

# Sun Fire™ X4500 Server as Storage Node for Sun StorageTek™ Enterprise Backup Software 7.4

Ryan Arneson  
April 2008  
Sun Microsystems, Inc.

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements. Use is subject to license terms. This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd. X/Open is a registered trademark of X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Solaris, Sun Fire, and StorageTek are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

This product is covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

# Table of Contents

Summary.....	3
Introduction.....	3
Configuration.....	4
Configuring the Disks.....	4
Creating Zpools.....	5
Configuring ZFS.....	6
Installing EBS Options.....	7
Configuring the EBS Storage Node.....	7
Configuring Each ZFS File System as a Storage Device.....	7
Creating a Media Pool.....	9
Configuring Client Backups.....	12
Frequently Asked Questions.....	13
For More Information.....	13

## Summary

This guide is an introduction to configuring the Sun Fire™ X4500 server as a disk-cache storage node for Sun StorageTek™ Enterprise Backup Software (EBS). It provides examples of configuring the zpool and ZFS file systems on the Solaris™ 10 08/07 Operating System and configuring the EBS application to utilize the Sun Fire X4500 server.

## Introduction

A Sun StorageTek EBS storage node can be used to improve performance by offloading from the Sun StorageTek EBS server much of the data movement involved in a backup or recovery operation. Any Sun StorageTek EBS client can be configured to use the Sun Fire X4500 server as the storage node.

The Sun Fire X4500 server delivers massive storage capacity and remarkable throughput making it an ideal nearline storage platform for backup and restore applications. Combining the power of the Solaris 10 OS with the data integrity and simplified administration of ZFS, the Sun Fire X4500 server becomes an ideal candidate for streamlining and improving backup and restore operations.

Figure 1 shows an EBS configuration with the Sun Fire X4500 server acting as a storage node. The default NIC configuration in the Sun Fire X4500 server allows up to four physical Gigabit Ethernet connections to be attached. The four connections can back up four separate subnets or the connections can be aggregated into a single subnet for improved throughput using `d1adm(1M)`.

The Sun Fire X4500 server can also be connected to a tape library for image duplication or archiving purposes (using Fibre Channel or SCSI connections). In addition, the Sun Fire X4500 server can be used as the primary data cache for staging a tape library for long-term storage. Data is moved to tape after certain file system utilization or time-based thresholds are reached. As an alternative, the tape library can be attached to a separate storage node or even the EBS server itself.

