386  
lsb-release lsscsi mt-st mtx xinetd zlib1g:i386
```

Amanda server installation on 64bit Debian 8 will require the following packages:

```
# apt-get install bsd-mailx gcc-4.9-base:i386 gettext libc6:i386 libffi6:i386 libgcc1:i386  
libglib2.0-0:i386 libncurses5:i386 libpcre3:i386 libreadline5:i386 libselinux1:i386 lsb-release lsscsi  
mt-st mtx xinetd zlib1g:i386 libstdc++6:i386
```

Amanda Enterprise server installation on 64bit Ubuntu 14.04 will require the following packages:

```
# apt-get install gcc-4.9-base:i386 libc6:i386 libstdc++6:i386 libgcc1:i386 libglib2.0-0:i386 libffi6:i386  
libncurses5:i386 libpcre3:i386 libreadline5:i386 libreadline6:i386 libselinux1:i386 zlib1g:i386 lsb-  
release lsscsi mt-st mtx bsd-mailx mailutils xinetd
```

Note: To install 32bit packages on 64bit Debian 7 or Debian 8, please run

```
# dpkg --add-architecture i386
```

```
# apt-get update
```

Linux Debian/Ubuntu

The Amanda backup client requires 32-bit version of following Debian packages; Amanda 64-bit backup client requires the following 64-bit version Debian packages

- **libcurl3-gnutls**
- **libglib2.0-0**
- **libpcre3**
- **libidn11**
- **libssh2-1**
- **libcurl3 or libcurl4** version >=7.13.2
- **libreadline5 or libreadline6 (depending on the platform)**

The Amanda backup server requires following 32-bit package(s) on Ubuntu 10.04 (or earlier releases) and Debian 5 (or earlier releases):

- **ia32-libs or ia32-libs-multiarch** (only needed by 64-bit Debian/Ubuntu platforms)
- **libglib2.0-0 or ia32-libs-gtk** (if it is not installed by ia32-libs package)
- **libgcc1 or lib32gcc1** (if it is not installed by ia32-libs package)

Note for 64bit Debian 5 or earlier and Ubuntu 10.04 or earlier Amanda servers: The best way to install 32-bit version library packages on 64bit Debian/Ubuntu platforms is with **getlibs**, which is described in more detail in this [Zmanda Knowledgebase article](#). Alternatively, **libreadline5** and **libreadline6** 32bit version can be obtained by installing **lib32readline5** package and **lib32readline6** package respectively.

Redhat/CentOS/Oracle Enterprise Linux/SLES/Fedora/OpenSuse

The AE 32bit client requires the following 32-bit packages; AE 64-bit backup client requires the following 64-bit packages

- **curl**
- **readline**
- **glib2**
- **termcap (not required on SLES/OpenSuse and RHEL/CentOS 6)**

To run the Amanda Enterprise Backup Server (or a 32-bit Amanda client) on 64-bit RedHat/CentOS/Oracle Linux/Fedora/Open Suse requires these same packages in 32-bit versions (please note that these packages might depend on other packages and so use tools such as **yum**, **apt-get** to install these packages).

For RHEL/CentOS 5 and earlier versions and Fedora, please install the following packages:

- **glib2.i386 or glib2.i686**
- **curl.i386 or curl.i686**
- **readline.i386 or readline.i686**
- **termcap.noarch**
- **libidn.i386 or libidn.i686**
- **openldap24-libs.i386**
- **ncurses.i686**
- **libgcc.i686 (for fedora platforms only)**

For RHEL/CentOS 6, please install the following packages:

- **glib2.i686**
- **ncurses-libs.i686**
- **libstdc++.i686**
- **readline.i686**
- **glibc.i686**
- **zlib.i686**

In case of 64bit RHEL6/CentOS 6/Oracle Enterprise Linux, following yum command line should be used (this list includes direct Amanda Enterprise and indirect dependencies)

yum install glibc.i686 nss-softokn-freebl.i686 glib2.i686 libseltlinux.i686 gamin.i686 readline.i686 ncurses-libs.i686 zlib.i686 libstdc++.i686 libgcc.i686

On 64bit SLES 11 platforms, these packages are called

- **glibc-32bit**
- **glib2 (includes both 32 a 64-bit libraries)**
- **libcurl4-32bit**
- **libreadline5-32bit**
- **libidn-32bit**
- **libncurses5-32bit**

Solaris 10/11 and Open Solaris

32-bit packages for dependencies are required for all Amanda Solaris server and clients.

It is easier to install dependencies by using Solaris **pkgutil** tool available at the <http://www.opencsw.org/get-it/pkgutil/>. It automatically get all indirect dependencies. Use the following pkgutil command to install dependencies on Solaris server (you can install gnupg2 instead of gnupg):

```
pkgutil -i gtar gnupg gawk libgcc_s1 mtz
```

Use the following pkgutil command to install dependencies on Solaris Client (you can install gnupg2 instead of gnupg):

```
pkgutil -i perl gtar gnupg libcurl4
```

Mac OSX (Client only)

Download and install the package dependencies available from the [Zmanda Network download page](#) when you select Mac OSX as the client platform.

Windows (2003 Server, XP, and Vista) clients

The **amandabackup** user must exist, and be added to the **Users**, **Backup Operators**, and **Administrators** groups. Windows clients must open inbound TCP ports 10080 and 10081, and outbound TCP ports 700:800. Windows application backups (licensed as separate agents) require that the Volume Shadow Copy Service be enabled and started. To prevent excessive memory consumption on application backup clients, Zmanda also recommends that you install the following VSS hotfix from Microsoft:

[Link to Microsoft Knowledgebase](#)

See the [The Zmanda Windows Client User's Manual](#) for further details on Windows client installation.

Post installation changes on Linux/Mac OS X/Solaris Amanda Clients

- This section is not applicable for Windows clients.
- Make sure there is a reverse name lookup record in the DNS for the Amanda server (FQDN if FQDN is used in the Amanda configuration files).
- The **.amandahosts** file for the **amandabackup** user (~**amandabackup/.amandahosts**) on the client must authorize the Amanda server to run a backup.
- The **/etc/amanda/amanda-client.conf** file must be edited on the client (if the client and server do not reside on the same machine). Specifically, the following entries must be edited:

```
index_server "localhost"          # your amindexd server
tape_server  "localhost"          # your amidxtaped server
```

Change **"localhost"** to match to the hostname of the Amanda backup server.

- If secure communication is required for restoration, **ssh** must be used for restoration. **ssh** is also required when restoring to a MacOS X/Unix/Linux system that is not running the Zmanda Client software. To force the restore process to use **ssh**, edit **/etc/xinetd.d/zmrecover** on Linux systems to include the following line:

disable = yes

On MacOSX systems, run the following command: **launchctl unload -w /Library/LaunchDaemons/org.amanda.zmrecover.plist**. To re-enable the standard client restore mechanism, enter: **launchctl load -w /Library/LaunchDaemons/org.amanda.zmrecover.plist**.

Zmanda Management Console Browser Requirements

The following browsers have been tested and verified to work with the Zmanda Management Console:

- Latest version of FireFox
- Latest version of Chrome
- Recent versions of other browsers may also work

In all cases you must have JavaScript enabled. Note that if JavaScript is enabled during a session either in the browser itself or in a plug-in such as NoScript, you must log out of the ZMC and then log in again to avoid potential problems with the interface. If you are browsing over a slow connection such as dial-up, loading of the JavaScript files may time out after 15 seconds with the following error:

```
ZMC has not yet loaded required javascript files. Do you want to continue waiting
```

Clicking OK will cause the ZMC to try again for 60 seconds. If that fails, another error. If it takes more than 60 seconds to load the JavaScript files, logging out and then in again will usually give the browser enough time to cache all necessary JavaScript files and allow the ZMC to load.

Downloading the Software

To install the ZMC, you must first download Amanda Enterprise Edition with the Zmanda Management Console. Packages for all Zmanda products are available from the [Zmanda network downloads page](#). The Amanda Enterprise Edition with ZMC package includes:

- Amanda Enterprise Server
- Zmanda Client for the server
- The Zmanda Management Console
- all of the dependencies for Zmanda Management Console

Upgrading from earlier version of 3.3.x release

You can upgrade from earlier version of 3.3.x by running the installer. The installer detects the older version of 3.3.x release and automatically upgrade to newer release. Please check 3.3.x release notes for any caveats.

During upgrade, the installer will prompt the user to preserve configuration and catalog information. It is important to preserve configuration and catalog if you want to restore from older backup images and use backup set policies created using earlier version of 3.3.x release.

Installing Amanda Enterprise server on the Same Server as Zmanda Recovery Manager for MySQL 3.6 server

ZRM and Amanda Enterprise must be installed in a specific order to work. Please perform the following procedure as superuser on the backup server. If you need to install ZRM 3.6 on top of Amanda Enterprise 3.3.x, please contact Zmanda Support team.

- Install ZRM 3.6 server. Do NOT stop ZRM ZMC services at this stage.
- Install Amanda Enterprise 3.3.x server and Amanda Enterprise patches. Use different port from ZRM services for web server (both http and https ports) at the time of Amanda Enterprise server installation.
- Stop ZRM ZMC services
`/etc/init.d/zmc_zrm stop`
- Stop Amanda Enterprise ZMC services
`/etc/init.d/zmc_aee stop`
- Start ZRM ZMC services
`/etc/init.d/zmc_zrm start`
- Start Amanda Enterprise ZMC services
`/etc/init.d/zmc_aee start`

Acquiring and Installing License file

After you have purchased a base license and any feature licenses (such as application module licenses), the [Zmanda Network Downloads page](#) will include an option to download a license key (**zmanda_license**). Download this file and save the file as **/etc/zmanda/zmanda_license** on the Amanda Enterprise Linux server, then make sure that the file permissions are set to 644 and owner is **root**. On Solaris servers, The license file location is **/opt/zmanda/amanda/etc/zmanda/zmanda_license**.

Secure Socket Layer (SSL) Certification

Although the Zmanda Management Console is shipped with pre-packaged Apache SSL certificate to get you started, Zmanda recommends you purchase (or create your own self-signed) SSL certificates and distribute them to all the browsers from which you wish to access the ZMC. The pre-packaged certificates are not secure (as they are shared by all Zmanda customers). These generic certificates will also generate security warnings on some browser versions.

Zmanda recommends that you either 1) [Create self-signed certificates](#) and distribute them to all the client machines that require access to the ZMC, or 2) Distribute certificates from a recognized Certificate Authority. Option 1 (self-signed certificates) is free, and is adequate for most organizations that deploy ZMC servers and the machines that access them "behind the same firewall".

If using a certificate from a recognized Certificate Authority, your browser will automatically create the secure connection with no errors or warnings.

If using a self-signed certificate, you must then deploy a mechanism to get the relevant browser(s) to accept this new root CA. One method is to generate the certificate using a special format that can be directly imported by common web browsers, and then providing a link on a secure intranet for ZMC users to download (web browsers automatically display the import dialog if the file is in the correct format and sent by the intranet web server using the correct mimetype). [PKCS12](#) (now part of OpenSSL, provides a mechanism to distribute self-signed private key certificates in a number formats recognized by different browsers.

Another approach is to manually add the new self-signed root CA to the root CA list of the client system, which will automatically provide access to the new CA for all web browsers on the client system. [This article](#) covers the procedures for doing this in a Microsoft Windows server environment.

For more details on certificate validation issues, see [this article from OpenSSL](#).

The ZMC Rapid Installer

For installation using a method other than the Rapid Installer, please see the sections that follow.

1. Copy the Rapid Installer binary file to the host where the given component will be installed.
2. Log in to the host as the superuser.
3. Make sure that the Rapid Installer binary file (it may have .bin extension also) is executable. For example:

```
# chmod +x amanda-enterprise-3.3.2-installer.run
```

4. Run the installer by double-clicking on it, or enter the following command line:

```
# ./amanda-enterprise-3.3.2-installer.run
```

5. The Rapid Installer then starts. Follow the on-screen instructions.

Important Note: When prompted to choose the **Zmanda Web Server protocol**, we strongly recommend that you choose **https** for security reasons. Even if you choose **http** for browser/ZMC communication, the ZMC still requires HTTPS for internal communication purposes, and will therefore prompt you for an SSL port during installation in all cases.

Note: The installer performs several tasks after creating and populating the Zmanda directories. These are completed after the progress bar (which only tracks the archive extraction) shows 100% completion. These tasks take time. Please wait till they complete.

6. After the Amanda Enterprise binaries have been extracted and installed, the Zmanda Management Console is launched and the **readme** file is displayed. The **readme** file includes the default Zmanda Management Console username and password. You can now login to the console using any supported browser and begin configuring backup sets.

Rapid Installer Command Line Options

Run the installer with the **--help** option to see what command line parameters are available.

--help

Display the list of valid options.

--version

Display product version information.

--optionfile *optionfile*

Use command line parameters specified in *optionfile*.

--mode *mode*

Choose the installation mode, where *mode* is **gtk** (the default), **xwindow**, **text**, or **unattended**. The **unattended** mode allows users to install the product without user interaction in an automated manner.

--debugtrace *debugtracefile*

Installer debug logs will be in *debugtracefile*. This option is not required for installation.

--installer-language *installer-language*

Choose installer language. This is not supported by Zmanda installer. Only *en* (English) is supported.

--apache_server_port *apache_server_port*

The ZMC Web Server port (default is **80**).

--apache_server_ssl_port *apache_server_ssl_port*

ZMC Web Server SSL port (default is **443**).

--mysql_port *mysql_port*

Specify the ZMC MySQL Server port (default is **3037**).

Uninstalling Amanda Enterprise

Important Note to Customers of ZRM for MySQL: Before running the uninstall script for Amanda Enterprise, you must first stop the ZRM for MySQL services from running (enter **/etc/init.d/zmc_zrm_init stop** as root), then restart it manually after the uninstall script completes (**/etc/init.d/zmc_zrm_init start**).

You can uninstall Amanda Enterprise by running the uninstall script located at **/opt/zmanda/amanda/uninstall** on the backup server. Using this script, you can remove the Amanda Enterprise binaries, with the option of leaving configuration files intact. Follow the on-screen instructions after running the script.

The uninstaller program does not remove the backup scheduled in the crontab. They will have to be removed manually.

Accessing Zmanda Management Console

Zmanda Management Console is accessed using a web browser. Amanda Enterprise is supported on latest version of Internet Explorer, Mozilla FireFox and Google Chrome.

Please access <https://<host name of the Amanda server>/> if you are using default port. If you had changed the port number using installation, please use URL <https://<host name of the Amanda server>:<port number>/>

The default port is 443. Please make sure access to this port is allowed by network configuration (such as firewalls).

Amanda Enterprise File Locations

These are the file locations on Linux server.

Amanda Enterprise files are installed in following directories:

Amanda Server File(s)	Location
CLI Executables	/usr/sbin
Man pages	/usr/share/man
GUI/Web Components (including GUI log files)	/opt/zmanda/amanda, /opt/zmanda/common
License key file	/etc/zmanda/zmanda_license
Backup Set and other configuration files	/etc/amanda, /etc/cron.d, /etc/logrotate.d
Libraries	/usr/lib
Amanda-controlled executables	/usr/lib/amanda
Amanda Perl Libraries	/usr/lib/amanda/perl

Amanda Debug log files	/var/log/amanda/
Backup images (vtapes) and staging area (holding disk)	/var/lib/amanda (default location for backup to disks)

On Solaris server, the Amanda Enterprise files are located at **/opt/zmanda/amanda** directory. For example: The configuration files under **/etc/amanda** in Linux server will be found under **/opt/zmanda/amanda/etc/amanda** on Solaris server.

Do not directly delete or change any of these files or directories unless directed to do so by the Zmanda Support Team. Changing any of these files directly can result in failed backups and other problems.

Zmanda provides Amanda server disaster recovery solution to protect Amanda server. Please contact Zmanda Sales or Zmanda Support Team for more details.

You should monitor free space availability in **/tmp**, **/etc/amanda** and **/opt/zmanda/amanda** directories. At least 10% free space must be available in these file systems. Amanda stores the backup index, Zmanda Management Console database and temporary files in these locations, and thus these directories can fill up over time. The **/tmp** directory is also used to store backup images temporarily during the restore process, which can fail if it runs out of space. Lack of space in these directories can affect Amanda functionality and performance.

Amanda Debug log files are collected by Zmanda support tool for troubleshooting Amanda problems. Amanda debug file for a process or command instance are kept for four days for each backup set.

Accessing Zmanda Management Console

Zmanda Management Console is accessed using a web browser. Amanda Enterprise is supported on latest version of Internet Explorer, Mozilla FireFox and Google Chrome.

Please access <https://<host name of the Amanda server>/> if you are using default port. If you had changed the port number using installation, please use URL <https://<host name of the Amanda server>:<port number>/>


The default port is 443. Please make sure access to this port is allowed by network configuration (such as firewalls).

If you are running Amanda Enterprise and ZRM Enterprise products on the same server, you will have to use port number configured during Amanda Enterprise installation to access both Amanda Management console and ZRM Management console.

ZMC Login Page

This is the first page of Zmanda Management Console. If you are running Amanda Enterprise and ZRM Enterprise products on the same server, you will get a product selection page.


The default user is **admin/admin**. Please change the password in the ZMC **Admin Users** page. This user is different from Zmanda Network user and operating system user.



The screenshot shows the Zmanda Management Console Login page. At the top, there is a header bar with the text "Zmanda Management Console Login". Below this, the "zmanda Management Console" logo is displayed. The login form includes fields for "Username: *" and "Password: *", each with a text input box. To the right of the password field is a "Login" button. Below the password field is a checkbox labeled "show password". Further down, there are two checkboxes: "Check Server Installation?" (checked) and "Sync Backup Sets?" (checked). At the bottom, there is a link that says "Can't access your account?" in blue text.

If you are making changes to Backup Sets manually, you should enable **Sync Backup Sets** at the time of login process. Otherwise, ZMC will present incorrect information.

ZMC can check the Amanda server installation and dependencies at the time of login. This process can take time. You should enable **Check Server Installation** when you are logging in the first time or if you logging in after hardware or OS upgrade. If there are errors in server installation, please resolve them before a backup or restore job is started. Please note that these checks are not the same as the client checks performed by the **Check Hosts** action in the **Backup What** page.



Lost Password

If you do not remember the password, please click **Can't access your account?** link. You will see the following section as shown below. Please enter the ZMC user name in the **Lost Password** section.



zmanda
Management Console

Enter your Username and a new password will be assigned to your account. Your new password will be emailed to you, if email service has been configured on this server for remote delivery (or local delivery, if your ZMC account email address is local to this server).

Alternatively, to reset the ZMC admin user password to "admin", run the command:

```
/opt/zmanda/amanda/apache2/htdocs/zmanda-ae/ZMC/scripts  
/reset_admin_password.sh
```

Username: *

Please enter the ZMC user name in the **Lost Password** section. Please note that password will be mailed to the mail address registered to the ZMC user account. Please note that email service should be configured on the Amanda backup server to receive the lost password email.

Zmanda Network Authentication

ZMC will ask [Zmanda Network](#) User name and password if you login as **admin** user. This authentication allows ZMC to connect to Zmanda Network to authenticate the user as well as obtain alerts including security alerts.

A screenshot of a Windows-style dialog box titled "Zmanda Network Authentication". It features a yellow warning icon and a message: "Please enter your Zmanda Network user name and password to connect to the Zmanda Network and receive update alerts, including security fixes." Below the message is a link for "(Advanced Details)". There are two input fields: "Zmanda Network Username:" and "Zmanda Network Password:". At the bottom right are "Cancel" and "Save" buttons.

Zmanda Network Authentication

! Please enter your [Zmanda Network](#) user name and password to connect to the Zmanda Network and receive update alerts, including security fixes.

[\(Advanced Details\)](#)

Zmanda Network Username: *

Zmanda Network Password: *

Cancel Save

Zmanda Network Authentication is not supported when Web proxy server is in use. If there is a web proxy server or there is no internet connection on the Amanda server, please select **Cancel** button.

Zmanda Network Authentication is performed every time ZMC **admin** user logs in.

ZMC Work Flow

Amanda Backup server configuration is done in terms of backup sets. The backup set is a grouping mechanism that simplifies and optimizes backing up many systems with different backup requirements. It lets an administrator define a set of backup policies (what, how, where and when) to different backup runs.

All ZMC actions (backup, restore, reporting, and monitoring) are performed in the context of backup sets.

Backup sets cannot share target media volumes. For example, two backup sets can use the same tape changer simultaneously, but each should have a separate set of slots.

Multiple backup sets are useful for protecting a large number of systems with different backup requirements, but many organizations with less complex backup requirements can define a single backup set to meet their needs. For example, on a network that includes servers with high rate of data change along with desktop systems that change more slowly, you would probably want to create one backup set for the servers, and another backup set for the desktops.

All data to be backed up can be organized in backup sets based on:

1. Backup policy: Backup schedule, how often a full backup is performed and how long the data is retained.

2. Backup media: All data from a backup set are backed up using same device (disk storage or tape drive).
3. Backup window: Data in different backup sets can be backed up at the same time to different devices. There can be parallelism within the backup set from the client to the backup staging area.

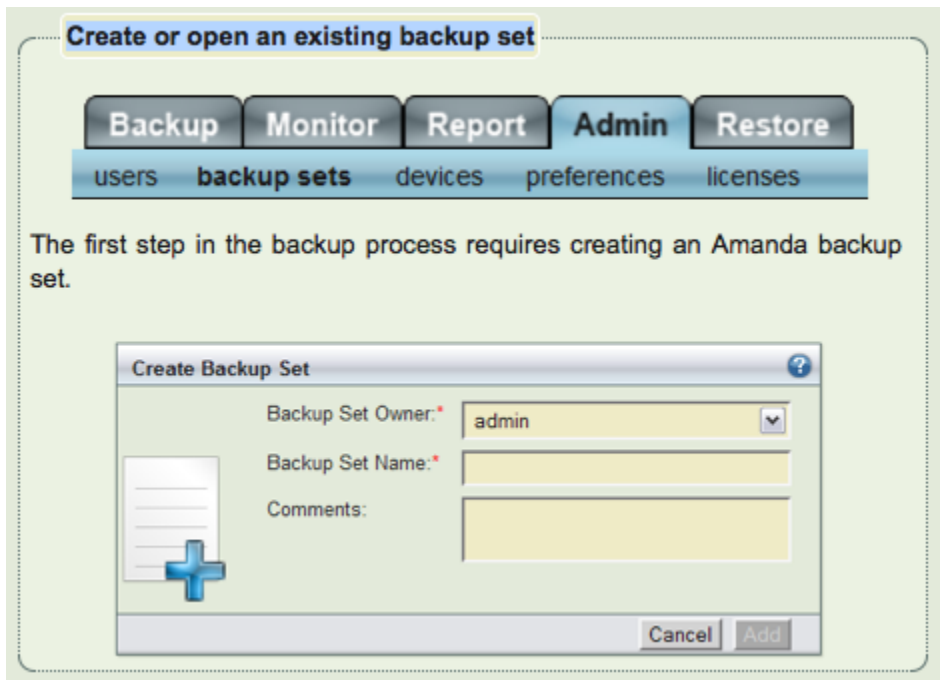
What a Backup Set Contains

A backup set is a uniquely-named record of backup policies, including:

- hosts, directories, and files to exclude.
- the backup target, which can be a tape device or disk (via holding disk or virtual tape)
- the type of backup to perform (full or incremental); schedules are automatically configured by Amanda.

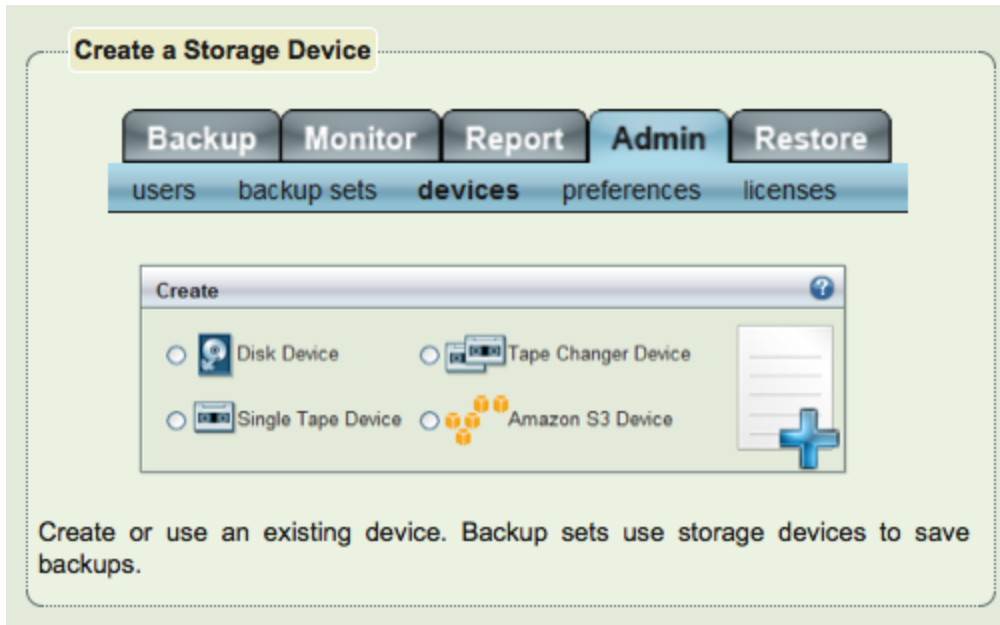
Amanda ensures that two backup sets do not write to the same tape (or to the same directory if disk media is used for backup). Amanda also enforces that it does not accidentally overwrite a tape/disk belonging to another backup set.

Backup sets are created in **Admin Backup sets** page as shown below.



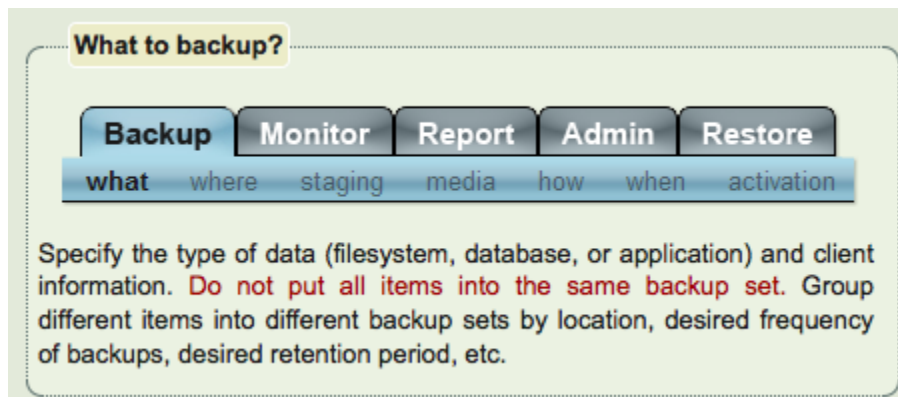
The screenshot shows the Zmanda Admin interface. At the top, there's a navigation bar with buttons: Backup, Monitor, Report, Admin, and Restore. Below these are sub-tabs: users, backup sets, devices, preferences, and licenses. The 'Admin' button and 'backup sets' sub-tab are highlighted. Below the navigation bar, a text box states: "The first step in the backup process requires creating an Amanda backup set." Below this text is a dialog box titled "Create Backup Set". The dialog box has a question mark icon in the top right corner. It contains three fields: "Backup Set Owner:" with a dropdown menu showing "admin", "Backup Set Name:" with an empty text box, and "Comments:" with a larger empty text box. There is a blue plus icon on the left side of the dialog box. At the bottom of the dialog box are "Cancel" and "Add" buttons.

You can create Storage devices for backups in **Admin devices** page. Generic information about the devices are configured in this page. All devices supported by Amanda Enterprise can be configured in this page. After creating device configuration, you can bind the device to the backup set.



The **Backup** tab lets you configure and edit backup sets. The six sub-tasks under the **Backup** tab let you specify the parameters of the currently selected backup set:

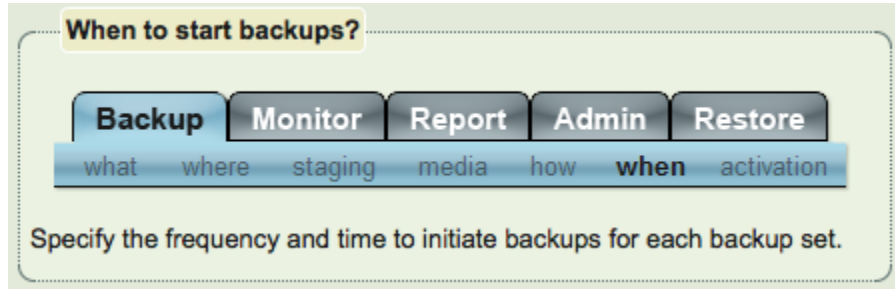
- [Backup What](#) specifies what host system and directories to include in the backup set.



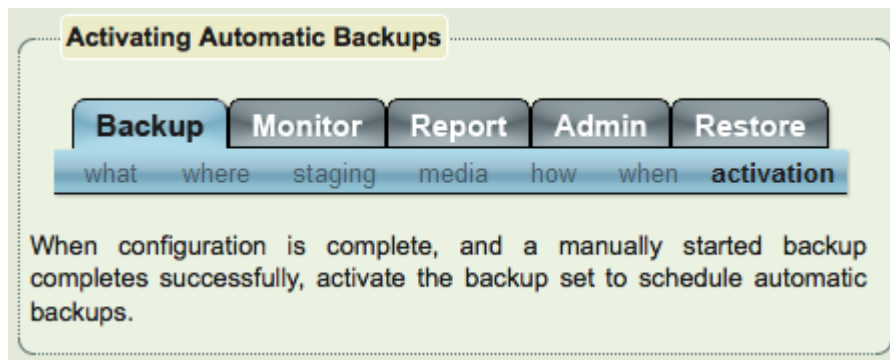
- [Backup Where](#) specifies the target device for the backup set(i.e., tape, hard disk, etc.)



- [Backup When](#) specifies the scheduling parameters for Amanda's intelligent scheduling.



- [Backup How](#) specifies the type of backup to be performed (full or incremental).



- [Backup Staging](#) specifies the staging area (Amanda holding disks) for the backup set.
- [Backup Media](#) manage media volumes used by the backup set.

Once this information has been specified,

- [Backup Activation](#) enables the backup jobs to be activated.

Accessing Zmanda Management Console

Zmanda Management Console is accessed using a web browser. Amanda Enterprise is supported on latest version of Internet Explorer, Mozilla FireFox and Google Chrome.

Please access <https://<host name of the Amanda server>/> if you are using default port. If you had changed the port number using installation, please use URL <https://<host name of the Amanda server>:<port number>/>

The default port is 443. Please make sure access to this port is allowed by network configuration (such as firewalls).

If you are running Amanda Enterprise and ZRM Enterprise products on the same server, you will have to use port number configured during Amanda Enterprise installation to access both Amanda Management console and ZRM Management console.

ZMC Login Page

This is the first page of Zmanda Management Console. If you are running Amanda Enterprise and ZRM Enterprise products on the same server, you will get a product selection page.

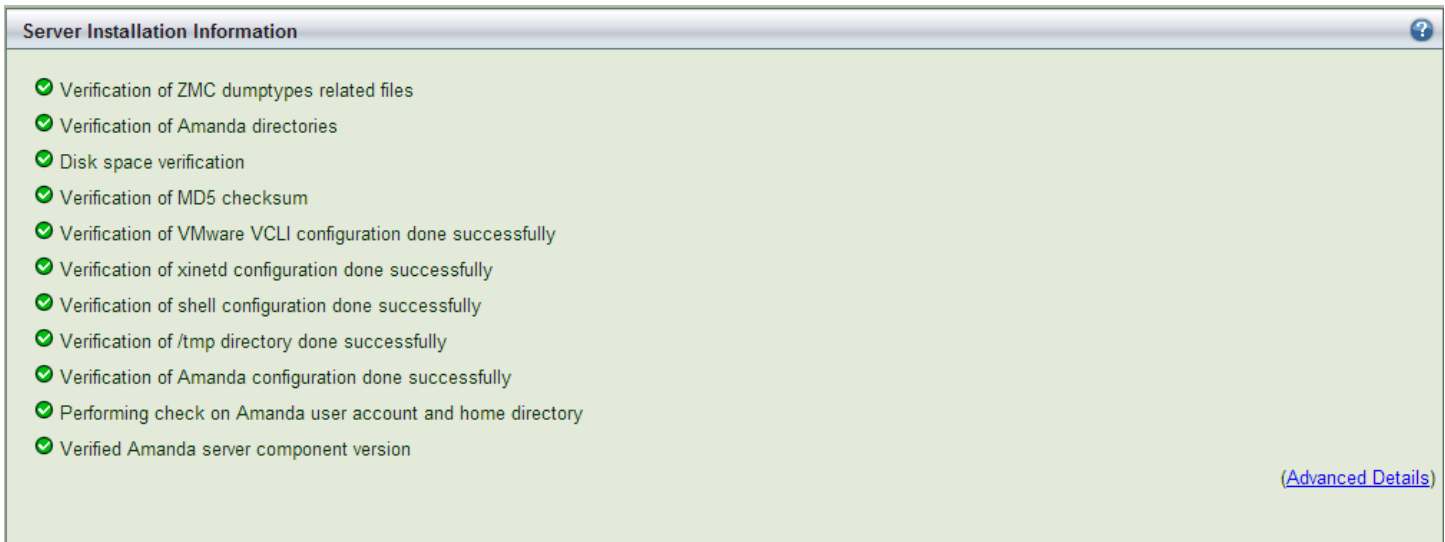
The default user is **admin/admin**. Please change the password in the ZMC **Admin Users** page. This user is different from Zmanda Network user and operating system user.



The image shows the Zmanda Management Console Login page. It has a title bar that says "Zmanda Management Console Login". Below the title bar is the Zmanda logo and the text "Management Console". There are two input fields for "Username: *" and "Password: *". To the right of the password field is a "Login" button. Below the password field is a checkbox labeled "show password". There are two checkboxes: "Check Server Installation?" and "Sync Backup Sets?", both of which are checked. At the bottom, there is a link that says "Can't access your account?".

If you are making changes to Backup Sets manually, you should enable **Sync Backup Sets** at the time of login process. Otherwise, ZMC will present incorrect information.

ZMC can check the Amanda server installation and dependencies at the time of login. This process can take time. You should enable **Check Server Installation** when you are logging in the first time or if you logging in after hardware or OS upgrade. If there are errors in server installation, please resolve them before a backup or restore job is started. Please note that these checks are not the same as the client checks performed by the **Check Hosts** action in the **Backup What** page.



Lost Password

If you do not remember the password, please click **Can't access your account?** link. You will see the following section as shown below. Please enter the ZMC user name in the **Lost Password** section.



zmanda
Management Console

Enter your Username and a new password will be assigned to your account. Your new password will be emailed to you, if email service has been configured on this server for remote delivery (or local delivery, if your ZMC account email address is local to this server).

Alternatively, to reset the ZMC admin user password to "admin", run the command:

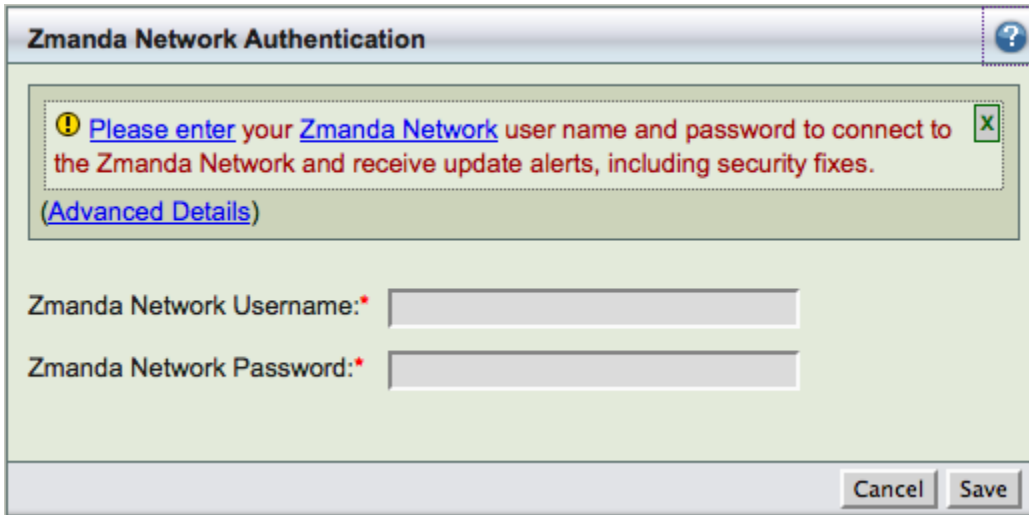
```
/opt/zmanda/amanda/apache2/htdocs/zmanda-ae/ZMC/scripts  
/reset_admin_password.sh
```

Username: *

Please enter the ZMC user name in the **Lost Password** section. Please note that password will be mailed to the mail address registered to the ZMC user account. Please note that email service should be configured on the Amanda backup server to receive the lost password email.

Zmanda Network Authentication

ZMC will ask [Zmanda Network](#) User name and password if you login as **admin** user. This authentication allows ZMC to connect to Zmanda Network to authenticate the user as well as obtain alerts including security alerts.

A screenshot of a Windows-style dialog box titled "Zmanda Network Authentication". It features a yellow warning icon and a message: "Please enter your Zmanda Network user name and password to connect to the Zmanda Network and receive update alerts, including security fixes." Below the message is a link for "(Advanced Details)". There are two input fields: "Zmanda Network Username:" and "Zmanda Network Password:". At the bottom right are "Cancel" and "Save" buttons.

Zmanda Network Authentication is not supported when Web proxy server is in use. If there is a web proxy server or there is no internet connection on the Amanda server, please select **Cancel** button.

Zmanda Network Authentication is performed every time ZMC **admin** user logs in.


ZMC Backup What page

The **Backup What** page specifies what is to be backed up: what clients, and what directories (or applications) on each client. ZMC can back up the whole network or a portion thereof, all from one central server.

To organize the backup in an efficient manner, ZMC divides the Enterprise into backup sets, and subdivides each backup set into Host/Directory pairs called Disk List Entries (DLEs). The DLEs can also specify directories and files to be excluded from the backup. Encryption and compression options can also be applied at the DLE level.

Warning: You should not change backup set parameters while a backup run for that set is in progress. You can check the status of backup runs for a backup set by going to the [Monitor](#) page.

The top portion of the of the page lets you create and edit *backup objects*, which define the file system, database(s), or applications you intend to back up. After you select a type from one of the dropdown menus (or select an existing object from the list at the bottom of the page), appropriate options for that backup object are displayed:



Selecting What to Back Up

Supported object types include the following. All objects are licensed. Note that the ZMC indicates how many licenses have been purchased and how many remain available for those object types that require licensing. Application Agents are described in more specific detail [here](#).

Linux/UNIX/Mac OSX file system

Lets you select a Linux, UNIX, or Mac OSX host name and directory path for backup. Also lets you specify files to exclude, and whether encryption and compression should be used. Amanda does not cross file system boundaries and thus each file system/partition should be entered as its own entry (separate DLE object).

Network/CIFS file system

Lets you select a Common Internet File System by host name, share name, and domain (you will need the username and password to access the share). For further details on CIFS backups, see ["Backing Up and Restoring Common Internet Filesystem Shares."](#)

NDMP

Lets you perform backup of NDMP appliances/filers using NDMPv4 protocol. Netapp filers, Sun Unified Storage and BlueArc Titan Storage are supported. For further details on NDMP backups, see ["NDMP appliances"](#).

VMWare ESX

Lets you select a VMWare ESX server to backup guest VM images. For further details on VMware ESX backups, see ["VMware Vsphere and ESX"](#).

Microsoft Sharepoint

Lets you select a host running Microsoft Sharepoint 2007 server or WSS 3.0 server to be backed up. For further details, see ["Backing Up and Restoring Microsoft Sharepoint Servers."](#)

Oracle on Window

Lets you select a Windows Oracle server by hostname for backup, and whether compression should be used. For further details, see ["Backing Up and Restoring Oracle Servers \(Windows\)."](#)

Microsoft SQL Server

Lets you select a Windows SQL server by hostname for backup, and whether compression should be used. For further details, see ["Backing Up and Restoring Microsoft SQL Servers."](#)

Microsoft Exchange

Lets you select a Windows Exchange server by hostname for backup, and whether compression should be used. For further details, see ["Backing Up and Restoring Microsoft Exchange Servers."](#)

Windows File System (NTFS/ReFS)

Lets you select a Windows file system for backup, and enable compression if desired.

Windows System State

Allows you to back up the MS Windows System State, and to enable compression as desired.

Windows Template

Lets you select a template that you create on the Windows backup client that defines what is to be backed up. See [Using the Zmanda Windows Client Configuration Utility](#) for details on template creation. It also allows you enable compression as desired.

Oracle on Linux/Solaris

Lets you select an Linux or Solaris Oracle server for backup, and to enable compression if desired. You must also specify an SID List Name for the Oracle database. For further details, see ["Backing Up and Restoring Oracle Servers \(Linux/Solaris\)."](#)

Solaris File System


Lets you select a Solaris file system by hostname and directory path for backup. Also lets you specify files to exclude, and whether encryption and compression should be used. For further details, see ["Solaris client"](#).

PostgreSQL

Lets you select a PostgreSQL database for backup by specifying a hostname and data directory, and to enable encryption and compression if desired. For further details, see ["Backing Up and Restoring PostgreSQL Servers."](#)


Click the **Add** button at the bottom of the page to create a new entry, or select an entry from the table to edit. You can also duplicate an existing entry (see below) and then edit the entry before using the **Add** button to add it to table of backup objects..



Create Linux/UNIX Object (DLE) in list: files-s3

Host Name*: Encryption: [Download](#) 

Directory/Path*: Compression:

Exclude Files:

Advanced Options 

Estimate:  ☐ Reliably Accurate ☒ Usually Accurate  ☐ Historical Average

Alias:

Client Max Backups: maximum permitted simultaneous DLE backups for this Amanda client

Strategy*:

Disable Staging: ☐ (not recommended)

Amanda Backup Client Application

Extended Attributes: ☐

Recommended:

Override:

Managing Backup Objects

Objects to be backed up (also known as "Disk List Entries" or "DLEs" in Amanda Community Edition) are listed in the table at the bottom of the page.

Object Type

Object Type can be any of the object types described in the section above.

File System Options

The options for backing up file systems are essentially the same regardless of platform. Options pertaining to other object types are discussed in relevant sections of the [Zmanda Application Agents Guide](#)

Host Name

Specify or select a host name (or IP address) to back up. Note that a backup set cannot include duplicate objects (a.k.a. DLEs); if you attempt to add host/directory combination that already exists in the backup set, an error is displayed. For purposes of host and pathname collision detection, the characters :, \, and / are considered the same by the ZMC.

Directory/Path

Specify a directory on the currently selected host to back up. Unless excluded (see below), all directories/files below this directory are recursively included in the backup. There is a limit of 255 characters for a directory name. Note that Amanda will not cross file system boundaries when completing backups on Linux filesystems. For example, if \ is specified as the directory to back up, \tmp will not be included in the backup if it resides on a separate file system. Solaris filesystems backup will cross filesystem boundaries. So it is important to exclude network directories (such as NFS or CIFS mounted directories).

Exclude Files

Lists Files to be excluded from the backup. Space-separated, quoted shell/tar glob expressions (i.e. * and ? wildcards) are allowed to specify multiple files and paths. See [Exclude Specifications](#) for more details. For example: If you are backing up root file system in Solaris, you can exclude `"/platform" "/system" "/proc" "/tmp" "/dev"` directories/file systems.

Encryption & Compression Options

Lists Encryption and Compression choices. Encryption and Compression of data is described [here](#).

Host Check

Shows the verification status of the object. Adding an object triggers a verification of the Amanda client/directory combination. Depending on the result of the last verification, the icon will change to a red stop sign (error), a dash (meaning the entry has not yet been checked), or green check mark (all systems go). You can also check all of the entries at once by clicking the **Check All Hosts** button at the bottom of the page. The message box displays the status of the **Check All Hosts** process: when the check is completed, and whether any errors were found.

Add

Clicking this adds any backup set changes to the server.

If any of the backup objects (also known as disk list entries) have a status other than the green checkmark, a confirmation dialog is displayed. This allows you cancel committing entries that may have problems (or may not have been checked yet). If you commit changes that have errors or that are unverified, you must be sure to verify the problem objects and correct any problems, or else backup failure may result.

Exclude Specifications

Excluding files can optimize the performance of the backup set, especially one that would otherwise back up an entire host from the root directory down.

Exclude specifications depends on the object type. The patterns supported are different for Linux/Solaris/Mac OS X and Windows. Please see next two sections for the details.

Linux/Solaris/Mac OS X filesystems - GNU tar (default)

The Linux/Solaris/Mac OS X filesystems uses the [GNU-tar utility](#) (unless [extended attribute](#) backup is enabled) which supports [exclude patterns](#). If backup of extended attributes are enabled (**schily tar** is used), exclude pattern cannot be specified.

The ZMC can accept one or more explicit pathnames or wildcard patterns per backup object/disk list entry, separated by a space. Some simple examples for GNU tar, Windows clients and Schily tar below.

GNU tar

- You can explicitly exclude any file or directory by pathname. For example, it is recommended that you avoid backing up staging areas for backup sets, so if you are backing up a root directory (/)

that includes staging area `/var/lib/amanda/staging/`, the exclude specification would be `./var/lib/amanda/staging`. If the backup object/DLE is set to back up `/var`, the exclude specification would be `./lib/amanda/staging`. The pathname in exclude specification should be relative to the DLE directory.

- To exclude any file or directory that ends in ".log" e.g. ppp.log, and Xfree86.0.log, specify ***.log**
- To exclude any file or directory with the string "log" e.g. logfile, maillog, syslog, and ppp.log, XFree86.0.log, specify ***log***
- To exclude any file or directory that starts with string "cron" and ends in ".gz" e.g. cron.1.gz, cron.2.gz, and log/cron.1.gz, specify ***cron*.gz**
- The question mark can be used to specify a single character. e.g. to exclude log.1 and log.2, specify log.?
- Multiple patterns are allowed, separated by spaces. For example, specifying following would exclude **this.doc**, **that.doc**, and **everythingelse.doc**, and would also exclude the **Misc Files** directory. The string must be quoted if it includes spaces (such as "**Misc Files**").
"/*.doc" "/Misc Files"
- Use the backslash to escape any double-quote characters included in the file or pathname itself. For example, specifying ***.sh" "foo bar" "some\"quote"** excludes ***.sh**, **foo bar**, and **some"quote**.
- The ZMC will save the exclude pattern in a standard format that includes explicit quotes (and escapes any characters that require it). For example, if you specify:

```
exclude "*.doc" *.txt
```

is saved as

```
exclude "*.doc" "*.txt"
```

- Shadowed excludes are automatically deleted whenever the user saves any edit to the backup object/DLE. Thus a backup object containing an exclude list of:

```
exclude "*.doc"  
exclude "*.txt"
```

will show only ***.txt** in the **Exclude** form field, and the first exclude (***.doc**) is removed from the backup object when any edits are saved.

- Although exclude patterns are case sensitive (except for Windows; see the next bullet item), the Amanda Enterprise catalog displays all pathnames in lowercase (even though it restores the original case). Because of this, it will appear that files you intended to exclude are being included in the backup because there were duplicate names with different cases included in the backup.

Schily tar

- Format should be just the name of the directory or file in the first level directory (of the directory being backed up) with no leading nor trailing slashes, dots, etc. To exclude `/path/to/top-level/foo/` from a backup of `/path/to/top-level/`, `"foo"` should be specified. This will only match `"foo"` in the first level directory and thus does not match `"foo"` anywhere else in the directory tree. Thus, there is no way to just exclude `/path/to/top-level/foo/a` because this is a second level down.
- Globbing characters and simple regex, for instance `[A-Z]`, `*` character can be used to match files and directories in the first level.

Windows filesystems - Zmanda Windows Client

Windows filesystems support wildcards in the exclude specification. Wildcards `"*"` (match one or more character) and `"?"` (match exactly one character) are supported. The pathname in exclude specification can be absolute or relative to the DLE directory. For example: We are backing up `C:\Data` directory. We would like to exclude all files with `*.jpg` extension. The exclude specification should be `"*.jpg"`. You can specify the directory name with the exclude pattern. For example: To exclude `*.exe` from `C:\Data\Test` and `C:\Data` is being backed up, specify `"Test*.exe"`, To exclude a folder, specify `"C:\users\all users\"`.

The list of patterns in the exclude specification for Windows file systems should be separated by space character.

When specifying exclude patterns for Windows clients, pathnames are case-insensitive. The `./` or `.\` characters do not work.

Reasons to Exclude Files

There are a number of different reasons to exclude files from a backup set.

1. You do not want the excluded data to be part of any backup set. Such data tends to be quite small in quantity and does not save much.
2. You want to exclude the data from Compression or Encryption options, saving CPU cycles, network bandwidth, and total time to run backups.
3. Lastly, if the data on the host is organized in such a manner that implementing multiple top level directories in different disk lists makes no sense, exclude lists may allow users to organize the backup with possibility of faster restore for some files. The files would be excluded from the general backup and kept in separate backup object/DLE(s).

Compression and Encryption

This section describes the compression and encryption options common to most of the backup object types.

Data Compression

ZMC compresses the data on the Amanda server or client. Amanda client compression resulting in more efficient use of bandwidth when the backup image is sent to the server. ZMC supports compression using **gzip**, which creates archives that can be extracted almost universal. Compression levels - **fast**, **best** and

custom can be specified. The **fast** compression provides smaller backup window. The **best** compression likely provides smaller backup sizes. The **custom** compression provides user an option to specify different compression command. Most common custom compression program used is `/bin/pigz` - this compression command provides faster, parallel compression that uses multi-core CPUs better.

Many Tape drives have built-in hardware compression. There are many advantages in allowing such hardware to handle the compression task.

TIP: Devices that use a proprietary compression can fail or otherwise becomes unusable, presenting difficulties in restoring from backups that were written to it. For images and other pre-compressed files, do not consume the backup window by pointlessly re-compressing them.

Data Encryption

Backup data should be secured as carefully as you would protect the live version. Encrypting backup data adds a layer of protection against misuse.

Backup encryption can be performed on the server or on the client. It is important to store encryption key passphrases and certificates securely. Backups cannot be retrieved if the passphrase files or certificates are lost.

Encryption passphrase or keys must be provided during the restore process for the backup images to be decrypted. Amanda Enterprise does not provide encryption key management. Customers should make sure the encryption keys are backed up.

Note: Encryption is a CPU intensive task. Enable it with care.

Linux/Solaris/Mac clients

ZMC encrypts the data on the Amanda server and Linux/Solaris clients using the [amcryptsimple\(8\)](#) program, which uses Gnu Privacy Guard (GPG) to perform symmetric data encryption.

Encryption passphrases are stored in the **amandabackup** user directory on the server. It is important to keep the encryption passphrases (default passphrase - `/var/lib/amanda/.am_passphrase`) safely and securely. The data **cannot** be restored without the passphrase. It is important to backup the passphrases on regular basis by adding `/var/lib/amanda` directory as an Amanda DLE **without enabling encryption** for that DLE. Also keep a backup of the passphrase in another secure location (for example: printed hardcopy).

ZMC uses encryption passphrase from the Amanda server to do restoration of Linux/Solaris/Mac encrypted backup images. You should copy the appropriate encryption passphrase file to the Amanda server - `/var/lib/amanda/.am_passphrase` (Linux/Mac) or `/opt/zmanda/amanda/amanda/.am_passphrase` (Solaris) before doing restoration of client encrypted backups using ZMC.

You can use command line tool - `amrecover` running on the Amanda client to restore client encrypted backups. The command line tool uses the encryption passphrase from the Amanda client to decrypt client encrypted backup images.

Windows clients

Backup encryption on Windows client is performed using AES encryption keys. Please see [Windows Client](#) manual for more details.

Advanced Options

This section describes all of the possible Advanced Options that may be displayed for any of the object types. Some of these options may not apply (or be displayed in the dialog) depending on the type of backup being configured.

Extended Attributes

This is the default for Windows and Mac OS X filesystems.

This is an optional feature for Linux, UNIX, and Solaris filesystems. Enabling this option in these cases selects a different archive program used for backing up the given object type. When this option is enabled, Amanda Enterprise uses (and requires) Schily tar instead of GNU tar as the archive program. Schily tar is required on the Amanda client and is not installed by default on Linux. Schily tar package is available for download from Zmanda Network.

If the backup object is a Solaris ZFS file system, Extended Attributes refer to ZFS Access Control Lists (ACLs); see [Solaris Client](#) for details.

Please be aware that exclusions are not supported when selecting Extended Attributes for Linux, UNIX, and Solaris file systems.

Estimate

Selects the method used for estimating the backup window. You can choose from a number of options that balance the requirements of accuracy vs. speed. The "fastest" method can be accurate enough if the backup source remains relatively constant in size; the "Always Accurate" option may be too slow given the backup window, or may not be available from the given backup client. Options not supported for a given backup client are grayed out so that you cannot select them.

Alias

Alternate name can be provided for the backup object. This is useful when you are backing up the same file system with different exclude lists. You may have to do this to divide a large filesystem into multiple backup objects.

Client Max Backups

The maximum number of backup threads on the client. This is usually specified for each host. Specifying large value may not decrease the backup window because the backup client may not be able to run multiple threads. For example: Client compression can become the CPU bottleneck for multiple backup threads on the client.

Strategy

You can specify the backup level restrictions for the backup object such as perform only full backups, perform only incremental backup. You can also skip backup of backup object using this field.

Disable Staging

Some backup objects such as NDMP backups do not support Staging Area. They can be disabled on a backup object basis using this field.

Amanda Backup Client Application

Amanda can perform backup of objects using multiple applications. There is a default application for each backup object. You can use different backup application for Linux, Solaris file systems and Oracle.

The override field should be used only when Zmanda Support instructs you to do so. This field is usually used for custom applications.

List of the backup objects in the backup set

The list of backup objects are configured in the backup set can be seen in a table.


View, add, edit, and delete list of objects (DLEs) to backup with: files-s3						
All	Type	Host Name / DLE Check Status ▲	Directory/Path	Encrypt	Compress	Last Modified
<input type="checkbox"/>	Linux/UNIX	✓ aetest-client.dhcp.zmanda.com	/etc	none	none	2010-05-12 11:59
<input checked="" type="checkbox"/>	Linux/UNIX	⚠ amandaclient.company.com	/home	server	client fast	2010-05-15 07:46
<input type="checkbox"/>	Linux/UNIX	✓ localhost	/etc	none	none	2010-05-12 11:59
<input type="button" value="Invert"/> <input type="button" value="Refresh Table"/> <input type="button" value="Edit"/> <input type="button" value="Duplicate"/> <input type="button" value="Check Hosts"/> <input type="button" value="Move"/> <input type="button" value="Copy To"/> <input type="button" value="Delete"/>						Row Legend: Skipped

If you have skipped the backup of an object, it will appear in yellow color as shown above. If you delete a DLE, it will appear in the list of backup objects with ~~strike-through~~ in the list of backup objects.

ZMC Backup Where page

Once the backup source for a backup set has been defined on the **Backup What** page, a backup *device* must be selected on the **Backup Where** page. If no devices have been configured (as will be the case with a new installation), the ZMC will take you directly to the [Admin Devices](#) page to define one or more backup devices. If backup devices are defined, Backup Where allows user to define a backup device for the backup set. You will have to provide backup set specific information for each device depending on the type. Sections below provide more information on how to configure the device for the backup set.

The physical object that stores a backup is referred to in Amanda as a *medium*. Thus, Amanda organizes all backup data in the following heirarchy (note that *slots* only apply to a tape changer device):











Regardless of the device type, the ZMC also allows you to define an optional write-cache mechanism called a staging area or holding disk, which stores the backup image on the server's hard disk. Because backups can be written in parallel to the holding disk, backups can be completed in much smaller windows than would be possible if writing directly to the device (even if the device is a virtual tape). The staging area can be configured in the [Backup Staging](#) page.

Note that while a backup set can include more than one host as a backup source, there can be only one target device per backup set. If you want to send backups from different host systems to different backup media, you must create multiple backup sets to do so. The **Duplicate** feature of the [Admin Backup sets](#) page is useful for this. Also, once a device has been associated to a backup set it cannot be changed (this would invalidate the restore catalog). To start sending backups from an existing backup set to a new device, duplicate the existing set and bind the resulting backup set to the new device. Note that **Backup When** settings will have to be manually set to match the old backup set; these settings are not duplicated.

Backup Where Page Overview

This table shows how the devices are used by different backup sets.

View, add, and edit how backup sets use devices >>									
All	Backup set	Device	Device Name	Changer Path	Device Path	Media Per Backup	Auto Label	Last Modified	Comment
<input type="checkbox"/>	 ch		STK_changer	/dev/changer		1	on	2010-05-15 06:41	
<input type="checkbox"/>	 files-s3		s3_storage	-	akiaj2jpk6mrgcux7yq-files-s3	1	on	2010-05-15 06:41	
<input type="checkbox"/>	 postgres		postgres_vtapes	/dev/null	/var/lib/amanda/vtapes/	1	on	2010-05-15 06:41	
<input type="checkbox"/>	 samba		samba_vtapes	/dev/null	/var/lib/amanda/vtapes/	1	on	2010-05-15 06:41	

Device Name is the name of device as specified in Admin Devices page.

Changer Path and **Device Path** are device file names for tape changer and device. In case of disk backup, **Changer Path** appears as `/dev/null` (i.e., changer entry is not valid). In case of Amazon S3 backup, the S3 access key is shown as **Device Path**.

Media Per Backup defines the number of media volumes that can be used in a backup run. For Amazon S3, tapes and tape changers, the value can be changed in the **Backup Where** page.

Comments are the notes that were specified when the backup device was configured with the backup set.

Advanced Options

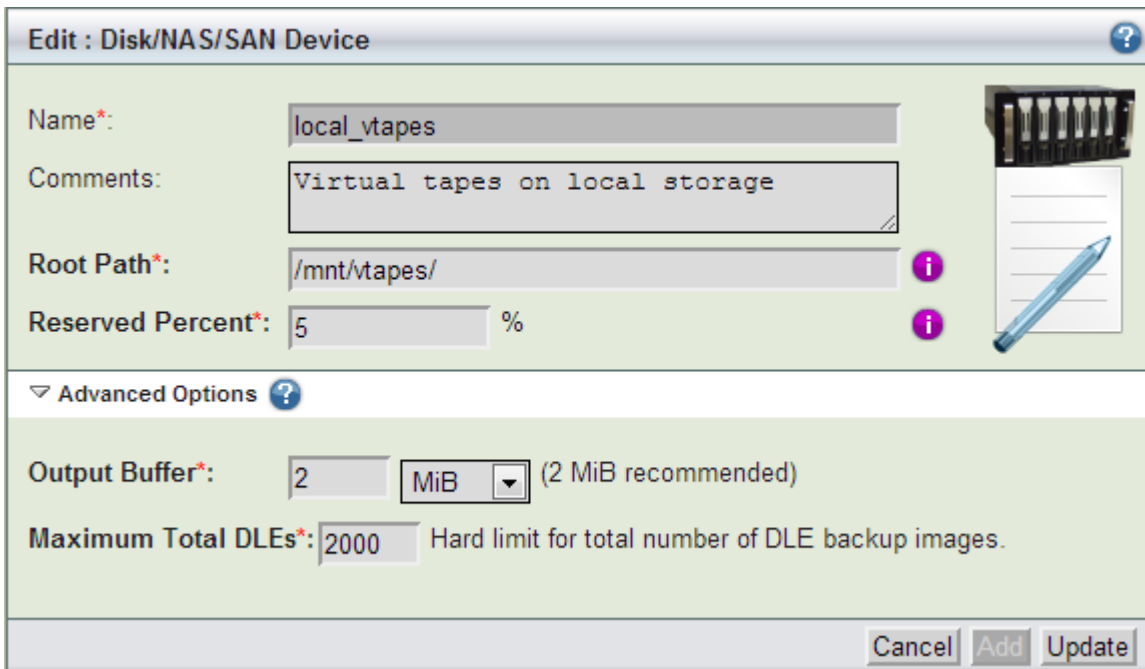
Each Device has advanced options that are required to be modified only under special circumstances or when Zmanda Support team requests you to do so. Please see specific device section about the Advanced options available.

Disk backups

Using hard disk space for a backup provides a number of important benefits:

- The backup window (in other words, the total time to execute backups) can be reduced.
- Files can be conveniently restored directly from the disk, speeding the restore process.

Disk can be locally attached disk or NAS or SAN device. Disk must accessible from the Amanda server.



Comments

Enter a descriptive comment, such as the physical location of the device or any operational notes.

Root Path

The directory where the backup images will be stored, specified when the device was configured in the **Admin Devices** page.

Tip: The amount of free space available to hold backup images is an important consideration for fast retrieval of data. To ensure effective restore capabilities, set aside sufficient disk space to hold more than one full Backup Set worth of Data. Just how many full Backup Sets you should keep on disk depends on the data and your sites requirements for quick restores of accidentally deleted data. The more full backup images stored, the longer the retention of accidentally deleted files. Start with enough space to hold three full Backup Set images and adjust this number as experience dictates.



Tip: Because there is no value in creating a backup of a backup on the same media, Zmanda recommends that the drive that holds the vtapes be excluded from the Backup Set that points to the vtapes.

Reserved Percent

Most file systems reserve certain amount of free space. This value is used for free space calculations to provide ZMC warnings when the backup volume is running out of space.

Advanced Options for Disk should not modified unless Zmanda Support Team recommends you to do so.

Tape Changer

Device Type:



Device Name*:

Comments:

Taperscan:

Auto Label Tapes:
☒

Slot Range*:

Total Tapes Allocated*:

Changer Device*:
MTX

Tape Drive Device*:

Identify Drive Elements to Use

Skip?	Drive Slot	OS Drive Path
<input type="checkbox"/>	0	/dev/nst0
<input type="checkbox"/>	1	/dev/nst1
<input checked="" type="checkbox"/>	skip	<input type="text" value="/dev/"/>
<input checked="" type="checkbox"/>	skip	<input type="text" value="/dev/"/>
<input checked="" type="checkbox"/>	skip	<input type="text" value="/dev/"/>

Advanced Options ?

Ignore Bar Codes:
☐
Not recommended.
Only use if tapes are missing barcode labels.

Comments

Enter a descriptive comment, such as the physical location of the device or any operational notes

Taperscan

The order in which tapes should be searched in the tape changer. Amanda has to load a tape to identify correct tape unless bar code reader is available.

Auto Label Tapes

The 'Auto Label Tapes' radio button allows users to specify whether they want to use the facility of automatic labeling of tapes. The default value is set to No, meaning ZMC expects you to use the [Backup Media](#) page to pre-label the tapes you intended to use for backup. Used tapes cannot be labeled automatically. Amanda will label the tape automatically based on the policy that is specified in the drop-down box. The list of policies are shown below

Auto Label Tapes:	<input checked="" type="checkbox"/>	Only if Amanda label not found
Use Bar Code Reader:	<input checked="" type="checkbox"/>	Only if empty (read returned 0 bytes), or error
Starting Slot Number*:	3	Only if Amanda label not found
		Only if already labelled for a different backup set
		Always

Setting the policy to "Always" is dangerous and is not recommended. If you are using tape changer for multiple backup products including Amanda Enterprise, do not use "Only if Amanda label not found" policy.

Slot range

You can specify the list of slots that can be used for the backup set. Only tapes in these slots will be looked to find tapes for backups. The ranges can be specified as "2-5", "9-13".

Total Tapes Allocated

This value determines the backup retention policy. Amanda will rotate backups among the number of tapes in rotation. You can take tapes out of rotation by archiving it in **Backup Media** page. Please see [Backup When](#) page for more information on this field.

Changer Device

This information is provided for information only and cannot be changed. It can be changed in the **Admin Devices** page.

Tape Drive Device

The tape drive that should be used by the backup set. Please note that Amanda will use **mt** command to manage the device. You can specify multiple tape drives for a backup set. This will allow parallel backups to the tape changer. Users have to match the tape device name to drive slot number. If the match is not done correctly, correct tapes may not be loaded and written to.

Ignore Bar Codes (Advanced Option)

You can disable bar code if the tape do not have bar code labels. This is not recommended because it is difficult to keep track of Amanda media labels and tape. Users will have track this on their own if the bar codes are not enabled.

On Linux servers, the output of **lsscsi** and **mtx** output is provided as a guide to fill out the values as shown below. On Solaris servers, **mtx** command output is provided (**lsscsi** command is not available on Solaris).

```


lsscsi      mt/mtx
[0:0:0:0]   disk    VMware   Virtual disk    1.0    /dev/sda
[0:0:1:0]   disk    VMware   Virtual disk    1.0    /dev/sdb
[0:0:2:0]   disk    VMware   Virtual disk    1.0    /dev/sdc
[0:0:3:0]   disk    VMware   Virtual disk    1.0    /dev/sdd
[0:0:4:0]   disk    VMware   Virtual disk    1.0    /dev/sde
[2:0:0:0]   cd/dvd  NECVMWar VMware IDE CDR10 1.00    /dev/sr0
[6:0:0:0]   mediumx ATL      1500          3.01    /dev/sch0
[6:0:0:1]   tape    HP       Ultrium 1-SCSI  E30W    /dev/st0
[6:0:0:2]   tape    HP       Ultrium 1-SCSI  E30W    /dev/st1

Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: VMware   Model: Virtual disk    Rev: 1.0
  Type:   Direct-Access      ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 01 Lun: 00
  Vendor: VMware   Model: Virtual disk    Rev: 1.0
  Type:   Direct-Access      ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 02 Lun: 00
  Vendor: VMware   Model: Virtual disk    Rev: 1.0
  Type:   Direct-Access      ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 03 Lun: 00
  Vendor: VMware   Model: Virtual disk    Rev: 1.0
  Type:   Direct-Access      ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 04 Lun: 00
  Vendor: VMware   Model: Virtual disk    Rev: 1.0
  Type:   Direct-Access      ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: NECVMWar Model: VMware IDE CDR10 Rev: 1.00
  Type:   CD-ROM            ANSI SCSI revision: 05
Host: scsi6 Channel: 00 Id: 00 Lun: 00
  Vendor: ATL       Model: 1500          Rev: 3.01
  Type:   Medium Changer    ANSI SCSI revision: 03
Host: scsi6 Channel: 00 Id: 00 Lun: 01
  Vendor: HP        Model: Ultrium 1-SCSI  Rev: E30W
  Type:   Sequential-Access ANSI SCSI revision: 03
Host: scsi6 Channel: 00 Id: 00 Lun: 02
  Vendor: HP        Model: Ultrium 1-SCSI  Rev: E30W
  Type:   Sequential-Access ANSI SCSI revision: 03
  
```

Amazon S3

Use ArchiveS3 configuration for device: S3_device

Device Type:



Amazon Simple Storage Service (S3)

Device Name*:

S3_device

Comments:

Taperscan:

Oldest

Location Restriction*:

US-Standard (Virginia or Wash)

Backups stored at*:

zmc-akiajzu3ysv2xo42a2gq-archives3

Advanced Options

Secure Communications:

☒ (recommended, unless running AE on EC2 at Amazon)

Use Full Upload/Download Bandwidth:

☒

Speed Per Thread

Max Upload Speed:

1024

KiB/second x

Max Download Speed:

1024

KiB/second x

Threads Per Upload/Download

Backup Parallelism:

4

Threads

Restore Parallelism:

4

Threads

Cancel

Add

Comments

Enter a descriptive comment, such as operational notes, amazon account owner.

Location Restriction

The default is **US Standard** (Closest Amazon US data center). European customers should use **European Union**(Amazon Ireland data centers) to reduce latency and backup Windows. Other S3 locations can be **US-West**(Northern California data center), **Asia Pacific** (Singapore data center) and other S3 locations. The Storage and Data transfer pricing for each S3 location is different. For details

on the storage and data transfer charges associated with different locations, see the Amazon S3 web site at <http://aws.amazon.com/s3/>.

Backups stored at

The S3 bucket where the backup images will be stored, which was generated when the device was configured in the [Admin Devices](#) page.

Secure Communications (Advanced Options)

Enable this to perform secure SSL transfers of data to and from Amazon S3 cloud. Zmanda recommends enabling secure communication.



Speed Per Thread

Users can throttle the network bandwidth for each upload and download thread. The speed limits apply only for the backup set. You can control the backup and restore from the cloud device.

Threads per Upload/Download

Users can control the number of threads used to upload (backup) and download (restore) for the backup set.

Google Cloud Storage


Device Type:

Google Cloud Storage


Device Name*:

Comments:

Taperscan:

Location Restriction*:

Backups stored at*:


Advanced Options ?

Durable Reduced Availability Storage:
☐

Secure Communications:
☒ (recommended)

Use Full Upload/Download Bandwidth:
☒

Speed Per Thread

Max Upload Speed:
 KiB/second x

Max Download Speed:
 KiB/second x

Threads Per Upload/Download

Backup Parallelism:
 Threads

Restore Parallelism:
 Threads

Comments

Enter a descriptive comment, such as operational notes, Google account information

Taperscan

Algorithm used to identify the next media volume to use. This parameter is not useful for Cloud Storage.

Location Restriction

Backups can be stored in US or EU (Europe) in Google Cloud Sotrage

Backups stored at

The bucket where the backup images will be stored. This value cannot be changed.

Durable Reduced Availability Storage

You can use storage that has less availability. This storage is less expensive. This value cannot be changed after it is set. Changing this value does not migrate data in the cloud.

Secure Communications

Enable or Disable secure backup and recovery to and from the cloud.

Bandwidth throttle

You can use all available bandwidth or throttle the upload (backup) or the download (restore) speed. The throttling is done per thread (see below).

Parallel threads

Each cloud object is uploaded or downloaded in separate thread. You can specify how many threads you want to use. Using too many threads might slow the system down.









ZMC Backup Staging

Regardless of the device type, the ZMC also allows you to define an optional write-cache mechanism called a staging area or holding disk, which stores the backup image on the server's hard disk. Because backups can be written in parallel to the holding disk, backups can be completed in much smaller windows than would be possible if writing directly to the device (even if the device is a virtual tape).

In case of device/media failures (such as an unavailable S3 connection or running out of tapes in the changer), backups are stored in the staging area. The backup images stored in the staging area can be used for recovery and will be moved to the secondary media in next backup run if the media problem is resolved.

The staging area is used for Amanda DLE only if the whole backup image will fit in the staging area. Otherwise, DLE is directly written to the backup media volume without the use of staging area. Zmanda recommends the staging area size should be at least the size of one full backup image of all DLEs in the backup set.

View Staging Area

(8) View and edit how backup sets use staging areas								
(8/8) Rows: 10								
All	Backup Set	Device	Flush	Directory	Staging Partition Statistics	Staging Usage Statistics	Time Last Modified	By
<input type="checkbox"/>	br001_vtape		yes	/var/lib/amanda/staging/br001_vtape	Disabled	Disabled	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br002_cloudena_cld		yes	/var/lib/amanda/staging/br002_cloudena_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br003_as3_cld		yes	/var/lib/amanda/staging/br003_as3_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br004_goo_cld		yes	/var/lib/amanda/staging/br004_goo_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br005_hp_cld		yes	/var/lib/amanda/staging/br005_hp_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br006_lij_cld		yes	/var/lib/amanda/staging/br006_lij_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br007_essex_cld		yes	/var/lib/amanda/staging/br007_essex_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin
<input type="checkbox"/>	br008_aplira_cld		yes	/var/lib/amanda/staging/br008_aplira_cld	16.0G 21.7G	0.0G 5.4G	2012-07-13 16:06:45	admin

The above table shows staging areas used by each backup set. The **Device** column whether it is a Disk or Tape or Amazon S3 backup. **Flush** column indicates if the previous backup images in the staging area are moved to backup media before the current backup run automatically. The location of staging area, disk space used, amount of data used in staging area are also displayed.

Staging Area Configuration

(8) Edit Staging Configuration for: br007_essex_cld

Device Name*:

Auto Flush: ☒ automatically flush dumps from staging to media

Staging Size Limit: ⓘ

Recommended: Staging size is larger than the largest DLE.

Reserved for root: % (5% recommended; 10% for Solaris)

Backup runs staged at*: ⓘ


Live Statistics

Partition total space*:

Partition free space*:

Partition used space*:

Note: Space for each backup is pre-allocated before starting the backup, unless staging has been disabled.



Advanced Options ⓘ

Reserve for Incrementals: % (0% recommended)

Cancel Update

Device Name

The name of the ZMC device used by the backup set. This information cannot be modified in this page.

Backups runs staged at

The absolute path to the file staging area where backup images are collected before being archived. The default is `/var/lib/amanda/staging/<backup set name>`. The staging area directory must be accessible only by the `amandabackup` user. The staging area is used to provide better performance and better media fault tolerance. All backups from the client can be copied to the staging area in parallel. The amount of parallelism depends on the Clients in parallel and other options on the Backup How page.

Auto Flush

Specifies whether Amanda should flush the backup images from the staging/area holding disk to the backup media before a backup run. The default value is enabled. You can change the default globally for all backup sets by changing the [amanda.conf\(5\)](#) file's `autoflush no` parameter to `autoflush yes`. You can monitor flush activity from the ZMC Monitor page.

Enabling Auto Flush may require additional space on the media volumes to account for previous backup runs, especially if you have set the Backup When Tapes/Virtual Tapes Per Run value conservatively. If the Backup When Tapes/Virtual Tapes Per Run value is greater than 1, then the additional media will be used if it is required; if not the staging area must hold any backup images that exceed the limit. For example, consider a setup where a full backup usually fills up 3/4 of the medium. If Tapes/Virtual Tapes Per Run is set to 1, you will usually end up with a separate pieces of media for every backup, each filled to 3/4 capacity. If there happens to be a day where there are two backup images in a run, the staging area/holding disk must store the backup until there is a backup run with only one image that is less than one quarter of the media size. So if staging area space is at a premium and/or you wish to ensure that backup images always get saved to secondary media rather than the holding disk, you should set Tapes/Virtual Tapes Per Run accordingly.

Staging size

The controls allow you to allocate all space, all space except a given amount, or no more than a given amount of space. Be sure to allocate enough space to hold at least a full backup.

Partition total space/Partition free space

The total space and free space in the partition that contains the staging areas. This information is only for information.

Reserve for Incrementals


When space on a staging area falls below a threshold size (the value of this parameter), Amanda limits itself to performing only incremental backups. A threshold of 20% causes Amanda fall back to incremental backups when holding disk(s) free space falls below 20%.

When there is a media error during a backup run, the backup images are still stored in the staging area. Amanda allows recovery from holding disk and backup media.

Important Note:

Do not attempt to back up the staging area. Use Exclude options if necessary to prevent any backup set from including a staging area for itself or any other backup set.

Staging Area Contents



This pane shows the space used for staging area for this backup set and list of files in the staging area.

ZMC Backup Media Page

Amanda use labels to identify backup media. Labels ensure that the correct backup media is loaded for the backup set, and that no backup is overwritten before the planned cycle is complete. When ZMC finds a labeled piece of media (or media that appears to already contain backup data), it will not write a backup to that tape. When the ZMC finds unlabeled media, it:









1. Writes a label for subsequent identification
2. A record of the new media is catalogued. This allows the ZMC to use new media before used media is recycled. Media is used in the order that it is labeled.

When the ZMC finds media that belongs to some other backup set (or media that has been written to by another application), it is not used, and an error message is returned. You can manually "force" labeling of such media by setting the **Overwrite Media Option** described below to **Yes**.

ZMC will not overwrite an existing backup tape until it has first written a number of fresh tapes equal to the number of tapes in the planned **Backup When** page.

ZMC Backup Media page allows users to view backup sets media information, manage media and label media. They are described in the following sections.

View Media

View Media Used by Backup Sets >>											
(page 1 of 2) Next >> Go To Page: 1 Go ?											
All	Backup set	Device	Schedule Type	Archived	Backup Cycle	Retention	Hours	Minute	When	Last Modified	Status
<input type="checkbox"/>	 glob		Every Weekday	0	7 days tapecycle:6	D7 days H4 days E7 days	0 (full:0)	0 (15)	S M T W R F S - A A A A A A	2010-05-18 14:52	
<input type="checkbox"/>	 mac_dle		Every Weekday	0	7 days tapecycle:6	D7 days H7 days E7 days	0 (full:0)	0 (15)	S M T W R F S - A A A A A A	2010-05-18 14:52	
<input type="checkbox"/>	 macmini		Every Weekday	0	7 days tapecycle:6	D7 days H4 days E7 days	0 (full:0)	0 (15)	S M T W R F S - A A A A A A	2010-05-18 14:52	
<input type="checkbox"/>	 samba		Every Weekday	0	7 days tapecycle:6	D7 days H11 days	0 (full:0)	0 (15)	S M T W R F S - A A A A A A	2010-05-18 14:52	

This table shows the list of backup sets, the backup schedule, the list of media volumes archived. This table can be used to select a backup set for media management.

Manage Media

The Manage Media table (as shown below) shows the list of media volumes in the backup set and volume labels. It shows when the media was used for backups. You can sort the table using **Last Used** column. The **Archived?** column shows the media volume has been archived. Archived volumes are retained forever and are not used for backups (i.e, never overwritten).

L0 in the Media Labels column shows that the media contains level 0 backup of a backup object (DLE). You can find more information about the data in the media volume in ZMC Report pages.

Manage Media for: samba >>			
Row Legend: <ul style="list-style-type: none"> ■ Part of current retention period. ■ Probably used next. ■ Desired retention period violation (unless new tapes provided). ■ L0 Media contains at least one DLE with a level 0 backup image. 			
All	Last Used	Media Labels / Barcodes	Archived?
		2 New Tape(s) New tape(s) needed.	No
<input type="checkbox"/>	2010-05-07 10:09	L0 samba-001	No
<input type="checkbox"/>	2010-05-07 11:59	L0 samba-002	No
<input type="checkbox"/>	2010-05-07 12:16	L0 samba-003	No
<input type="checkbox"/>	2010-05-07 12:24	L0 samba-004	No
<input type="button" value="Invert"/> <input type="button" value="Recycle"/> <input type="button" value="Drop"/> <input type="button" value="Archive"/>			

Yellow color on media volume row(s) indicate that this volume(s) is likely to be used in next backup run. This information can be used to put the tapes in the tape changer before the next scheduled backup run. The green color in a row indicates that this volume will be overwritten as part of the media rotation by Amanda. The red color indicates that tape is likely to be overwritten unless new tapes are available. This information is useful in tape changer configurations to make sure fresh tapes are available.

The Manage Media table below shows a tape changer configuration that has multiple labelled empty tapes available for future backup runs (yellow color). The table also shows the tapes labels that include the tape bar codes.

Manage Media for: gchanger >> ?			
Row Legend: <ul style="list-style-type: none"> ▪ Part of current retention period. ▪ Probably used next. ▪ Desired retention period violation (unless new tapes provided). ▪ L0 Media contains at least one DLE with a level 0 backup image. 			
All	▲Last Used	Media Labels / Barcodes	Archived?
<input type="checkbox"/>	2010-04-12 14:42	L0 gchanger-001	No
<input type="checkbox"/>	2010-04-15 11:11	L0 gchanger-002	No
<input type="checkbox"/>	2010-04-16 19:47	gchanger-021-007007L3	No
<input checked="" type="checkbox"/>	2010-04-22 18:00	gchanger-019-007006L3	No
<input type="checkbox"/>	empty	gchanger-020-007005L3 gchanger-012-007007L3 gchanger-011-007006L3 gchanger-010-007005L3 gchanger-009-007007L3 gchanger-008-007006L3 gchanger-007-007005L3	No
<input type="button" value="Invert"/> <input type="button" value="Recycle"/> <input type="button" value="Drop"/> <input type="button" value="Archive"/>			

Invert

Select all entries except the ones that are selected using check boxes.

Recycle

Recycle identifies the media as immediately available for re-use (in other words, existing data on the media will be overwritten and lost). Recycle operations cannot be performed while a backup is in progress.

Drop

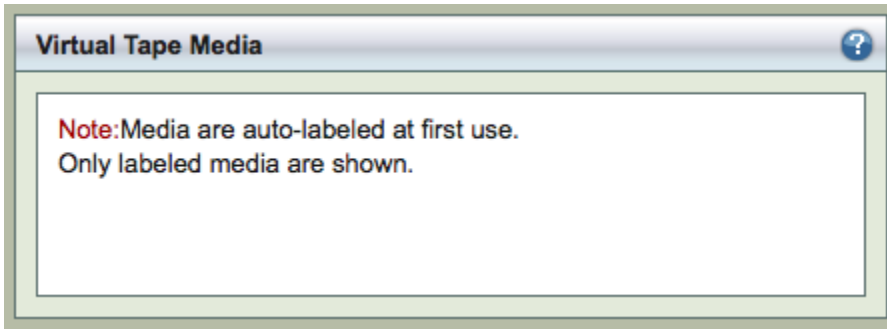
Removes the media and all backups it contains from the Amanda server. Do not drop media that you wish to restore from in the future unless you first click the **Archive** button for that media, and then manually copy the data to a different location. Drop operations cannot be performed while a backup is in progress.

Archive

Amanda will take the media volume out of rotation. The media volume will not be overwritten. You should use the preserve the backup data for longer duration. This media volume is still tracked by Amanda catalog and is available for backup restoration.

Media Labelling

Amanda media must be labelled before it is used for backups. Disk backup media volumes are labelled automatically. If the backup sets uses disk backups, following message will appear. Auto Labelling can be enabled in **Backup Where** page.



It should be noted auto labelling will work only with new tapes and auto labelling will not work on tapes have any data (not just other Amanda labels).

You can label media volumes (tapes) manually. In a backup set that uses a tape changer, use **Scan All Slots** to read the labels from the tapes (if any) that are on the slots reserved for the backup set (See Starting/Last Slot number in **Backup Where** page). Scan all slots before starting manual labelling process.

The list of tapes with barcode is displayed with slot numbers. It also shows the **Last Used** date when a backup run used the tape. If there is no tape label, the table shows **unknown**. If there is no tape in the slot, **empty slot** is displayed as shown for slots 6 to 10 in the table below.

Label Media for: gchanger >>

Advanced Options

Include Barcode: ☒
Scan All Slots

Overwrite Media: ☐

All	Slot	Barcode	Last Used	Current/New Label
<input type="checkbox"/>	1	00A50000		unknown <input type="text"/> <input checked="" type="checkbox"/>
<input type="checkbox"/>	2	00A50001		unknown <input type="text"/>
<input type="checkbox"/>	3	00A50002		unknown <input type="text"/>
<input type="checkbox"/>	4	00A50003		unknown <input type="text"/>
<input type="checkbox"/>	5	00A50004		unknown <input type="text"/>
	6	empty		empty slot
	7	empty		empty slot
	8	empty		empty slot
	9	empty		empty slot
	10	empty		empty slot

Invert
Save Labels
Load New Media

In **Advanced Options**, you can include tape bar code in the tape labels. Please note that bar code reader must be enabled in **Backup Where** page. See the next section on how to configure Tape changer bar code information.

Use **Overwrite Media** option to label tapes that already have data in them. Use caution before overwriting label. The old data cannot be recovered.

To label a tape, enter the label name in the **Current/New Label** field as shown below and click **Save Labels** button. The label name must be in `<number(s)>-<alphanumeric>` format

Label Media for: gchanger >> ?

Advanced Options ?

All	Slot	Barcode	Last Used	Current/New Label
<input type="checkbox"/>	1	00A50000		unknown ✓
<input type="checkbox"/>	2	00A50001		unknown
<input type="checkbox"/>	3	00A50002		unknown
<input type="checkbox"/>	4	00A50003		unknown
<input checked="" type="checkbox"/>	5	00A50004		newlabel
	6	empty		empty slot
	7	empty		empty slot
	8	empty		empty slot
	9	empty		empty slot
	10	empty		empty slot

Invert Save Labels Load New Media

The label name will be <backup set name>-<text entered>-<bar code>.

Click **Load New Media** to display a dialog that prompts you to load new media in the selected slots. When you load media and click **Done**, the ZMC reads the tape labels and refreshes the display.

<input type="checkbox"/>	5	00A50004		unknown ✗
--------------------------	---	----------	--	-----------

If the labelling process fails, there is a red cross next to the entry and message appears in the message box as shown below.

✗ Label of slot 5 failed: Label 'gchanger-newlabel-00A50004' doesn't match labelstr '^gchanger-[0-9][0-9]*(-.*)?\$'.

Tape Changer Bar Codes

Amanda Enterprise edition supports tape changers with bar code readers. To use bar codes in a backup set, it is necessary to initialize the bar code database. Initialization steps must be performed using command line tools. Initialization is required only if tapes were not labeled with bar code enabled or not managed using Zmanda Management Console.

1. Remove the existing bar code database file (if any) on the Amanda server.

```
$ /etc/amanda/<backup set name>/changer-barcodes
```

2. Create bar code database file for the backup set on the Amanda server. This process can take a lot of time. The process involves loading of all tapes in the tape changer and bar codes are read.

```
$ /usr/sbin/amtape <backup set name> update
```

This operation can be performed only after saving a Backup Where configuration for the backup set, and must be performed as the **amandabackup** user on the Amanda server.

ZMC Backup When Page

Unlike most other backup software, ZMC does not require users to lay down detailed rules as to when the backup should be performed. ZMC works through Amanda's backup scheduler which tries to maintain a consistent backup window across all backup runs by controlling the size of the backup image for each backup run.









Amanda also simplifies the configuration of the backup run schedule and specification of backup level. Instead of users pre-specifying either a Full Backup or an Incremental Backup in a particular schedule, ZMC has Full backup and 9 levels of Incremental Backup. However, ZMC does not require Administrators to specify the backup levels for each backup run. It calculates, using an internal algorithm, the levels of each DLE within the backup set for every run on a dynamic basis.

The **Tapes In Rotation/Virtual Tapes in rotation** option lets you define a backup retention policy how much media is available for rotation (see below) for a single backup cycle, and therefore indirectly set a retention policy. By putting more media in rotation, you are effectively setting a longer retention policy because if more media is available, completed backups will not get overwritten as quickly by subsequent backup runs.

Backup Set Schedule table

The table shows the list of backup sets and backup schedules. Overall view on the backup schedules help understand how the backup windows are overlapping and how many backups are running at any time.

(8) View and edit how backup sets use staging areas (8/8) Rows: 10

All	Backup Set	Device	Schedule Type	Archived vs. Used	Backup Cycle	Retention%	Hours	Minute	When	Time Last Modified	By	Status
<input type="checkbox"/>	br001_vtape		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br002_cloudena_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br003_as3_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br004_goo_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br005_hp_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br006_lij_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br007_essex_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	
<input type="checkbox"/>	br008_apтира_cld		Every Weekday	0/0	7	.07 days No backups found. E49 days	0	15	SIMTWTWTFSS - A A A A A -	2012-07-13 16:06:45	admin	

Invert Selection Refresh Table Edit


The backup job for the backup set starts at the time **Hours:Minute** column and it is run based on schedule in **When** column. The **When** column shows two rows: The first row has the days of the week (Sun to Saturday). The second row shows the type of backup schedule is performed on each day.

The backup schedule can be **Amanda** (A), **Full** (F) and **Incremental** (I). Amanda backup schedule means Amanda decides backup level to use. Amanda tries to balance the amount of data backed up on each run and is the recommended schedule.

Changing Backup Set Schedule

(8) Edit Backup Schedule br001_vtp for: br001_vtape

Backup Start Time:


 AM 12 ☒ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10 ☐ 11 ☐
 PM 12 ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10 ☐ 11 ☐

Minutes after the hours selected above

Schedule Type:

Cancel Update

Backup Start Time

Specifies the time of day when the backup run starts. This time is the same for all days and for all the DLEs in the set.

Backup Cycle

The number of weeks within which at least one full backup for all DLEs must be performed.

Schedule Type

The Schedule Type can be Every Day, Every Weekday, Custom Weekday, Incremental Weekdays-Full Sunday, and Incremental Weekdays-Full Saturday. Amanda will schedule backup levels in Every Day and Every Weekday schedule on every day of the week and Monday-Friday respectively. You can see force specific backup levels using Increment Weekdays-Full Sunday and Incremental Weekdays-Full Saturday schedule.

Custom Weekday schedule is an advanced option. For each day of the week, you select from No Backup, Full (backup level 0), Incremental (backup level > 0) and Auto (Amanda decides backup level) as shown below.

Custom Weekday (Advanced) Backup:*

Type	Mon	Tue	Wed	Thur	Fri	Sat	Sun
No Backup	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>
Auto	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Full	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Incremental	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Custom Days of the Month (shown below) is another advanced option. For each day of the month, you can select Amanda schedule (auto), full backup or incremental (level > 0) or no backup. For days of the month that are not valid for the specific month (such 31st in Feb), the value is ignored.

Custom Days of the Month (Advanced) Backup:*

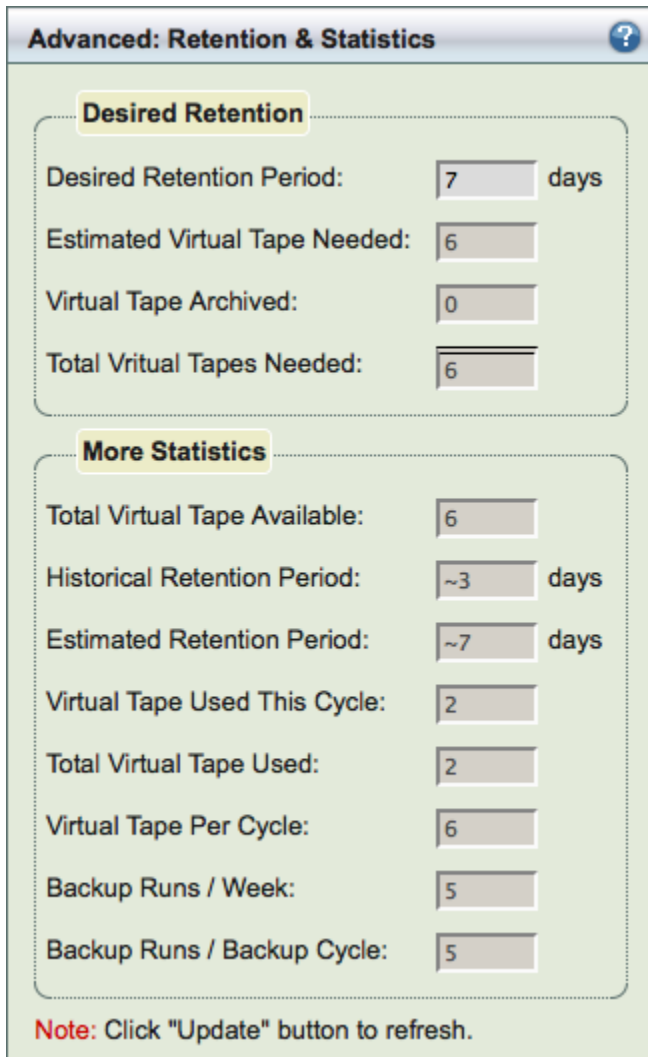
Day	Type	Day	Type	Day	Type
1	-	11	-	21	-
2	-	12	-	22	-
3	auto	13	-	23	-
4	-	14	-	24	-
5	-	15	-	25	-
6	-	16	-	26	-
7	-	17	-	27	-
8	-	18	-	28	-
9	-	19	-	29	-
10	-	20	-	30	-
				31	-

Virtual Tape/Backup or Tape/Backup

Specifies the maximum amount of media that Amanda will be allowed to consume in a single backup run. This value, coupled with **Virtual Tapes/Tapes in Rotation** above, effectively sets the data retention policy for the backups.

Scheduling Advanced Options

Amanda scheduling and retention policies are inter-twined with the number of media volumes required (i.e, number of tapes needed for the backup set or amount of disk space required for backups). The following picture shows the values for disk backups.



Advanced: Retention & Statistics

Desired Retention

Desired Retention Period:	<input type="text" value="7"/>	days
Estimated Virtual Tape Needed:	<input type="text" value="6"/>	
Virtual Tape Archived:	<input type="text" value="0"/>	
Total Virtual Tapes Needed:	<input type="text" value="6"/>	

More Statistics

Total Virtual Tape Available:	<input type="text" value="6"/>	
Historical Retention Period:	<input type="text" value="~3"/>	days
Estimated Retention Period:	<input type="text" value="~7"/>	days
Virtual Tape Used This Cycle:	<input type="text" value="2"/>	
Total Virtual Tape Used:	<input type="text" value="2"/>	
Virtual Tape Per Cycle:	<input type="text" value="6"/>	
Backup Runs / Week:	<input type="text" value="5"/>	
Backup Runs / Backup Cycle:	<input type="text" value="5"/>	

Note: Click "Update" button to refresh.

You can enter the **Desired Retention Period** in days and click **Update** to obtain various values. It provides an estimate of number of virtual tapes. It also provides the number of virtual tapes archived so far and total number of virtual tapes for the backup set on the Amanda server.

It also provide comparison with current number of Virtual tapes and current backup schedule.

The information provided in this pane is advisory and is not used by Amanda. It can be used to fill in values for Tape/Virtual Tape in rotation, backup schedule and backup cycle in the **Backup When** page.

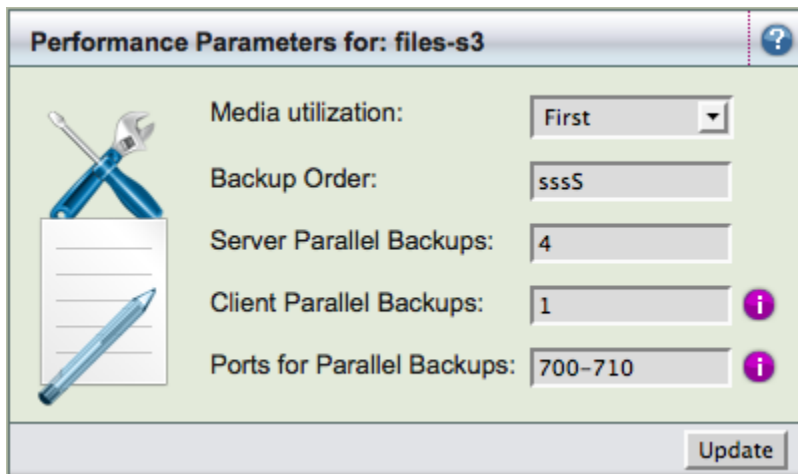
ZMC Backup How Page

After you have defined what, where and when for the backup set, you can use the **Backup How** page to define some of the key internal parameters that control how the backup set will run after it has been activated. In most situations, the default values set for these parameters will work well for most backup sets and therefore need not be changed.

Advanced users intending to modify the default values should study the logs and reports of a few past backup runs in light of the discussion that follows.

The parameters should be specified for each backup set. Select the backup set using the table at the bottom of the page by clicking the table row.

Performance Parameters



Performance Parameters for: files-s3

Media utilization: First

Backup Order: sssS

Server Parallel Backups: 4

Client Parallel Backups: 1

Ports for Parallel Backups: 700-710

Update

Media Utilization

Staging area holding disks are written in parallel. By the time the first completed DLE backup has been written to the media, many more DLEs backups may also be ready to be moved. This parameter sets the algorithm that determines the order in which the completed backup images are moved from holding disk to the backup media.

Only completed backup images on the holding disk are considered for moving; the ordering of the images can help in better utilization of backup media.

Options are:

- First: (Default value) First in - first out.
- Firstfit: The first backup image that will fit on the current media volume.
- Largest: The largest backup image first.
- Largestfit: The largest backup image that will fit on the current media volume.
- Smallest: The smallest backup image first.

- Last: Last in - first out.

Backup Order

Whenever two or more parallel streams need simultaneous access to the same media volume, there is a potential conflict: Which stream should be allowed to write first?

Using Backup Order, you can assign priority to each of the parallel backup processes, the number of which is determined by the value of the "Parallel Backups" field above.

The possible option choices and what they mean is indicated below:

- **s** -> smallest size first
- **S** -> biggest size first
- **t** -> smallest time first
- **T** -> biggest time first
- **b** -> smallest bandwidth first
- **B** -> biggest bandwidth first
- Try **BTBTBTBTBTBT** if holding disks have not been assigned a threshold (see below)

Example 1 :

A string like "sssS" which represents the priority order for four parallel backup processes. Three dumpers will seek the smallest size hosts while one dumper will seek the biggest size host to backup.

Example 2 :

A string like 'BbsTt' indicates there are five parallel backup processes of which:
One backup process(B) is looking for the Biggest Bandwidth occupying backups.
Another backup process(b) is looking for the smallest Bandwidth occupying backups.
Another backup process(s) is looking for the smallest size backups.
Another backup process(T) is looking for biggest Time occupying backups.
Another backup process(t) is looking for smallest Time occupying backups.

Server Parallel Backups

Sets the number of parallel data backups that are performed from the Amanda clients to the holding disk in a backup run. The default value is 10. How many parallel backups happen to the staging area/holding disk depends on

- CPU and Write speed on the Amanda server
- Network configuration between the Amanda server and Amanda client

Parallel Backup Clients

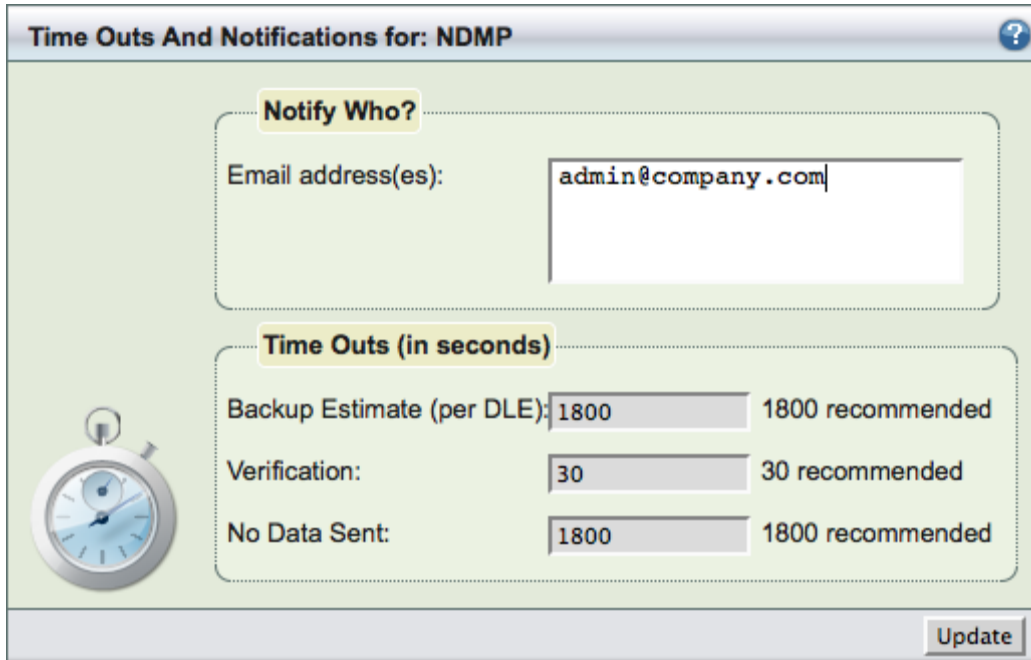
Sets the number of parallel backups performed from an Amanda client. The default value is 1, which specifies that all DLEs on a client are backed up sequentially.

If an Amanda client has sufficient CPU and network resources, multiple DLEs on the client could be backed up in parallel.

Ports for Parallel Backups

The port range to use on Amanda server to connect to clients. Amanda will use all the ports in the range that are not bound by another process or not reserved in the `/etc/services` file. You will have increase the range (or use different range) if you are increasing the value of **Server Parallel Backups**.

Time Outs and Notifications



Time Outs And Notifications for: NDMP

Notify Who?

Email address(es):

Time Outs (in seconds)

Backup Estimate (per DLE): 1800 recommended

Verification: 30 recommended

No Data Sent: 1800 recommended

Email address

Sets the e-mail address of the Amanda backup administrator. Though this is an optional parameter, it is important to set the Notification parameter to the email address(es) of the Amanda administrator. Amanda backup run notifications will be sent to the email address. The email address can be different from the email address set in the Admin page for a particular user. ZMC also provides content of the notification in the Report Summary page for the backup set. You can enter multiple email addresses separated by comma or white space characters.

The Mail User Agent (MUA) such as *Mail* should be configured on the Amanda server before any email can be sent. ZMC does not install or configure the MUA. Amanda uses `/usr/bin/mail` by default to send email notifications. To configure alternate mail program, please see [Configuring Alternate Mail Program](#) section.

Backup Estimate (per DLE/Backup Object)

Limits the amount of time the ZMC will spend on the first phase (Planning phase) of the backup process. The Amanda backup process first estimates backup size for each DLE configured in the backup set. The backup size estimation for all DLEs in the backup set is done in parallel.

Specify a time out (in Seconds) by which each Amanda client has to respond with a backup size estimate.

If the time out occurs for a particular DLE, that DLE will not be backed up during the backup run. Amanda will perform a full backup for the DLE in the next backup run.

The time out specified should take into account factors such as clients with limited network bandwidth, limited CPU, multiple DLEs on one host, and clients with large amounts of data to be backed up.

Some idea of the time being taken to perform size estimates can be had from the **Monitor** page where details of the current backup run are displayed. A much better idea can be had from the **Report Timeline** page where DLE-wise details are displayed.

If a DLE has timed out, it will be displayed in red. Other DLEs are displayed in green.

Verification

ZMC **Backup What** verifies the configuration of the client/Backup object that is being added or modified. The client verification is controlled by this time out.

The time out may be adjusted for clients with limited network bandwidth.

If the time out expires, the verification of the client fails.

This time out has no impact on the backup process. If the client is properly set up, it will be backed up.

The time out applies to all clients in the backup set.

No Data Sent

Limits the time (in Seconds) that the Amanda server will wait (during the backup data transfer phase) for a client to begin to respond to a backup request.

If data time out occurs for a client, the backup of all directories for the client is not done in the backup run. The directories on the client will be backed up in the next backup run.

The time out may be adjusted for clients with limited network bandwidth.

The time out applies to all clients in the backup set.

Configuring an Alternative Mail Program

By default, Amanda assumes the local mail program found during initial configuration will be used to send notifications. You may wish to override this. For example, if your site uses an SMTP server, you can configure Amanda to use that instead of the program found during initial configuration. Use the **mailer** directive in the **amanda.conf** file:

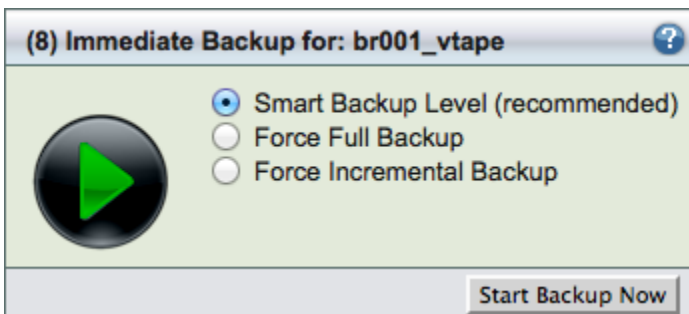
```
mailer mail_program_path
```

The mail program must be able to process messages using the syntax:

```
MAILER -s "subject" user < message_file
```






ZMC Backup Activation Page

A backup set must be activated for its automatically-scheduled backups to execute. backup sets must be activated individually. The **Backup Activate** page also includes an option for immediate execution of the backup set.



Activate Backup lets you select backup sets to put in production (in other words, begin executing the backups as scheduled). You can also **Deactivate** an already activated backup set. When a backup set is deactivated, ZMC does not execute backup runs associated with that backup set. You can also trigger an **Immediate Backup** to be executed at any time. When an Immediate Backup is run, you can observe the progress of the backup run in the **Monitor** page.

Before a backup set is activated, make sure the backup schedules are configured in the **Backup When** page, the backup device is configured in the **Backup Where** page and clients/applications are configured in the **Backup What** page for the backup set.

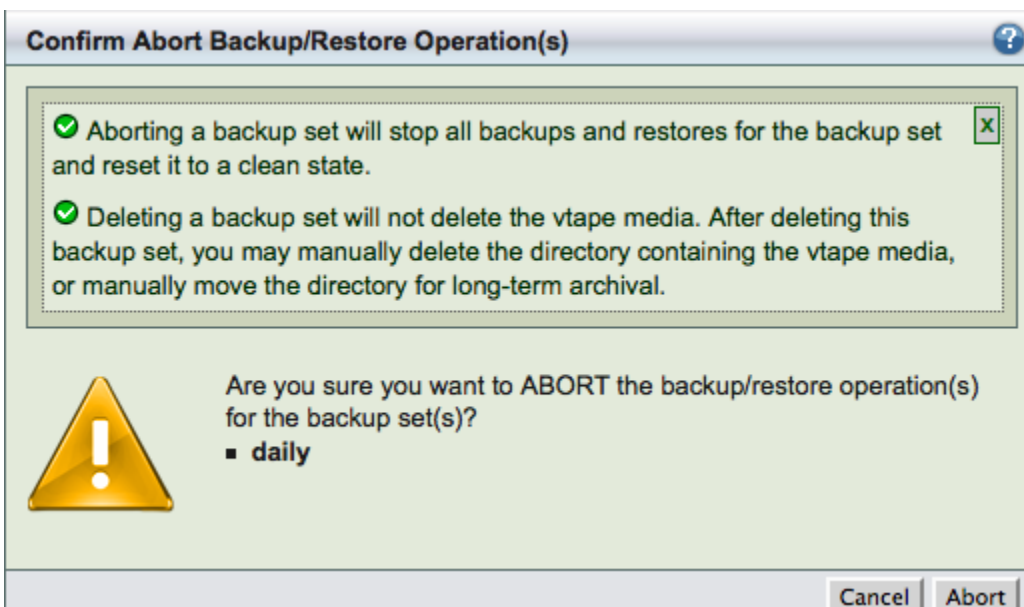
Select a backup set >>										
All	Active	Backup Set Name▲	Schedule	Code	✓	✗	Version	Status	Creation Date	Comments
<input type="checkbox"/>		ch	Every Weekday	-	1	0	3.0	OK (device: STK_changer)	2010-05-13	backup set that uses a VTL changer
<input type="checkbox"/>		files-s3	Every Weekday	-	3	1	3.0	OK (device: s3_storage)	2010-05-12	
<input type="checkbox"/>		new-backup-set		-	0	0	3.0	OK (device: NONE)	2010-05-15	
<input type="checkbox"/>		postgres	Every Weekday	-	1	0	3.0	OK (device: postgres_vtapes)	2010-05-12	
<input type="checkbox"/>		samba	Every Weekday	-	1	0	3.0	OK (device: samba_vtapes)	2010-05-12	

You view status of each backup set. The **Active** column in the above shows state of activation. A green "power" icon indicates the backup set has a schedule configured and installed using the "cron" program. Red icons indicate the backup set has a schedule, but the schedule has not been activated. A grey icon indicates the backup set can not be activated. Backup sets lacking device bindings or having other problems can not be activated.

Activating a backup set causes ZMC to create a **crontab** entry for the **amandabackup** user corresponding to the backup set schedule specified in the **Backup When** page. Some complex ZMC schedules will require multiple crontab entries. ZMC will create and manage these crontab entries automatically, when the backup set is activated using the "Activate" button.

All crontab entries added by ZMC have the option argument `--zmcdev`. These ZMC-created crontab entries are automatically regenerated by ZMC. To manually edit and alter these crontab entries, first remove the argument `--zmcdev`. Since ZMC ignores crontab entries lacking `--zmcdev`, ZMC will not replace the manually edited entry, and ZMC will show the backup set as inactive.

You can **Activate/Deactivate/Edit** backup sets after selecting multiple backup sets. You can also abort the backup run for a backup set using the **Abort** button. Aborting a backup stops the backup job and cleans up all the backup processes on the server.



Above image shows the confirmation window for the abort operation.

Monitoring Backup runs





















Once backup sets have been configured and activated, the only operator intervention remaining is to manually load media as required. The **Monitor** page lets you track the progress of backup sets and fix any problems should they arise.

If no backup is running for the currently selected backup set, the **Monitor Backups** page will be empty.

If no details are shown for a backup you think is in progress, please check the **Change Backup Set** drop down box at the top of the page to ensure that the correct backup set is selected to monitor.

The **Monitor** page is not updated automatically. To see updated details, please click **Refresh** button.

Use **Toggle Details** button get information about a backup phase as shown below.

(10) Timeline Monitor Chart							(10/48) (page 1 of 5) Next >> Page: <div><div></div><div></div></div> Rows: <div><div></div><div></div></div>				
▼Date and Time	Backup Set	Type	Level	Host Name	Directory/Path	State	Clearing Staging Area	Checking Backup Plan	Transferring Backup to Server	Writing to Backup Media	
2012-07-13 11:00	 br007_essex_cld	Linux*nix	<div><div></div></div>	10.5.52.2	/opt/zmanda/amanda/ZMC/testtools/test_backup		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>4m finished (11:01:15)</div>	
2012-07-13 11:00	 br007_essex_cld	NFS	<div><div></div></div>	127.0.0.1	/mnt/opt		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>9m finished (11:01:36)</div>	
2012-07-13 11:00	 br007_essex_cld	Share	<div><div></div></div>	127.0.0.1	\\10.5.52.2\opt		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>0m finished (11:01:41)</div>	
2012-07-13 11:00	 br007_essex_cld	Linux*nix	<div><div></div></div>	localhost	/opt/zmanda/amanda/ZMC/testtools/test_backup		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>4m finished (11:00:56)</div>	
2012-07-13 11:00	 br007_essex_cld	Windows	<div><div></div></div>	zwc-qa.dhcp.zmanda.com	C:/zmc-test		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>0m finished (11:01:22)</div>	
2012-07-13 11:00	 br007_essex_cld	Windows Template	<div><div></div></div>	zwc-qa.dhcp.zmanda.com	ZMC-Test		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>0m finished (11:01:03)</div>	
2012-07-13 10:59	 br006_iii_cld	Linux*nix	<div><div></div></div>	10.5.52.2	/opt/zmanda/amanda/ZMC/testtools/test_backup		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>4m finished (11:00:15)</div>	
2012-07-13 10:59	 br006_iii_cld	NFS	<div><div></div></div>	127.0.0.1	/mnt/opt		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>9m finished (11:00:39)</div>	
2012-07-13 10:59	 br006_iii_cld	Share	<div><div></div></div>	127.0.0.1	\\10.5.52.2\opt		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>0m finished (11:00:50)</div>	
2012-07-13 10:59	 br006_iii_cld	Linux*nix	<div><div></div></div>	localhost	/opt/zmanda/amanda/ZMC/testtools/test_backup		<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	<div><div></div></div> <div>4m finished (10:59:56)</div>	
<div>Refresh Table</div>											

The **Monitor Chart** as shown above displays the progress of backups that have already completed as well as those that are in progress. The chart is labeled with the time at which the selected backup set run was started. Each Disk List Entry is displayed as a separate row. Each row displays the following information:

Host Name

The client Host name as entered in the Disk List.

Directory

The starting Directory from which the backup begins.

Level

ZMC schedules its own backup levels and automatically selects them.

- Level 0, Full Backup, is indicated by a solid green bar as in second row above.
- Level 1 and above, incremental backup, is indicated by a partially-filled-in green bar.

Clearing Staging Area

Clearing the staging area flushes the data from the disk staging area to the secondary media. ZMC first checks that the data of last backup run has been flushed or not. If that has been done, a solid green bar will appear under this section. If flushing is in progress, a striped green bar will appear under this section.

Checking Backup Plan

ZMC requests each client to furnish an estimate of the size of backup. This process takes some time. While in progress, striped green bars are displayed. When it is complete, a solid green bar is displayed.

Transferring Backup to Server

When Backup image is to be written to a staging area, the Timeline will show striped or solid green bars depending upon the status.

Writing to Backup Media

Whether the Data has been written to the Backup Media with success or not.

If the activity has not yet been started, the position is blank. While it is being written, there is a flashing striped green bar. If it has been written, there will be a solid green bar against the DLE.

Report Timeline Monitoring

To see reports on previously-completed backup runs, go to the [Report Timeline](#) page.

Backup Status Summary

A summary of the Backup Status is displayed in left panel.

Total Backups

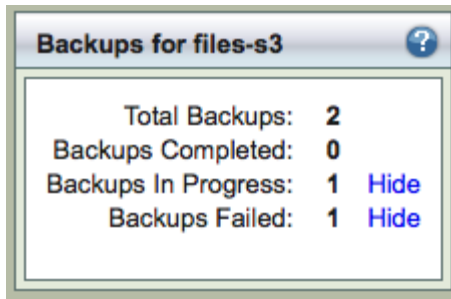
The Total Number of Disk List Entries(DLEs) that are scheduled to be backed up.

Backups Completed

Number of DLEs whose backups has been completed.

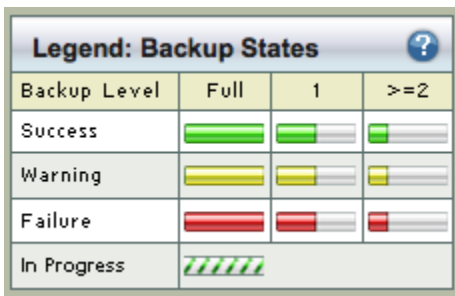
Backups in Progress

The number of backups that are in progress. The **Hide/Show** link lets you toggle the Time Line Monitor display in the Right Hand panel.



Legend

The Legend Box displays the three legends that are used in the Time Line section of the Monitor chart, and shows what the various colors and patterns used in the timeline mean.

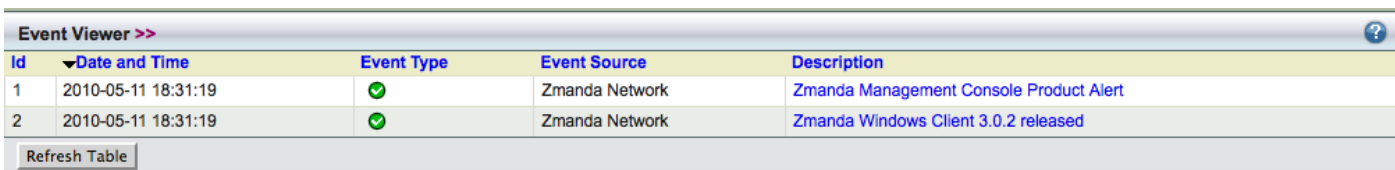


Legend: Backup States			
Backup Level	Full	1	>=2
Success			
Warning			
Failure			
In Progress			

ZMC Alerts

This page describes ZMC **Monitor Alerts** page.

Alerts are information from Zmanda Network. Usually you will get new product updates and alerts. These messages are also posted in your [Zmanda Network home page](#).



Event Viewer >>				
Id	▼Date and Time	Event Type	Event Source	Description
1	2010-05-11 18:31:19	✓	Zmanda Network	Zmanda Management Console Product Alert
2	2010-05-11 18:31:19	✓	Zmanda Network	Zmanda Windows Client 3.0.2 released

[Refresh Table](#)

Monitoring Amanda Events

ZMC **Monitor Events** page provides information about all events on the backup server. This includes backup and restoration.

Event Log table

Event Viewer >> ?					
Id	Date and Time	Event Type	Event Source	Backup Set	Description
32	2010-05-16 00:09:52	✓	Tasks	files-s3	13: Task Completed
31	2010-05-16 00:09:49	✓	Tasks	files-s3	13: Task is running
30	2010-05-16 00:07:48	✓	Tasks	samba	12: Task Completed
29	2010-05-16 00:07:44	✓	Tasks	samba	12: Task is running
28	2010-05-14 19:02:02	✓	Tasks	files-s3	11: Task Completed
27	2010-05-14 19:01:48	✓	Tasks	files-s3	11: Task is running

The table columns are:

Timestamp

When the event occurred, in *yyyy/mm/dd hh:mm:ss* format.

Source

The Amanda/ZMC component that generated the event, such as **planner**, **amdump**, **taper** or the Zmanda Network.

Severity

failure (which require immediate attention), **warning**, or **info**.

Backup Set

Backup Set name (if the message was generated by Amanda)

Description

The event description includes links to the Zmanda Network Knowledgebase or the Amanda wiki that describe the event and any corrective actions if required.

Events can be generated by Amanda Enterprise modules and the Zmanda Network.

- Amanda Enterprise Edition generates lots of events during a backup process and configuration process.
- Zmanda Network provides security and product alerts.
- Events about all backup sets are provided in the same view irrespective of the backup set selected.

All events are retained indefinitely, letting you track error trends across backup runs over time.

Amanda Enterprise log files

Amanda Enterprise log files are stored in **/var/log/amanda**. ZMC UI log messages are stored in **/opt/zmanda/amanda/logs** and **/opt/zmanda/amanda/apache2/logs** directories. The **zm-support** tool captures these logs and sends them to Zmanda Technical Support to assist troubleshooting.

Log Rotate Utility

ZMC works with the Linux **logrotate** utility to prune active logs. If desired, use **crontab** to set up log rotation.

Here is an example **crontab** entry:

```
0 1 * * 1,5 logrotate /etc/logrotate.d/zmc_logrotate (For 1 AM on Monday and Fri
```

Amanda Backup Reports

Amanda automatically generates backup reports after the backup run is completed (or if it fails for some reason). This backup summary report shows statistics about the backup and summarizes any error, warning, or informational messages that were generated during the backup.

If the Amanda server is configured to send mail, reports can be automatically e-mailed after every backup run. See **ZMC Backup How** for email notification configuration.

The Summary Report page is divided into two panels.

- A panel shows a calendar control from which you can select report dates, and a legend that explains the report icons.
- A panel displays the report for date selected on the calendar.

You can select and navigate reports on each backup run using any of the following:

- Browse buttons and timestamp links along the top of the report itself.
- You can enter a date on the left panel and click the **Go** button.
- You can pick a date off the calendar.

When you identify a backup that you want to restore from, clicking the timestamp link in the report will take you directly to the **Restore What** page with the information for that backup already filled in.

All of these methods are described below.

The Summary Report page displays the progress and final status of previous backup runs of a given backup set.

Backup Summary Panel

The Report Summary panel on the right has the following navigational aids:

Date Browse Buttons

The date browse buttons allow to conveniently move back and forth one day (using < or >) or one week (using << or >>) at a time.

Time Stamp links

If you select a date with multiple backup runs, timestamp links with status icons are displayed below the browse button. These links allow you to scroll the report display directly to that timestamp.

The timestamp link within the report window shows the time at which the backup run was initiated along with a status icon. Clicking on a Timestamp link go to the [Restore What](#) page with the date and time automatically filled in.

Displaying Long Field Names in Reports

The widths for various columns in reports are defined in the **amanda.conf** file. If the data for a field exceeds the maximum width defined in **amanda.conf**, the display is truncated. You can change the maximum width. For example, to display longer hostnames, edit the **amanda.conf** file to define a wider column width for hostnames. If you edit amanda.conf to include the following entry:

```
columnspec "HostName=0:15"
```

See the description of the **columnspec** parameter in the [amanda.conf man page](#) for further information.

The **Summary Report** page will display hostnames of up to 15 characters.

Backup Summary

◀◀ 2010-05-16 ▶▶

?

Timestamp: 00:09:52 ✖

TimeStamp: ✖ 00:09:52

DUMP SUMMARY:

DUMPER STATS

TAPER STATS

HOSTNAME	DISK L	ORIG-MB	OUT-MB	COMP%	MMM:SS	KB/s	MMM:SS	KB/s
aetest-client.dhcp.zmanda.com	/etc	MISSING						
localhost	/etc	0	83	83	--	17:17 82.2	17:42 80.3	

STATISTICS:

	Total	Full	Incr.
Estimate Time (hrs:min)	0:01		
Run Time (hrs:min)	0:18		
Dump Time (hrs:min)	0:17	0:17	0:00
Output Size (meg)	83.3	83.3	0.0
Original Size (meg)	83.3	83.3	0.0
Avg Encryption'>Compressed Size (%)		100.0	100.0
Filesystems Dumped	1	1	0
Avg Dump Rate (k/s)	82.2	82.2	--
Tape Time (hrs:min)	0:18	0:18	0:00
Tape Size (meg)	83.3	83.3	0.0
Tape Used (%)	8.1	8.1	0.0
Filesystems Taped	1	1	0
Parts Taped	1	1	0
Avg Tp Write Rate (k/s)	80.3	80.3	--

USAGE BY TAPE:

Label	Time	Size	%	Nb	Nc
files-s3-003	0:11	83M	8.1	1	1

The report is in format described in `amreport(8)` man page. The Compressed size is determined by $(\text{new size}/\text{original size} * 100)$. Dump rate is the rate at which backups are transferred from the client to backup server. Tape rate is the rate at which backup images are written to the media.

Date and Legend Panel

The Summary Report page opens with the calendar set to the current month. Icons (described by the Legend) identify the status of the backup (Success, Warning, Failure, In Progress) as shown below.

Backup Runs	
Success	✓
Warning	!
Failure	✗
In Progress	⌛

◀ May 2010 ▶						
S	M	T	W	T	F	S
25	26	27	28	29	30	1
2	3	4	5	6	7	8
9	10	11	12 !	13 ✗	14	15
16 ✗	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5
2010-05-16						Go

Selecting a Report Date

You can either enter the date (`mm/dd/yy[yy]`) and click **Go**, or click on any icon in the calendar. When you click **Go** (or click on a different date), the report shows any summary data (if any is available) for that date.

Calendar control

The calendar control offers a graphical way to choose a report date. Dates that include backup data have a status icon; click the icon to display the given report (you cannot select dates with no backups to report on). When you click a report icon, it changes color to show it has been visited.

The example above shows April 8 as having been selected and visited.

Note that the calendar icons indicate when backup runs have occurred, which is not necessarily the same as what backups are currently available to Amanda for verification/restore. For example, if you have **Dropped** the media from the [Backup Media](#) page, clicking the calendar icon will result in the following message:

No complete backup available on timestamp *timestamp_value*. Error in Verifying:

This message can also occur after the backup for a particular date has been overwritten because the backup cycle/retention policy (or options selected on the **Backup Media** page) allow that media to be recycled.

ZMC Timeline Reports

The **Report Timeline** tab provides a timeline view of the backup objects/DLEs in the backup set. Note that the most recently completed backup run is not listed here; for the status of the last completed run, see the [Monitor](#) tab.

Report Timeline							
◀◀ 2010-05-12 ▶▶ ?							
HostName	Directory	Time	Level	Clearing Staging Area	Checking Backup Plan	Transferring Backup to Server	Writing to Backup Media
aetest-client.dhcp.zmanda.com	/etc	12:04:58	<div></div>	<div></div>	<div>82m waiting for dumping</div>	<div></div>	<div>82m finished (12:18:55)</div>
localhost	/etc	12:04:58	<div></div>	<div></div>	<div></div>	<div></div>	<div>83m finished (12:12:29)</div>
Toggle Details							

The **Report Timeline** page displays the backup level and progress of previous backup runs, organized by backup object/DLE. The display uses the same legends that the **Monitor** page would use for backups currently in progress.

- One panel displays the calendar controls and a Legend (see next section).
- The other panel as shown above displays color-coded progress bars that report how the backup process happened for each Backup Object/ DLE from the backup set. Hovering the cursor over a progress bar displays a tooltip with details such as duration or amount of data transferred during that phase. As with all **Report** tab pages, a control at the top (the **Calendar Turner**) lets you change the date that is displayed.

The **Report Timeline** panel includes a Calendar turner control. Click either of the two inner single arrows to move backwards and forwards, one day at a time.

Click either of the outer two arrows to move back and forth by a week (seven days) at a time.

Date and Legend Panel

The **Report Timeline** page opens with the current month selected for display. Use the Calendar controls to change as desired.

Date and Calendar Controls

The top portion consists of a **Backup Date** input box along with a **Go** button. Clicking **Go** displays the given date in the Calendar. The input box accepts dates in both mm/dd/yyyy and mm/dd/yy formats.

The Calendar control at the bottom of the panel works the same as it does for the [Summary Report](#).

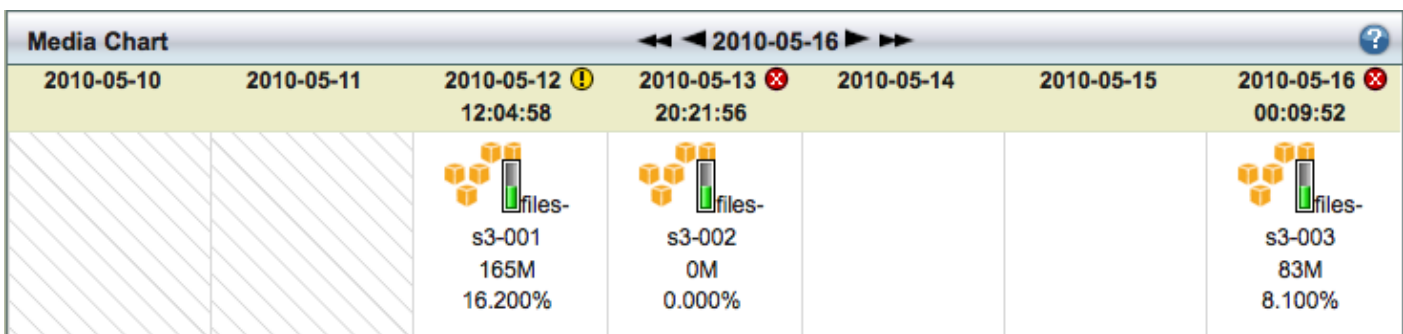
The Legend is same as the [Monitor Backups](#) page.



ZMC Media Reports

The **Report Media** tab provides a media-based view of completed backup runs on a particular date for a particular backup set.

Note that multiple backup objects/DLEs can be backed up to a single media volume. To backup up different backup objects to separate pieces of media, assign the backup objects to separate backup sets. If the objects are small in size, a number of them could fit on a single piece of media.



The **Report Media** page presents a "backup media-centric" view of a past backup run of a backup set. It is especially useful for examining media utilization over time.

The percentage utilization reported for each volume is based on the actual data stored / estimated length provided in the [Backup Where](#) page. The length provided (denominator) is accurate for disk backups. In the case of tape backups, data is written till EOT (end of tape), thus utilization can exceed 100%. In the case of S3 backups, this value not useful.

The Media Chart panel has the a Calender Turner Pressing either of the two inner single arrows allows the users to change reports one day at a time. When either of the outer two arrows signs are pressed, the Calender changes by a week (seven days) at a time.

The Chart also has three controls as shown below: Various dates are spread across the top. Each date carries with it the same legend that is shown in the Summary report. A drop down box below it allows users to choose different Backup runs on the same day, if they exist. On changing the selection from the drop down box changes the Legend and Text summary details displayed below the date. Finally, below the Legend is a short Text summary. The number after the Backup Set name (like -004 or -001) is that of the backup run (of multiple runs, if any) on that day on the same backup set.

Date and Legend Panel

By default the **Report Media** page opens to the current month.









Calendar/Date Controls

These work exactly the same as for the [Summary Reports](#) page.

◀ May 2010 ▶ ?						
S	M	T	W	T	F	S
25	26	27	28	29	30	1
2	3	4	5	6	7	8
9	10	11	12 !	13 x	14	15
16 x	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5
2010-05-16						Go

Legend Sub Panel

The Legend differentiates between two media i.e. Tape and Disk. It also differentiates between Full and partial utilization of the media.

Legend: Media ?		
Media Type	Full	Partial
Disk		
Tape		
Tape Changer		
Amazon S3		

ZMC Client Reports

The **Report Data** tab reports on backup runs for a week of a particular backup set.

The **Report Data** page is divided into two sections:

One panel contains a date/calendar controls and a legend that explains the report status icons that is described in the next section.

The Data Chart list results for each backup object/DLE, using icons that show the type of backup (full or incremental) and whether any errors or warnings were returned. The browse buttons at the top lets you scroll back and forth in time.

Data Chart		◀◀ 2010-05-16 ▶▶ ?						
HostName	Directory	2010-05-10	2010-05-11	2010-05-12 ! 12:04:58	2010-05-13 ✖ 20:21:56	2010-05-14	2010-05-15	2010-05-16 ✖ 00:09:52
...ent.dhcp.zmanda.com	/etc			■	■			■
localhost	/etc			■	■			■

The **Report Data** page provides a **Backup object-by-backup level** view of backup runs for a week of a particular backup set. The graphical display also incorporates the backup status in its display.

Common Navigational Aids

The Right hand Report Data panel has the Calendar Turner (see above). Pressing either of the two inner single arrows allows the users to change reports one day at a time. When either of the outer two arrows signs are pressed, the Calendar changes by a week (seven days) at a time.

Other Controls

The Chart has additional controls shown below: Various dates are spread across the top. Each date carries with it the same legend that is shown in the Summary report. A drop down box below it allows users to choose different Backup runs on the same day, if they exist.

Data Displayed

All the DLEs that exist in the Backup Set are displayed in the two left most columns labeled 'Hostname and 'Directory' respectively. The seven dates are displayed in the next seven columns with the selected date shown in the right most position. In the cell that is formed by the inter-section of DLE and date, one of the nine legends is displayed. The display allows users to see the historical trend of a **Backup Set** broken into its DLEs

Date and Legend Panel

Date and Calendar Controls

The top portion consists of a **Backup Date** input box along with a **Go** button. Clicking **Go** displays the given date in the Calendar. The input box accepts dates in both mm/dd/yyyy and mm/dd/yy formats.

The Calendar control at the bottom of the panel works the same as it does for the [Summary Report](#).



Backup Level

ZMC automatically decides what level of Backup (full or incremental) is appropriate for a backup run. Full backups are labeled level 0 and incremental levels are labeled 1 or 2 and higher levels. When ZMC runs the backup, users are not aware of the level at which it has run. NB: While, immediate backups can be enforced from the [Backup Activation](#) page, that does not mean that Immediate Backups are Full backups! ZMC takes its own decision as to the appropriate level while running the Immediate Backup. Level 0 is shown by a long bar. Level 1 is shown by bar which is about half the size of full bar. Level 2 or above are shown by a bar which is one fourth the size of full bar. These lengths have no intrinsic meanings. Their lengths are just to differentiate them from each other.

Legend: Data				?
Backup Level	Full	1	>=2	
Success	<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	
Warning	<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	
Failure	<div><div></div></div>	<div><div></div></div>	<div><div></div></div>	
In Progress	<div><div></div></div>			

Backup Status

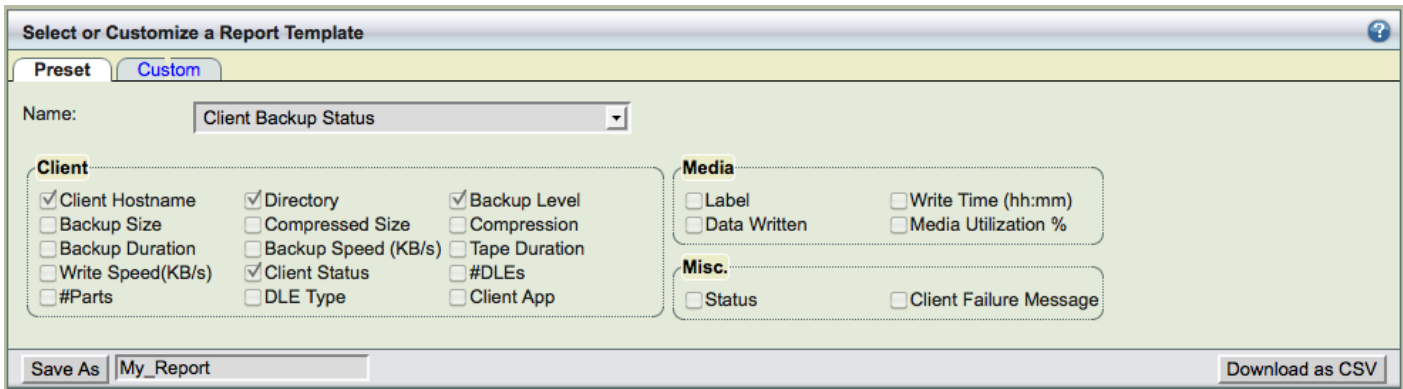
ZMC labels each DLE backup as having completed as Normal, or as with Warning or with Errors. Normal status is indicated by a green color. Warnings are indicated by a yellow color. Errors are indicated by a red color. The combination of level and status gives rise to nine legends shown in the Fig. 3.

ZMC Custom Reports

The various **Report** sub-tabs display backup run data from many perspectives. The **Report Custom** tab lets you create custom reports for a backup set and save them.

The **Custom Report** page display the results of backup set runs. In addition to the standard reports, you can configure and save your own custom report configurations. The results can be saved as a Comma Separated Value (CSV) file for use in other applications and to generate graphical reports.

Select Custom Reports Panel



There are three standard reports; these cannot be edited or deleted.

- **Client Backup Status**
- **Client Performance**
- **Media Performance**

Note also that you can mix Client and Media report options in the same report. Choose the appropriate tab from the **Customize Report** pane to select which type of report you would like to create.

To create your own report, check the desired parameters from the custom reports panel, enter a name in the **Report Name** field and click **Save Report**. The saved user-defined reports will then appear in the drop down box.

Custom reports can be overwritten. A confirmation message box will be displayed when a report is being overwritten.

Report Custom Page Results

Backup Report >>				
Date and Time	Client Hostname	Directory	Client Status	Backup Level
2010-05-16 00:09	aetest-client.dhcp.zmanda.com	/etc	✗	<div><div></div></div>
2010-05-16 00:09	localhost	/etc	✓	<div><div></div></div>
2010-05-13 20:21	localhost	/etc	✓	<div><div></div></div>
2010-05-13 20:21	aetest-client.dhcp.zmanda.com	/etc	✗	<div><div></div></div>
2010-05-12 12:04	aetest-client.dhcp.zmanda.com	/etc	✓	<div><div></div></div>
2010-05-12 12:04	localhost	/etc	✓	<div><div></div></div>

The **Report Results** displays the results of the report defined in the left panel in a table. Data can be sorted in ascending or descending order by clicking on the column headers. You can save the report as a CSV file by clicking the **Save as CSV** button in the lower right. You can then use the CSV file in any application (spreadsheet, wordprocessor, etc.) to generate print and PDF output.

If there are no statistics available for a given parameter because of a backup failure, "na" (not applicable) is displayed in the first column of the row. These "blank" lines are not included in CSV file output.

Displaying Long Field Names in Reports

The widths for various columns in reports are defined in the **amanda.conf** file. If the data for a field exceeds the maximum width defined in **amanda.conf**, the display is truncated. You can change the maximum width. For example, to display longer hostnames, edit the **amanda.conf** file to define a wider column width for hostnames. If you edit **amanda.conf** to include the following entry:

```
columnspec "HostName=0:15"
```

See the description of the **columnspec** parameter in the [amanda.conf man page](#) for further information.

The **Summary Report** page will display hostnames of up to 15 characters.

Backup Data Integrity

Checking Data Integrity of the backup media volumes ensures that the backup data has not been modified after being sent to the backup media.

During various transfers from one media to another, there is a remote possibility that such a change has happened and restoring from such altered data may not give the desired result.

The **Report Data Integrity** page reads and verifies that backed up data image on the media has not been modified. In the case of application clients (such as Windows Oracle), the check does not proceed beyond determining whether the tape is readable. In this case, the message "Could not determine validation for dumper" is displayed and logged. This message can be safely ignored.

If you verify by date, all the tapes associated with each backup run on the selected date are verified (in most cases there will be only one run per day). If you verify by label, only a particular tape is verified. The correct tape(s) must be loaded in the tape device. Verification of each backup run begins by listing the tape(s) required for that run, and then prompting you to load each tape as necessary until all have been checked.

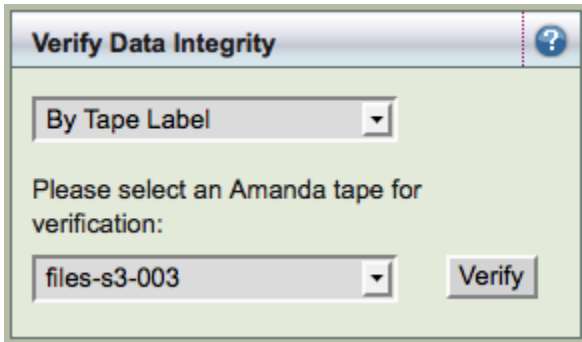
The **Data Integrity** page lets you check out that the already-written media (selected by either date or label) has not been altered after the backup. Run the **Data Integrity** before restoring a backup image.

Procedure to Check Data Integrity

The default method of invoking a **Data Integrity** report is by Tape label.

From the **Please select verification method** drop down box the method of verification can be changed to **By Date**. Selecting by date displays the appropriate interface controls.

Verify By Tape Label

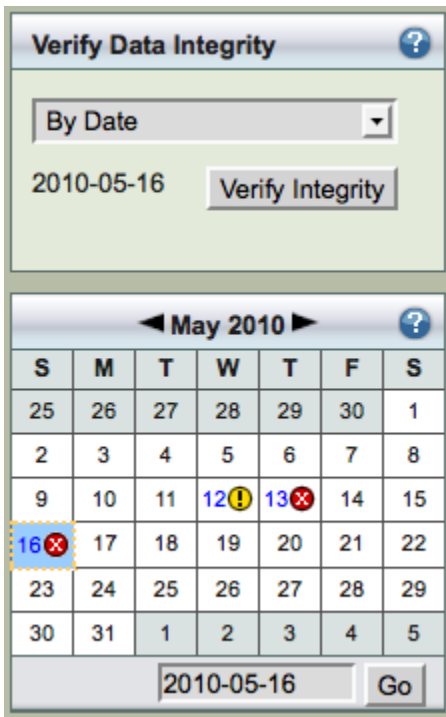
A screenshot of a 'Verify Data Integrity' dialog box. The title bar says 'Verify Data Integrity' with a question mark icon on the right. Inside the dialog, there is a dropdown menu currently showing 'By Tape Label'. Below this, the text 'Please select an Amanda tape for verification:' is displayed. Underneath, another dropdown menu shows 'files-s3-003'. To the right of this second dropdown is a button labeled 'Verify'.

The **Select by Tape Label** dialog is displayed in the left panel by default. It allows you to select an Amanda tape by volume from the dropdown menu. Click the **Verify** button to start the data integrity check of the selected tape.

Procedure - By Date

On choosing 'By Date' option the GUI opens to a page with the Calendar control in left hand panel, with the current month selected for display. Use the calendar controls to change as desired.

Date and Calendar Controls



Verify Data Integrity ?

By Date ▾

2010-05-16

◀ May 2010 ▶ ?

S	M	T	W	T	F	S
25	26	27	28	29	30	1
2	3	4	5	6	7	8
9	10	11	12 !	13 X	14	15
16 X	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

2010-05-16

The top portion consists of a **Backup Date** input box along with a **Go** button. Clicking **Go** displays the given date in the Calendar. The input box accepts dates in both mm/dd/yyyy and mm/dd/yy formats.

The Calendar control at the bottom of the panel works the same as it does for the [Summary Report](#). Click the **Verify** button to start the data integrity check for all the tapes used by the selected backup date.

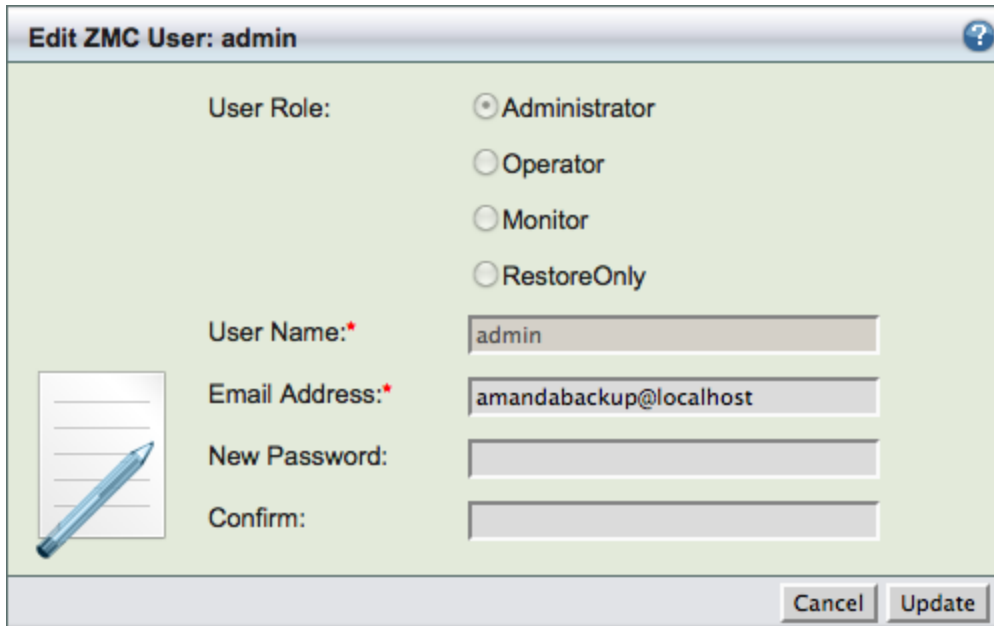
Status of Verification

The right hand panel displays the status of the verification process. You can Clear Output of the message panel by clicking the link.

It is possible to automate the media verification procedure using [amcheckdump](#) command. Please contact Zmanda Support team if you need more information.

ZMC Users

This page allows to you manage ZMC users. You can create, edit and delete ZMC users. ZMC users are different from the operating system/network users.

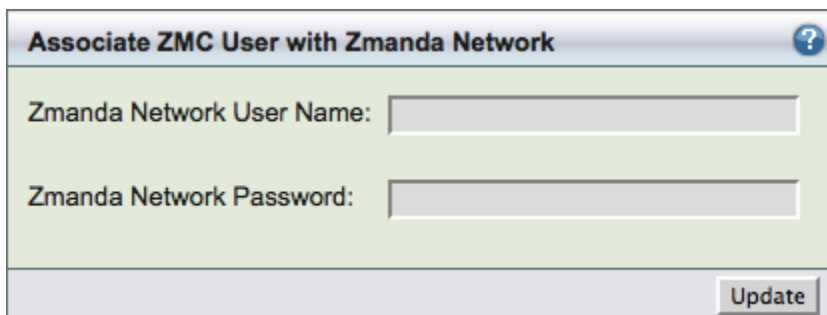


Each user should have a valid email address. This address is used to recover lost ZMC passwords.

There are four user roles as shown above

- Administrator - All backup sets can be accessed and all operations can be performed. New users can be created.
- Operator - Only the backup set that is owned by the operator can be accessed. All operations can be performed on the backup set.
- Monitor - Only Monitor and Reporting pages can be accessed. The user will not be able to make any modifications
- RestoreOnly - Only Restore pages can be accessed. User will not be able to access other pages or make modifications to the backup set configuration.

Associate ZMC User with Zmanda Network



All ZMC users with Administrator role should be associated with Zmanda Network. You can specify the same Zmanda Network authentication information for all users. When the information is entered, it is validated with Zmanda network. Internet access is required for this process.

View all users

The list of ZMC users are shown in a table. Users role and their email address are also displayed. Users with Administrator role can use this table to delete ZMC users.

View, edit and delete users				
All	User Name ▲	User Role	Email Address	Creation Date
-	admin	Administrator	amandabackup@localhost	2010-05-11 18:30:43
<input type="button" value="Invert Selection"/> <input type="button" value="Delete"/>				

Backup Sets

The **Backup Sets** tab allows you to create, edit, duplicate, and delete backup sets. All Amanda configuration is organized as backup sets.

Create Backup Set

Backup Set Owner:*

admin

Backup Set Name:*

new-backup-set

Comments:

(clear)

Report Display Unit:

Gigabytes

Cancel

Add

Backup Set Owner

Choose an owner from the dropdown. ZMC User management is described [here](#).

Backup Set Name

Specify a unique name for the backup set. The name can consist alphanumeric characters. Periods (.) and dashes (-) are also allowed. Backup Set name cannot be more than 255 characters.

Comments

Comments are optional and are intended to serve as a reminder as to why the backup set was created.

Report Display Unit

The unit used in Amanda reports for the backup set. Please note changing the value will not impact reports of the already completed backup runs.

User is warned if there are no storage device or backup clients or backup schedule is associated with the backup set as shown below. These operations must be completed before the backup set can be activated.

Backup Set Owner:

admin

Backup Set Name:

new-backup-set

Comments:

(clear)

Report Display Unit:

Megabytes

⚠ No backup schedule found.





⚠ No backup objects/DLEs found.

⚠ No storage device configured.

Cancel

Update

View/Edit/Delete Backup Sets

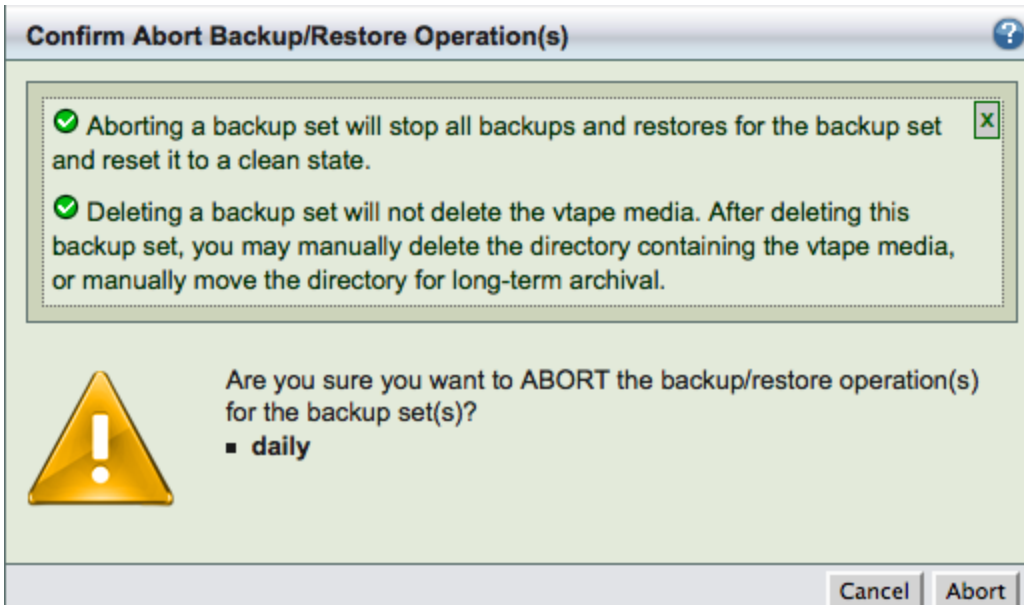
View, add, edit, and delete backup sets >>										
All	Active	Backup Set Name▲	Schedule	Code	✓	✗	Version	Status	Creation Date	Comments
<input type="checkbox"/>		ch	Every Weekday	-	1	0	3.0	OK (device: STK_changer)	2010-05-13	backup set that uses a VTL changer
<input type="checkbox"/>		files-s3	Every Weekday	-	2	0	3.0	OK (device: s3_storage)	2010-05-12	
<input type="checkbox"/>		postgres	Every Weekday	-	1	0	3.0	OK (device: postgres_vtapes)	2010-05-12	
<input type="checkbox"/>		samba	Every Weekday	-	1	0	3.0	OK (device: samba_vtapes)	2010-05-12	

The list of backup sets are displayed in the table. The state of the backup sets: active, inactive and not configured as displayed. The health **status** of the backup set are also displayed. If there are errors, the backup set cannot be used.

Deletion of the backup set removes the given backup set and all the settings associated with it from the ZMC. After deletion, the ZMC will have no record of backups completed by the backup set, but the backups themselves are retained regardless of storage media (staging area, tape, S3, etc.) unless manually deleted. To retrieve backup data from the staging area after the backup set (and hence its catalog) has been deleted, please see [this Zmanda Knowledgebase article](#).

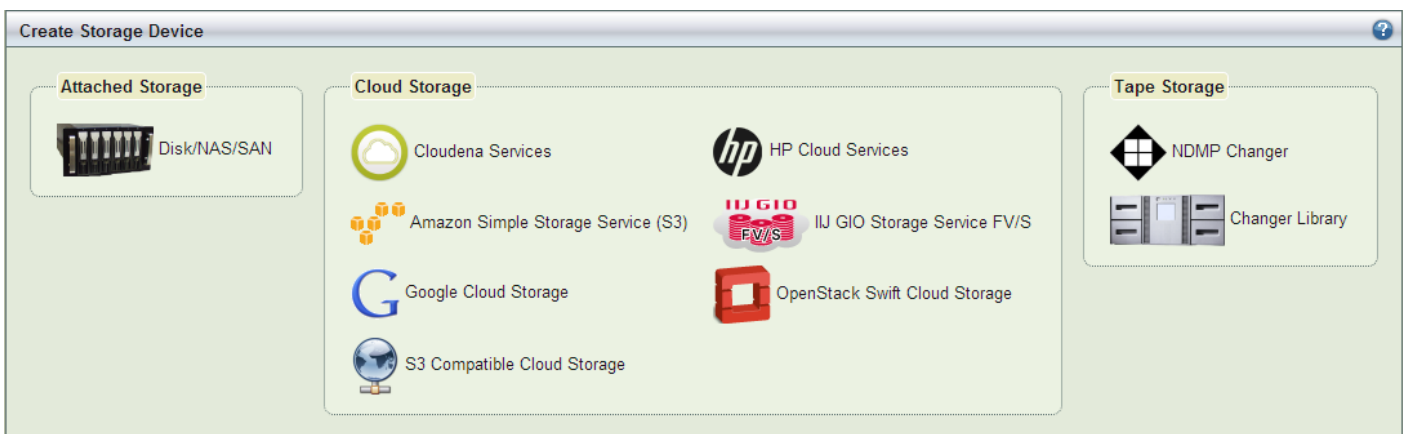
When you click the **Delete** button, the ZMC prompts for confirmation. Proceed with caution as there is no way to undo the deletion. If you attempt to delete a backup set that is active (i.e., selected from the **Backup Set** dropdown at the top right of the page), the ZMC will take you to the **Create a New Backup Set** page after confirmation, otherwise you are returned to the **Admin Backup Sets** page.

You can also abort the backup run for a backup set using the **Abort** button. Aborting a backup stops the backup job and cleans up all the backup processes on the Amanda server. Following confirmation window is displayed to confirm abort operation.



Device Management






The **Admin Devices** page lets you select and configure target devices for backup. Once configured, a device is then *bound* to the current backup set.



When you click on a device type (Attached storage, Cloud Storage or Tape Storage) and click **Add**, the options for that device are displayed in the top pane as shown above. Devices already configured are listed in the bottom pane along with buttons that allow you to **Edit**, **Delete**, and **Use** the device as shown below. When there are more items that can be listed one page, page navigation links at the bottom of the table allow you to move between pages.

NDMP Changer Device can be used only with NDMP application. For more information, please see [NDMP appliance application agent](#) configuration.

All devices must have a **Name**; a descriptive **Comment** is optional. Device-specific options are described in the subsections that follow.

View and edit backup set devices								
All	Type	Device Name	Status	Path	Comments	Used With	Last Modified	By
<input type="checkbox"/>		ATL_1500	OK	/dev/changer	ATL 1500 Tape Changer		2012-07-26 14:33:25	admin
<input type="checkbox"/>		S3_device	OK				2013-02-25 01:46:11	admin
<input type="checkbox"/>		local_vtapes	OK	/mnt/mtapes/	Virtual tapes on local storage		2012-07-26 14:27:24	admin
<input type="checkbox"/>		remote_vtapes	OK	/mnt/datadomain/mtapes/	Virtual tape space on remote volume		2012-12-19 15:42:33	admin
<input type="checkbox"/>		ss7110	OK	10.0.1.20@/dev/scsi/changer			2012-12-20 16:05:37	admin

You can select the specific device to edit from the backup devices table. Clicking on **List** button will list the buckets/containers and objects in case of cloud device. All files (even non-Amanda objects) are listed. Following figure shows an example of Openstack cloud device containers and objects.

List of Cloud Containers for ZMC device: br_openstack_tape								
No Cloud requests performed. All results from cache. Total Containers Found: 2 ZMC Containers: 1 ZMC Containers / Total Containers: 50% Total Containers Reported: 4 Total Objects: 54 Total Bytes Used: 0 GiB								
Show hidden containers that were not created by ZMC? <input type="button" value="List"/>								

Cloud Containers								
All	Backup Set	Creation (Server Time)	Location	ZMC	Bucket Name	Endpoint	Location Constraint	Creation Date
<input type="checkbox"/>	 br-openstack-backupset	2013-09-06 10:50:51	Regionone	Yes	zmc-admin-br-openstack-backupset	http://10.5.52.126:8080/v1/AUTH_6ae6c68201a04ed79c49722e3bc83384	regionone	Fri Sep 6 10:50:51 2013

Disk Device

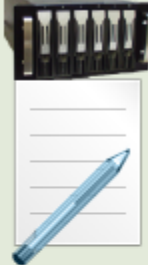
Edit : Disk/NAS/SAN Device
 ?

Name*:

Comments:

Root Path*:
i

Reserved Percent*: %
 i



▾ Advanced Options ?

Output Buffer*: (2 MiB recommended)

Maximum Total DLEs*: Hard limit for total number of DLE backup images.

Cancel Add Update

Root Path The directory to store backup images, specified as an absolute path.

By default, ZMC fills in the field with `/var/lib/amanda/vtapes/`. The **amandabackup** user must have permission to write to this directory, which should also be large enough to hold images for all intended backups. When you bind a disk device profile to a backup set on the **Backup Where** page, it will display this value as the storage location, creating a subdirectory of the root path that matches the backup set name to store images.

Tip: The amount of free space available to hold backup images is an important consideration for fast retrieval of data.

To ensure effective restore capabilities, set aside sufficient disk space to hold more than one full Backup Set worth of Data. Just how many full Backup Sets you should keep on disk depends on the data and your sites requirements for quick restores of accidentally deleted data. The more full backup images stored, the longer the retention of accidentally deleted files. Start with enough space to hold three full Backup Set images and adjust this number as experience dictates.

When using vtapes, you must first decide whether to oversubscribe before you can determine how much disk space to allocate. If oversubscribing, you should still determine by how much you plan to oversubscribe and set up processes on the host system to alert you when oversubscription threatens to prevent backups from occurring.

If you set the number of vtapes in rotation to match the number of vtape slots on the **Backup Where** page (i.e. normal, recommended configuration), then to avoid oversubscription issues you must allocate disk space equal to the number of vtapes times the size of each vtape. Changes to either of

these two parameters require adjusting the amount of disk space allocated for use by ZMC's backup set outside of Amanda, since currently the Amanda does not reserve space before usage.

Tip: Because there is no value in creating a backup of a backup on the same media, Zmanda recommends that the drive that holds the vtapes be excluded from the backup set that points to the vtapes.

Reserved Percent Percentage of the file system that should be free. This value depends on the file system in use. All file systems have reserve space. This value is used for ZMC disk space warning messages.

Output Buffer and **Maximum Total DLEs** need not be adjusted unless Zmanda support team recommends changes in these value for the backup set configuration.

Amazon S3 Device

Setting up an Amazon S3 storage device requires that you sign up for an Amazon S3 account and obtain the Access Key ID and Secret Key as described in [Setting up an Amazon S3 Account for Use with Amanda](#). If the S3 account is already set up, you can obtain the Key ID and Secret Key by logging in [here](#) (Amazon username and password required), from where you can cut and paste the Key ID/Secret Key.

Create : Amazon Simple Storage Service (S3) Device

Name*: S3Production

Comments:

Access Key*: ACCESSKEYSHOULDBEHERE

Secret Key*: SECRETKEYSHOULDBEHERE

User Token:

Storage Option: Reduced Redundancy (RRS)

Advanced Options

Use API Credentials*: ☒ (API keys recommended over username/password)

Secure Communications*: ☒ (recommended, but using IPSec/VPN is faster)(recommended, unless running AE on EC2 at Amazon)

Cloud URL Path: Leave blank if not sure.

Bucket Names*: ☒ Use DNS ☐ Use URL path

Cloud Object Size*: 256 MiB

Cloud Encrypt: ☒ none ☐ AES256

Reuse Connections: ☒

Maximum Total DLEs*: 2000 Hard limit for total number of DLE backup images.

1 TiB Backup

Size	# Objects	0.01/1000	Time
256MiB	~4,000	.04	100%
128MiB	~8,000	.08	102%
64MiB	~16,000	.16	107%
32MiB	~32,000	.32	116%
16MiB	~64,000	.64	135%
8MiB	~128,000	1.28	171%

Cancel

Add

Access Key

The Access Key ID you obtained when signing up for S3 storage.

Secret Key

The Secret Key you obtained when signing up for S3 storage.

User Token

This is included for compatibility with the Amanda Enterprise 2.6.4 payment model, which required an "Amazon S3 Certificate." If you have such an account, this field will be automatically filled with User Token for the account, which will be required if you need to access the data without using Amanda.

Amanda Enterprise 2.6.4 customers please note: Amanda will always support the old payment model, but if you wish to change to the new payment model, please contact the Zmanda Support team.

Account migration does not happen automatically, and the best migration method will depend on how Amanda/S3 was used at your particular site.

Storage Option

It can be either **standard** (Amazon S3 will complete redundancy and more availability) or **Reduced Redundancy** (Amazon S3 with reduced redundancy and less expensive storage).

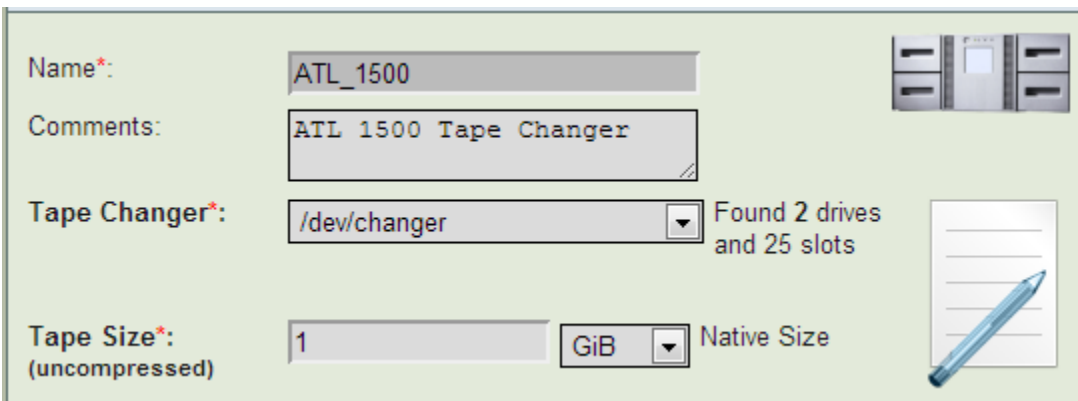
Advanced Options

You can disable **secure communications** to the cloud. This will increase performance at the cost of security. This is recommended only when you are not backing up sensitive data and backing up from Amazon EC2 to Amazon S3.

You can also enable **AES 256bit Cloud Encryption** provided by Amazon. The backup data is stored at rest in S3 encrypted. The keys are managed by Amazon. Please note that you can enable this option even if you are using Amanda client or server encryption with your private keys.

Other options should be changed only if they are recommended by Zmanda Support team.

Tape Changer Device



Tape Changer

- **Please select:** Click this option to display a list of tape changers sensed by the operating system and the ZMC. Choose one. Number of drives on the changer and tape slots are also displayed.

- **Other** Click this option to enter a path to a Tape Changer device that may be currently inaccessible or not detected by the ZMC. When the **Backup Where** configuration is saved, the Tape Changer name and its location are not validated. ZMC attempts to discover tape changer and tape drive device names automatically. The devices must be readable and writeable by **amandabackup** user.

Tape Size

Specify the size of the tape in MB (megabytes), GB (gigabytes) or TB (terabytes). This should be the uncompressed size of tapes used by the changer. It is important not to use the number quoted by marketing material. You can use the [amtapetype](#) command to allow Amanda to estimate the size of tape. The execution of **amtapetype** will probably take hours because it must perform I/O operations on the entire tape length.

Note on Block Size: Although you cannot change the block size of the device using the ZMC, you can manually edit the [amanda.conf](#) file to specify a different block size if necessary. See the description of **device_property** in the link for details.

There are Advanced Options available for tape changer device.

▼ Advanced Options ?

Misc.

Total Slots*: tape slots available

Output Buffer*: (2 MiB recommended)

Next Drive Strategy:

Bar Code Reader: ☒ Exists?

"mt offline": ☐ Does your robot require an "mt offline", before "mtx unload"?

Timing in seconds

Status Poll Interval: Delay period between checking status with MTX

Unload Delay: Delay period, after unloading a tape.

Eject Delay: Delay period, after ejecting a tape, before unloading to storage slot.

First Poll: Delay period, after loading a new tape, before polling for status.

Poll Frequency: Delay period, between polling for status.

Max Wait for Ready: wait up this many seconds for the drive to be ready

Bar Code Reader

Check this box if a barcode reader is attached to the changer so that the ZMC can use it to identify tapes.

All other advanced options should be used when advised by Zmanda Support Team.

Google Cloud Storage

You will have to set up Google Account and sign up for Google Cloud Storage. You will have get access and secret keys from [Google APIs console](#) under Legacy Access from Google Cloud Storage tab.

Create : Google Cloud Storage Device

Name*:

Google_Storage

Comments:

Google cloud storage used for file system backups

Access Key*:

GOOGSJKALDILLEKAL3E

Secret Key*:

889d8adaasdsad8asf9fas9f0f8s9f8s0d9f89s0f

Advanced Options

Secure Communications*:

☒ (recommended, but using IPSec/VPN is faster)

Cloud Object Size*:

256 MiB

Reuse Connections:

☒

Maximum Total DLEs*:

2000

Hard limit for total number of DLE backup images.

1 TiB Backup

Size	# Objects	0.01/1000	Time
256MiB	~4,000	.04	100%
128MiB	~8,000	.08	102%
64MiB	~16,000	.16	107%
32MiB	~32,000	.32	116%
16MiB	~64,000	.64	135%
8MiB	~128,000	1.28	171%

Cancel

Add

Access Key

The Access Key ID you obtained when signing up for Google Cloud Storage.

Secret Key

The Secret Key you obtained when signing up for Google Cloud Storage.

Secure Communications

Enable secure communication between Amanda server and Google Cloud Storage. It is recommended to use secure communication.

Cloud Object Size

Size of objects stored in Google Cloud. If there is a network transmission failure, the object is re-transmitted. For network connections that are reliable, higher cloud object size can be used.

Reuse Connections

When a cloud communication failure occurs, Amanda will reuse the existing communication handle instead of recreating a new one. The default is to reuse communication handles.

OpenStack Swift Cloud Storage

Amanda Enterprise talks to Open Stack Swift cloud using Keystone authentication mechanism and without using S3 compatibility layer.

Edit : OpenStack Swift Cloud Storage Device

Name*:

Comments:

Identity/Auth Service:

Use API Credentials*: ☐ (API keys recommended over username/password)

Tenant Name*:

Access Key*:

Secret Key*:

Advanced Options

Secure Communications*: ☐ (recommended, but using IPSec/VPN is faster)

Auth Plugin Path*:

Bucket Names*: ☐ Use DNS ☒ Use URL path

Cloud Object Size*:

Reuse Connections: ☒

Use Untrusted Certificate*: ☐

Trusted Site Certificate:

Maximum Total DLEs*: Hard limit for total number of DLE backup images.

1 TiB Backup

Size	# Objects	0.01/1000	Time
256MiB	~4,000	.04	100%
128MiB	~8,000	.08	102%
64MiB	~16,000	.16	107%
32MiB	~32,000	.32	116%
16MiB	~64,000	.64	135%
8MiB	~128,000	1.28	171%

Cancel Add Update

Identity/Auth Service

The IP address and port number used by Keystone authentication service.

Tenant Name

Tenant name or the project name that can be used to isolate resources for a group/project.

Access Key

Name of the user in the tenant or the access key.

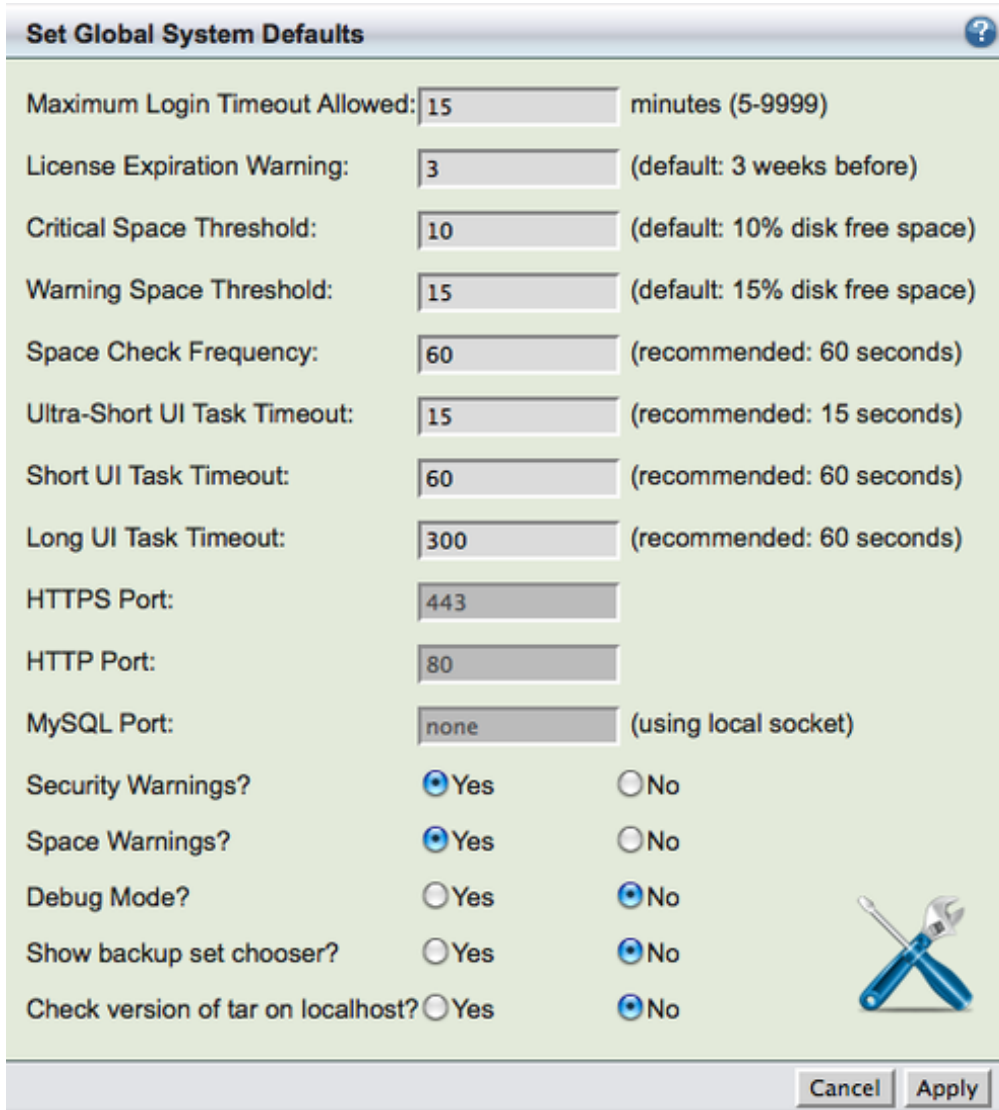
Secret Key

Password of the user or the secret key.

ZMC Preferences

This page allows users to manage ZMC preferences and as well as execute Amanda and some system command on the Amanda server.

Set Global System Defaults



Parameter	Value	Default/Recommended
Maximum Login Timeout Allowed	15	minutes (5-9999)
License Expiration Warning	3	(default: 3 weeks before)
Critical Space Threshold	10	(default: 10% disk free space)
Warning Space Threshold	15	(default: 15% disk free space)
Space Check Frequency	60	(recommended: 60 seconds)
Ultra-Short UI Task Timeout	15	(recommended: 15 seconds)
Short UI Task Timeout	60	(recommended: 60 seconds)
Long UI Task Timeout	300	(recommended: 60 seconds)
HTTPS Port	443	
HTTP Port	80	
MySQL Port	none	(using local socket)
Security Warnings?	<input checked="" type="radio"/> Yes <input type="radio"/> No	
Space Warnings?	<input checked="" type="radio"/> Yes <input type="radio"/> No	
Debug Mode?	<input type="radio"/> Yes <input checked="" type="radio"/> No	
Show backup set chooser?	<input type="radio"/> Yes <input checked="" type="radio"/> No	
Check version of tar on localhost?	<input type="radio"/> Yes <input checked="" type="radio"/> No	

Above panel shows the list of ZMC preferences that are available. These preferences are common to all ZMC users. Please contact Zmanda Support Team before making changes to these parameters. Some interesting preferences are

Maximum Login Timeout Allowed

The timeout after which an idle ZMC session times out. User will have to login again. You can resume the ZMC session by selecting the check box.

License Expiration Warning

The **Admin Licenses** page provides warnings about expiring licenses. Many features in Amanda Enterprise are licensed. When license expires, the backups will not be performed. It is important to take notice of warning and contact Zmanda Support about license renewal.

Usually Zmanda Support team will also contact you about the license renewal.

Critical Space Threshold

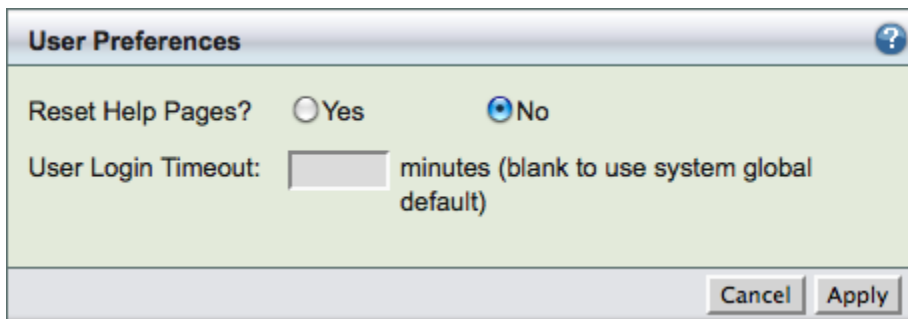
Warning Space threshold

Space Check Frequency

It is important to have sufficient space available in certain directories such as */tmp*, */opt* for ZMC to function properly. These three preferences control the messages regarding lack of disk space.

User Preferences

These are preferences specific to the ZMC user.

A screenshot of a 'User Preferences' dialog box. The title bar says 'User Preferences' with a question mark icon. Inside, there are two sections. The first section is 'Reset Help Pages?' with two radio buttons: 'Yes' (unselected) and 'No' (selected). The second section is 'User Login Timeout:' followed by a text input field containing '15', then the text 'minutes (blank to use system global default)'. At the bottom right, there are 'Cancel' and 'Apply' buttons.

User Login Timeout specifies the time after which the idle user session will be timed out. User will have to log in again.

Amanda Enterprise Licenses

The **Admin Licenses** page provides a view for the list of Amanda feature licenses in the license file */etc/zmanda/zmanda_license*. The Amanda licenses are counted based on the number of hosts protected. Amanda server is also a client and uses up a license.

Licensed Features Summary

The licensed feature summary shows the number of licenses available for each feature (**Licensed**), number of licenses that have been used (**Used**), number of license remaining(**Remaining**), number of licenses close to expiring (**Expiring**) and number of licenses that have expired (**Expired**).

Licensed	Used	Remaining	Expiring	Expired	Feature
1	-	1	-	-	Microsoft Exchange
1	-	1	-	-	Microsoft SQL Server
1	-	1	-	-	Microsoft Sharepoint
Buy now	-	-	-	-	NDMP
1	-	1	-	-	Oracle on Linux/Solaris
Buy now	-	-	-	-	Oracle on Windows
1	1	0	-	-	PostgreSQL
5	3	2	⚠ 2	-	Unix: Linux, Solaris, Mac
Buy now	-	-	-	-	VMWare ESX Virtual Machine
11	-	11	-	-	Windows: files, state
✓	1	-	-	-	Amazon S3 Device
✓	2	-	⚠	-	Disk Device
✓	-	-	⚠	-	Single Tape Device
✓	1	-	⚠	-	Tape Changer Device

When the number of licenses remaining is less than 2, the number appears in red color. Expiring licenses warning can be controlled using License Warning preference in **Admin Preferences** page.

When feature licenses are used up, it is not possible to add any more DLEs of that type. If Amanda server license (Unix file system) expires, all backups will not performed. So, it is important to renew licenses before the expiration date.

Zmanda team will contact you via email before the license expiration date.

Licenses Used

Licenses Used table shows how the licenses are used by various hosts and backup sets configured on the Amanda server.

Licenses Used by Feature			
Hosts	Feature	Total DLEs	Hosts
1	PostgreSQL	1	localhost
3	Unix: Linux, Solaris, Mac	4	rhel5-server-64, aetest-client.dhcp.zmanda.com, localhost
Licenses Used by Host			
Features	Host	Total DLEs	Features
1	aetest-client.dhcp.zmanda.com	1	Unix: Linux, Solaris, Mac
2	localhost	3	Unix: Linux, Solaris, Mac; PostgreSQL
1	rhel5-server-64	1	Unix: Linux, Solaris, Mac
Devices Used by Backup Sets			
Backup Sets	Device	Total DLEs	Backup Sets
1	Tape Changer Device	1	STK_changer
2	Disk Device	2	postgres_vtapes, samba_vtapes
1	Amazon S3 Device	1	s3_storage

Restoring from ZMC

ZMC **Restore What** page is the first step in restoring files using the Zmanda Management Console. This section covers generic file system restores; for details on application-specific restores, see the appropriate section of the Zmanda [Application Agents for Amanda Enterprise Edition User Guide](#).

The **Restore What** page specifies what to restore. It allows you to choose a **single file** or a **single directory** or **all directories/files under a single directory**.

Selecting what to restore is a two step process:

1. Use the **What would you like to restore from** panel to specify a broad definition of which backup image is to be restored.
2. Use the **Restore Preference** option to specify more detailed information about what is to be restored. You can use the **Restore All** option to restore the complete backup image or use **Explore & Select** option to view the restore tree and select files manually or use **Search Specific files** to search for specific files/folders to do the restoration.

What would you like to restore from: DailyLocalDisk

Restore To: ☒ Today, OR i
 (on or before)

Restore Device:*

Media Explored: 2014-06-02 01:15 DailyLocalDisk-AA-002
 2014-05-30 01:15 DailyLocalDisk-AA-054

Host Name:* Edit

Alias/Directory/Path:* Edit

Restore Preference: ☒ Restore All
☐ Explore & Select
☐ Search Specific files

File/Pattern to Restore:*

Search for:

Reset Next

Backup Date

This field is optional. The current date or current time is used, if no date or time is given. Numerous formats for date and time are acceptable, including "yesterday", "4 days ago", "2011/07/28", etc. Click the information **i** icon next to the field for Date/Time formats to see more examples. You can click the Timestamps in the lower portion of the **Backup Reports** page to automatically select the date and time for a particular backup.

ZMC will choose the most recent, successful full backup image created on or before the date and time selected. The most recent incremental backups (if any) for the full backup are automatically selected as well. All items in the full and incremental(s) are restored using "Express" restore, or shown by "Explore & Select" (see below). To exclude one or more of the incremental(s), choose a date and time prior to the backup date of the incremental backup.

Host Name

Required. The name of host whose backup has to be restored. The host name must exactly match the name given in the ZMC **Backup What** page for this backup set. You can click **Edit** button to select various possible options.

Alias/Directory/Path

Required. The directory/path/Alias that has to be restored. The value has to exactly match the name given in the ZMC **Backup What** page for the backup set. You can use the **Edit** button to get possible values for this field in the host specified in the **Host Name** field.

Search Specific Files Option

When you search for specific files to restore, you can provide specific pattern to search for and you can do exact match or the files/folders that begin with the pattern or ends with the pattern. All folders/directories with the pattern are displayed in the panel below for selection.

What would you like to restore from: DailyLocalDisk

Restore To: ☒ Today, OR ⓘ

Restore Device:*

Media Explored: 2014-06-24 01:15 DailyLocalDisk-AA-024

Host Name:*

Alias/Directory/Path:*

Restore Preference:

☐ Restore All

☐ Explore & Select

☒ Search Specific files

File/Pattern to Restore:*

Search for:

Select Directories / Files to Restore from: DailyLocalDisk ⓘ

C:/ (6)

Select: All | None | Invert

☐ (20) Program Files (x86)/

☐ (14) Program Files/

☐ (23) Users/

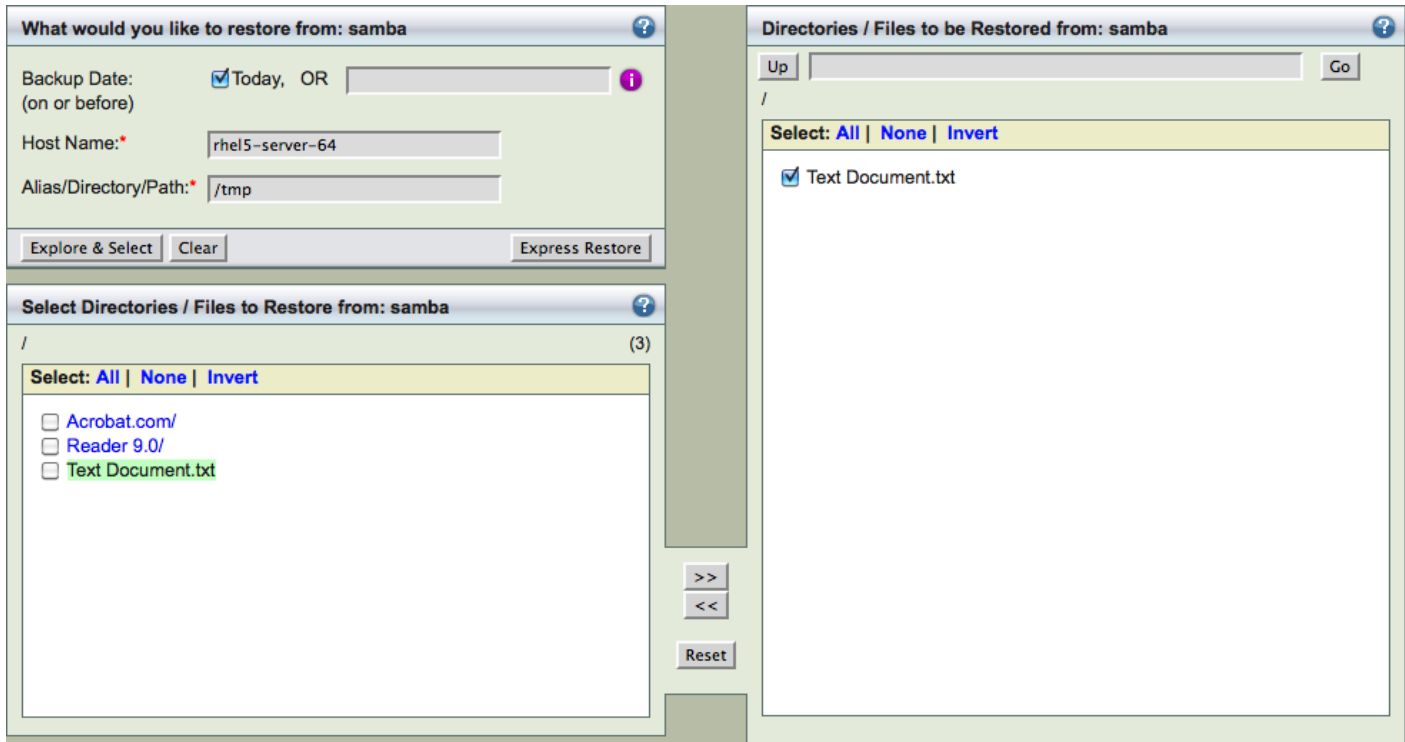
☐ (2975) Windows/

☐ (47) restore/

☐ (15) tmp/

Explore & Select Option

The **Explore** button acts as a link between the initial host and directory that are entered in the **Restore Setting** panel and the detailed file listing displayed underneath in **Select directories/files** panel as shown below. This panel is in the lower left hand side of the page. If there are any errors in entering the data in the **Restore** setting panel, they will be caught when you click **Explore**. If everything has been correctly entered, Explore Button will populate the 'Select directories/files panel below it.



The screenshot shows the ZMC Restore interface with two main panels. The left panel, titled "What would you like to restore from: samba", contains fields for "Backup Date: (on or before)" with a checkbox for "Today" and an "OR" option, "Host Name:" with the value "rhel5-server-64", and "Alias/Directory/Path:" with the value "/tmp". Below these fields are buttons for "Explore & Select", "Clear", and "Express Restore". The right panel, titled "Directories / Files to be Restored from: samba", shows a file list with "Text Document.txt" selected. Below the file list are buttons for ">>", "<<", and "Reset".

Select the directories/files to be restored by clicking left and right arrows. Selected entries should be in the right panel.

Explore process can take time and the time taken depends on the number of entries in the Amanda index for the **Host Name** and the **Directory**.

The number next to directory (in the above figure - **(3)**) shows the number of files in the explore window pane.

Restore All Option.

Use this button to restore complete backup image. This is short-cut to avoid the exploring the backup images for quick restores. This is useful in case of complete restoration of the application or the file system. For some applications, only Restore All is supported.

Note that although the ZMC allows multiple users to access the Restore pages from the same Amanda server, the ZMC server can only explore one backup object/DLE at a time. A warning message is displayed if multiple users are accessing the restore pages for the same backup set.

ZMC Restore Where

The **Restore Where** page specifies the destination directory for recovered files. Before attempting a restore, you may want to review the [Requirements for Restore Clients](#) section of the pre-installation checklist.

The **Restore Where** defines where the data is to be restored. It also lets you control how the file restore operation proceeds. For example, you can choose whether to permit overwriting of existing files.

When you are restoring to Amanda Linux, Solaris or Mac OS X client, configuration is required on the client. See [next](#) section. For Windows client, no configuration is required other than client software must be installed on the machine being restored to.

Restoring to a Amanda Linux/Solaris/Mac OS X Client

1. If you have not done so, install the Zmanda client software. The installation procedure is described in the [Linux](#), [Solaris](#) and [Mac OS X](#) client manuals.
2. Make sure that **zmrecover** is enabled on the client (enabled by default). Examine `/etc/xinetd.d/zmrecover` to make sure it includes the following line

disable = no

Or, on Mac OSX systems, run the following command:

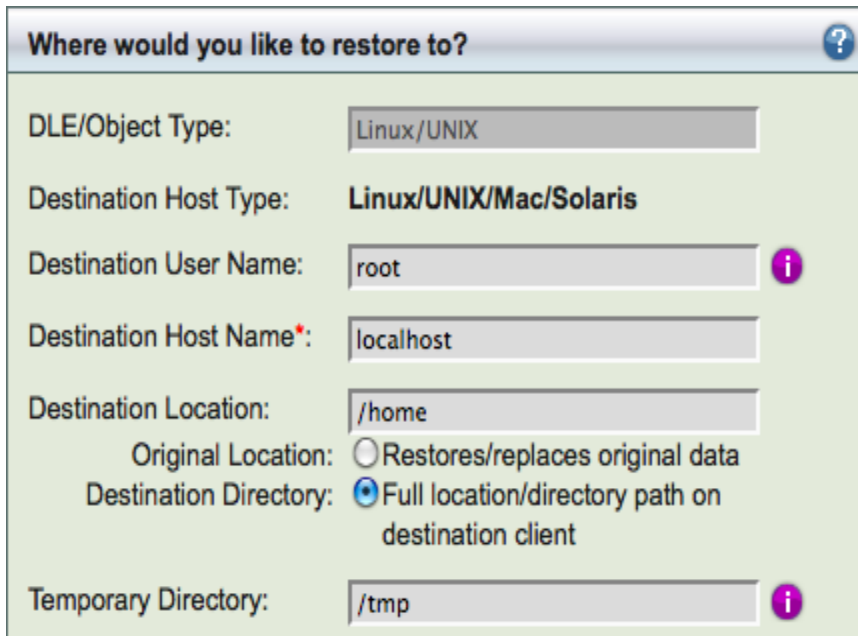
launchctl load -w /Library/LaunchDaemons org.amanda.zmrecover.plist.

For Linux/Solaris/Mac OS X clients, the `/etc/amanda/amanda-client.conf` file must be edited on the client (if the client and server do not reside on the same machine). Specifically, the following entries must be edited:

```
index_server "localhost"          # your amindexd server
tape_server  "localhost"          # your amidxtaped server
```

Change **"localhost"** to match to the hostname of the Amanda backup server.

Where would you like to restore to?



Where would you like to restore to?

DLE/Object Type:

Destination Host Type: **Linux/UNIX/Mac/Solaris**

Destination User Name: ⓘ

Destination Host Name*:

Destination Location:

Original Location: ☐ Restores/replaces original data

Destination Directory: ☒ Full location/directory path on destination client

Temporary Directory: ⓘ

DLE/Object Type and **Source Host Type** are non-editable fields and they are provided for information to fill other fields.

Destination Host Name

The Destination Host is the machine where you want restore the files. It need not be the same machine that originally contained the backed up data.

If no **Destination Host** is specified, the files are restored to the Amanda server machine.

Destination Host Type

Choose either **Linux/Unix/Mac/Solaris** or **Windows**.

- If the **Destination Host** is defined as **Linux/UNIX** the data will be restored as user **root**. You can specify a different user such as **amandabackup** to execute the restore. Please note restoring as a non-root user may not preserve directory/file ownership.
- If the Destination Host is running Windows, then the data will be restored as user **amandabackup**. This cannot be changed.

If you are restoring Windows backups to Linux/Unix/Mac/Solaris host, the data is restored as zip file which can be copied to Windows machine and extracted using pkzip or winzip tools that support ZIP64 format. Example message that appears in ZMC:

When restoring Windows backup image files created by ZWC to non-Windows systems, only the raw image file (ZIP formatted) can be restored. Please make sure adequate space exists for restoring the entire backup image file.

Linux/Unix/Mac/Solaris client backups cannot be restored to Windows.

Destination Username

Specifies the OS user on the destination machine that will provide access to the restore process. If the **Destination Host** has Amanda client installed, the user has to be **root**. For restoration of Windows backups, this field cannot be changed.

If you are restoring to a machine that does not have Amanda client installed, restoration will be done using **ssh**. **ssh** must be configured on the Destination Host and the **ssh** user can be specified as Destination Username. **ssh** user password will be required at the start of restoration process. Only file system (not applications) data can be restored using **ssh**.

Destination Directory

The directory on the **Destination Host** where the files will be restored. It can be either **Original location** from where backup was performed or an alternate location (**Destination Directory**).

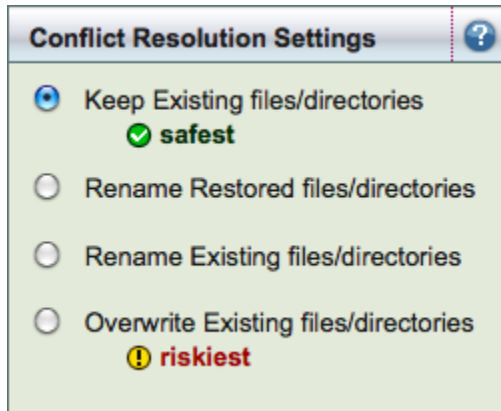
Warning: The **Destination Directory** MUST be specified as an absolute path on the **Destination Host**.

Temporary Directory

A directory on the **Destination Host** (not applicable for Windows hosts) that ZMC will use temporarily during the restore process. Please note that there should be sufficient space to store the whole backup image in this directory. The **amandabackup** user should have permissions to write to this directory.

Conflict Resolution

Conflict Resolution is displayed only for Windows, Linux, Solaris and Mac OS X file systems.



The radio buttons in the above panel lets you set how to handle filename conflicts in the destination directory (select one):

Keep Existing Files

If checked (the default), ZMC does not overwrite any existing files on the destination directory. Conflicting files are simply skipped.

Overwrite Existing Files

If checked, the entire pathnames of existing files on the destination will be overwritten by the backed-up versions, if any. **Use with extreme caution.**

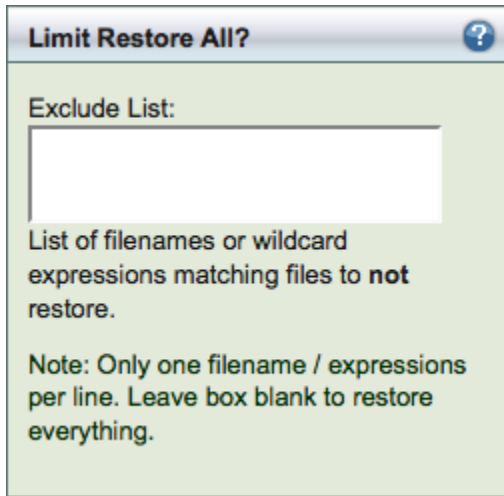
Rename Existing Files

If checked, name conflicts are resolved by renaming the existing file (and the entire directory path to that file) using the following convention: *original_filename.original.timestamp*.

Rename Restored Files

If checked, name conflicts are resolved by renaming the existing file (and the entire directory path to that file) using the following convention: *original_filename.original.timestamp*.

Exclude

A dialog box titled "Limit Restore All?" with a question mark icon in the top right corner. It contains a text area labeled "Exclude List:" with a light green background. Below the text area, there is a note: "List of filenames or wildcard expressions matching files to **not** restore." and "Note: Only one filename / expressions per line. Leave box blank to restore everything."

Limit Restore All?

Exclude List:

List of filenames or wildcard expressions matching files to **not** restore.

Note: Only one filename / expressions per line. Leave box blank to restore everything.

Express restores allows users to exclude file names or files/directories that match the exclude pattern. Multiple files or patterns (filename globbing syntax) can be specified.

Running Restore Task

The **Run Restore** page is the last step in the recovery process. In the **Run Restore** page review the restore options you have specified, start the actual restore process, and monitor its progress. Only one restore process can be performed for a backup set in ZMC.

Restore From, Restore To, Tapes needed panels (as shown below) provide information about the restore job. Make sure the required tapes are in the tape changer (in the slots reserved for the backup set) if the restoration is from a tape.

Click the **Restore** button to start the restore process.

Restore From ?

DLE/Object Type: **Linux/UNIX**

Initial Host:

Initial Directory/Path:

Restore To ?

Destination Host:

Destination Directory: **! Original Location**

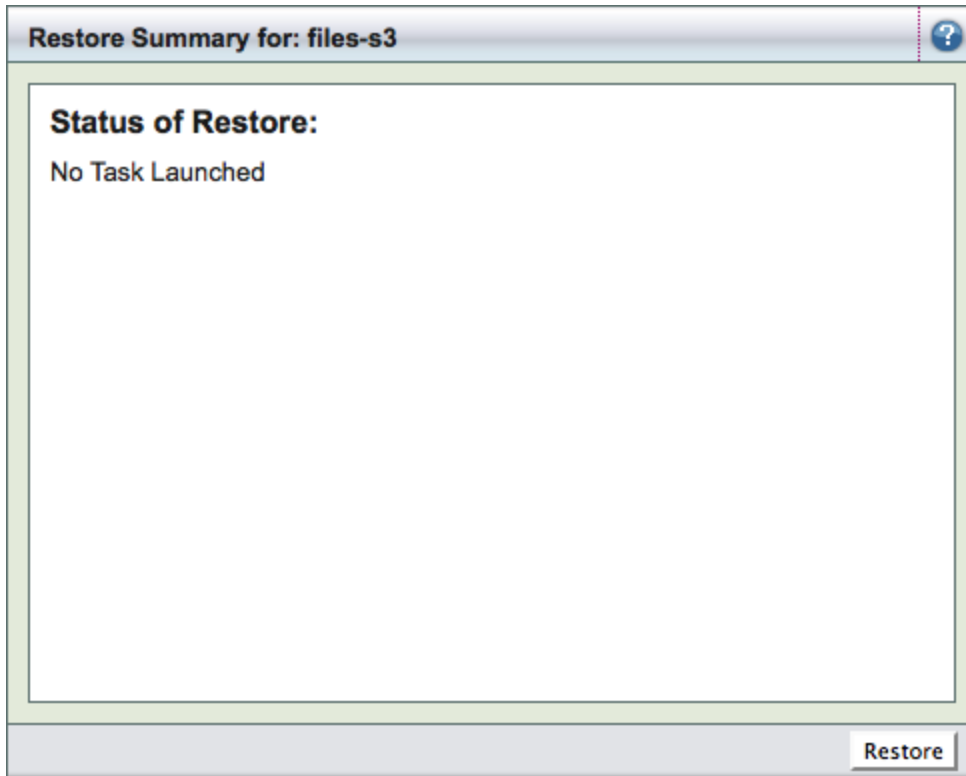
Conflict Resolution: **Keep Existing**

Tapes Needed ?

Date	Time	Label
2010-05-12	12:04:58	files-s3-001
2010-05-13	20:21:56	files-s3-002

The Restore Summary panel shows the restore process progress. The backup image is read from the media and it is transferred to the Destination Host. The relevant files/directories are extracted or the applications are restored on the Destination Host. If tapes are required, you will be prompted in this window.

If this restore is being performed using **ssh**, you will requested for the password.



Notes on Browsers

This section discusses some browser issues you should be aware of when using the Zmanda Management Console.

The Browser "Remember Input" Option

Most browsers have an option to "remember" what is entered in forms. In Firefox, this is set in the **Preferences->Privacy** tab. Turning this option on can result in fields being pre-filled with incorrect parameters, as the browser merely remembers the ZMC page and knows nothing about Backup Sets and other ZMC input.

Page Refresh Considerations

When refreshing a page after a POST operation, most browsers display a dialog asking if you want to repeat the operation. Avoid reposting form data when using ZMC. Re-posting form data can result in undesired consequences.