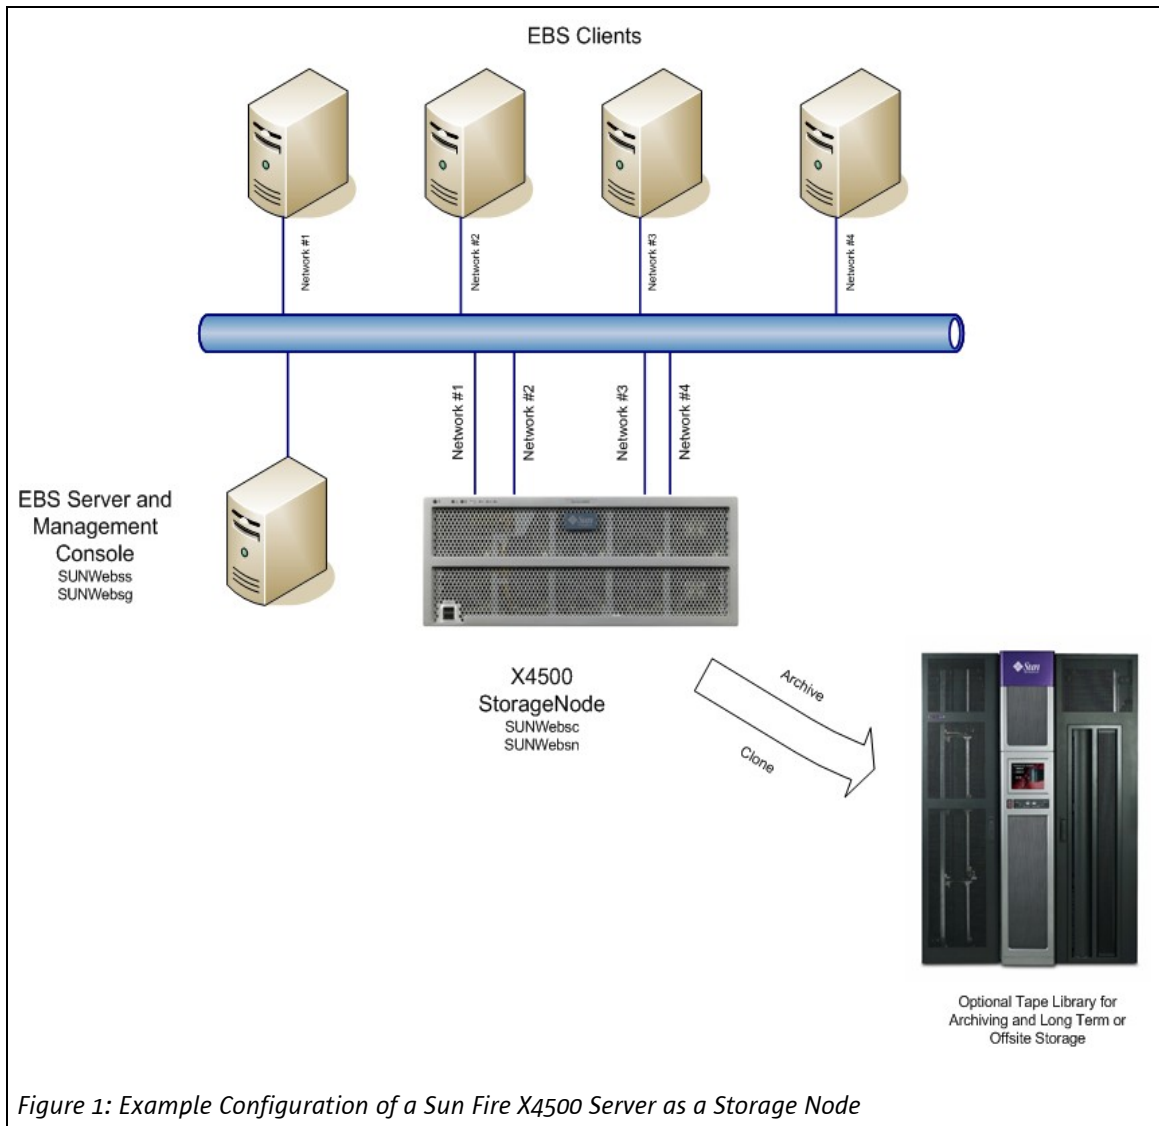


Figure 1: Example Configuration of a Sun Fire X4500 Server as a Storage Node

## Configuration

The minimum operating system for this configuration is the Solaris 10 08/07 OS because this version includes ZFS improvements.

### Configuring the Disks

The Sun Fire X4500 server has 48 500GB SATA drives. Use the Solaris Volume Manager to mirror the boot drives, `c5t0d0` and `c5t4d0`. The remaining drives are available for backup.

## Creating Zpools

The configuration must balance the competing needs of a particular site. Gather the following requirements:

- What is the required performance needed (IOPS/throughput)?
- What is the required data protection level needed (RAIDZ, RAIDZ2, or mirror)?
- What is the required capacity?

Do not configure the entire architecture (encompassing backup clients and network infrastructure) to optimize disk performance at the expense of space or data protection if the architecture is not able to send data at adequate levels. Likewise, do not optimize for space at the expense of performance or data protection.

When configuring network backups, the ability of the Sun Fire X4500 server to receive data over the network is also a factor. For example, a 450MB/s backup speed requirement cannot be met solely by using the four onboard Gigabit Network ports. Add network cards to meet this requirement.

In this guide, the Sun Fire X4500 server is configured as one zpool consisting of 8 x 5 disk RAIDZ vdevs. Additionally, six hot spare disks are configured for reliability. This gives the Sun Fire X4500 server sufficient disk performance to handle 300-350MB/s of backup traffic over the network.

```
#!/bin/bash
zpool create -f ebspool raidz c0t0d0 c1t0d0 c4t0d0 c6t0d0 c7t0d0 \
    raidz    c0t1d0 c1t1d0 c4t1d0 c5t1d0 c7t1d0 \
    raidz    c0t2d0 c1t2d0 c5t2d0 c6t2d0 c7t2d0 \
    raidz    c0t3d0 c1t3d0 c4t3d0 c6t3d0 c7t3d0 \
    raidz    c0t4d0 c1t4d0 c4t4d0 c6t4d0 c7t4d0 \
    raidz    c0t5d0 c4t5d0 c5t5d0 c6t5d0 c7t5d0 \
    raidz    c0t6d0 c1t6d0 c4t6d0 c5t6d0 c6t6d0 \
    raidz    c1t7d0 c4t7d0 c5t7d0 c6t7d0 c7t7d0 \
    spare c6t1d0 c5t3d0 c4t2d0 c1t5d0 c7t6d0 c0t7d0
```

To see the result, use the `zpool list` command:

The `SIZE` column shows the actual physical space in the pool. This might differ from the actual space that the data occupies because various RAIDZ methods and configurations are not taken into account. See the `zpool(1M)` man page for more details.

```
# zpool list
NAME          SIZE    USED  AVAIL    CAP  HEALTH   ALTROOT
ebspool      18.1T   161K   18.1T    0%  ONLINE   -
```

## Configuring ZFS

As a storage node, the Sun Fire X4500 server is configured as the EBS device type, `adv_file`. This device type, also called an Advanced File Type Device (AFTD), indicates that the storage device uses a volume manager to extend its disk space dynamically. For EBS, the `adv_file` device type allows concurrent backup and recover sessions to the same device. Also, while an `adv_file` device is being used for backups, it can also be used to perform a clone or stage operation.

Each `adv_file` device is represented by one file system and has an `nsrmmmd` process that writes the backup data to the file system. Adding more `adv_file` devices can affect system performance during peak load times while not improving overall throughput by any significant amount. Testing must be performed to determine the correct mix of performance and flexibility needed.

The number of file systems needed is a case-by-case decision depending on the following factors:

- Number of backup/recover streams needed
- Number of clone or stage operations needed
- Number of tape drives configured for clone or stage operations

In this example, four backup file systems are created (`/backup1`, `/backup2`, `/backup3`, `/backup4`). This allows four `adv_file` devices to perform four simultaneous cloning or staging operations either to other `adv_file` devices configured for clone/stage pools or to tape devices for long-term storage. In addition to the backup file systems, a `/default` file system is created to handle any default pool needs, such as user backups, and a `/index` file system is created to handle client index backups.

```
#!/bin/bash
for fs in 1 2 3 4
do
zfs create -o mountpoint=/backup$fs ebspool/backup$fs
done
zfs create -o mountpoint=/default ebspool/default
zfs create -o mountpoint=/index ebspool/index
```

To see the results, use the `zfs list` command:

```
# zfs list
NAME                USED    AVAIL  REFER  MOUNTPOINT
ebspool              509K   14.3T  39.1K  /ebspool
ebspool/backup1     39.1K   4.3T   39.1K  /backup1
ebspool/backup2     39.1K   14.3T  39.1K  /backup2
ebspool/backup3     39.1K   14.3T  39.1K  /backup3
ebspool/backup4     39.1K   14.3T  39.1K  /backup4
ebspool/default     39.1K   14.3T  39.1K  /default
ebspool/index       39.1K   14.3T  39.1K  /index
```

## Installing EBS Options

The EBS version is 7.4.

The Sun Fire X4500 server acts as a storage node. The architecture, as shown in Figure 1, includes a separate EBS Master Server and perhaps a separate Management Console server.

Two EBS packages must be installed on the Sun Fire X4500 server to allow the server to be configured as a storage node:

- `SUNWebssc` – Sun StorageTek Enterprise Backup Client package
- `SUNWebssn` – Sun StorageTek Enterprise Backup Storage Node package

Even though the Sun Fire X4500 system might not be used as an EBS client, the client package is a prerequisite for installing the storage node package. These packages are separate, licensed EBS options. Ensure you have adequate licenses to configure the Sun Fire X4500 server.

During the `SUNWebssc` installation, you must provide the following information:

- The location of the EBS log and data files (This is not the backup data.)
- The name of the EBS Master Server

The packages must be installed from the `solaris_amd64` subdirectory of the EBS installation media. Issue the following command:

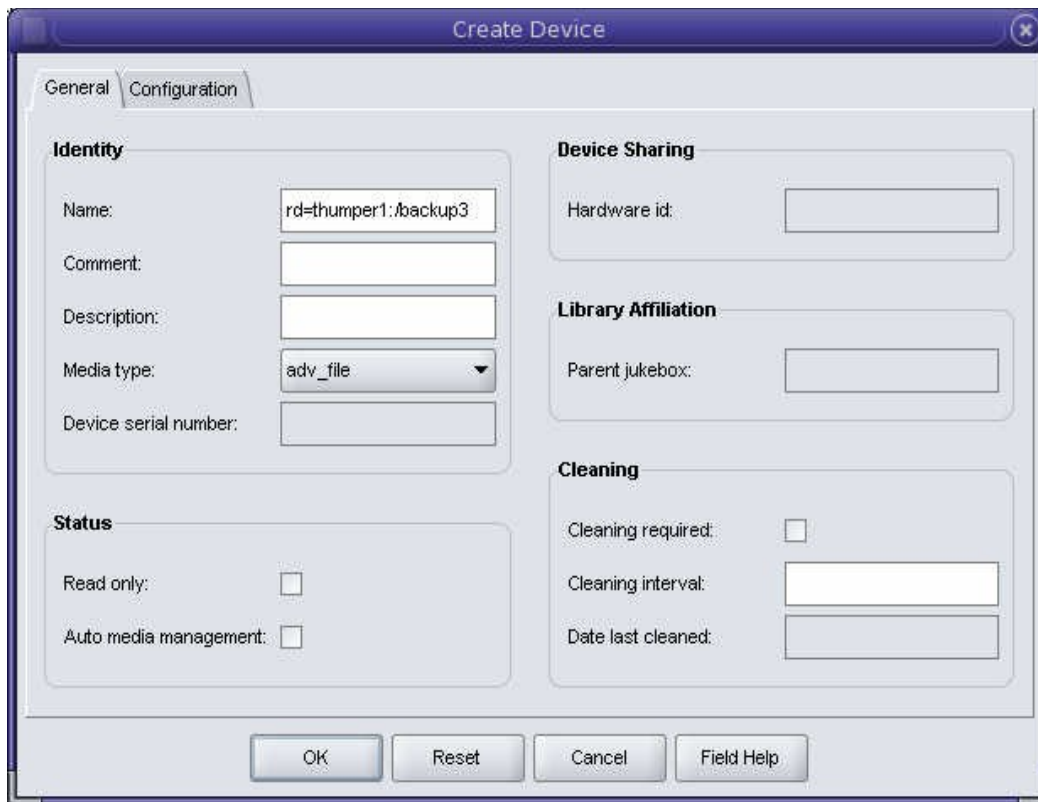
```
# pkgadd -d . SUNWebssc SUNWebssn
```

### Configuring the EBS Storage Node

1. In the server's Administration interface, click the Devices button.
2. In the navigation tree, right-click Storage Nodes and select New. The Create NSR Storage Node window appears at the General tab.
3. In the Identity area, in the Name field, enter the fully qualified domain name or short name of the Sun Fire X4500 server.
4. In the Type of Storage area, verify that the Node field is the default, `scsi`.
5. Click OK. The new storage node appears in the navigation tree.

### Configuring Each ZFS File System as a Storage Device

1. In the server's Administration interface, click the Devices button.
2. Right-click Devices in the navigation tree, and select New. The Create Device window appears with the General tab selected.



3. In the Identity area's Name field, replace the default name with the full path of the ZFS file system. Because the Sun Fire X4500 server is a remote storage node, the path name of the device includes `rd=` with the name of the Sun Fire X4500 server. In the sample configuration, the remote storage node is `thumper1` so the path is `rd=thumper1:/backup3`.
4. In the Media Type list, choose the `adv_file` type.
5. In the Status area, verify that Auto Media Management option is disabled.
6. On the Configuration tab, make changes to the sessions as your site requires.
7. Click OK to complete the definition of the new device.
8. Repeat the procedure for the remaining ZFS file systems.

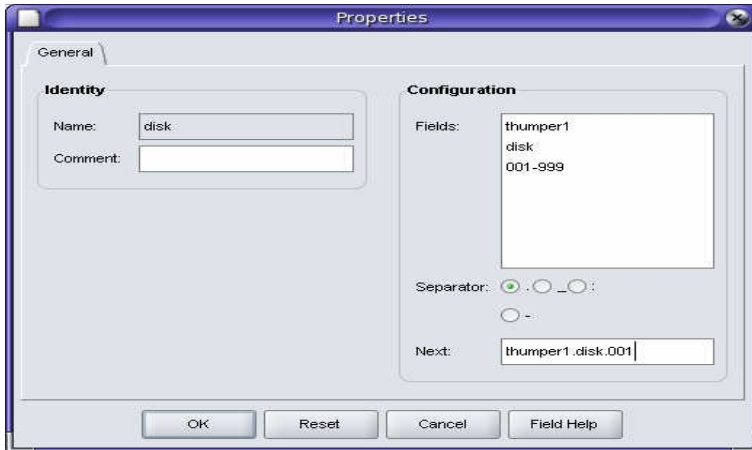
## Creating a Media Pool

The Sun StorageTek EBS server creates a unique label for each volume by applying a label template to sort, store, and track data on volumes. To direct backups to the storage node, designate that data from a specific media pool whose volumes share label components is backed up to the storage node.

To create the media pool, first create the label template for disk storage and then create the media pool with the new label.

To create a label template:

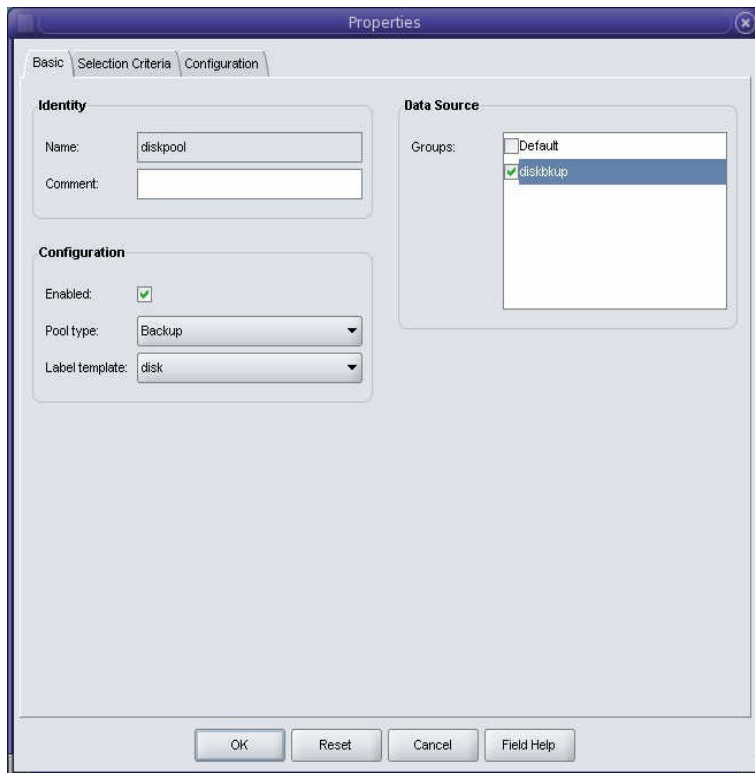
1. In the Administration window, click the Media button.
2. In the expanded left pane, select Label Templates.
3. From the File menu, select New.



4. Enter the components for the label template.
  - Name: The name of the new label template. This example uses the name disk.
  - Comment: Describe the purpose of media pools with this label.
  - Fields: A list of label components.
  - Separator: The character to be inserted between label components.
  - Next: (Optional) Enter the next label to be generated by the template.
5. Click OK.

To create the media pool:

1. In the Administration window, click the Media button.
2. In the left pane, select Media Pools.
3. From the File menu, select New. The Properties dialog is displayed with the Basic tab selected.



4. In the Name field, type a name for the media pool that also identifies its label template. This example uses `diskpool`.
5. In the Comment field, type a description of the purpose for the media pool.
6. Select the Enabled field.
7. For the Pool Type field, select Backup.
8. In the Label Template field, select the new label template, `disk`.
9. In the Data Source field, select the backup groups that are eligible to store data in this media pool.
10. Select the Selection Criteria tab to further restrict which data can back up to this media pool.
11. Complete the other fields as necessary, and click OK.

To label and mount a volume in the new storage node:

1. In the Administration window, click the Devices button.
2. Select Devices in the navigation tree. The Devices Detail table appears.
3. Right-click the Sun Fire X4500 server and select Label. The Label window appears.



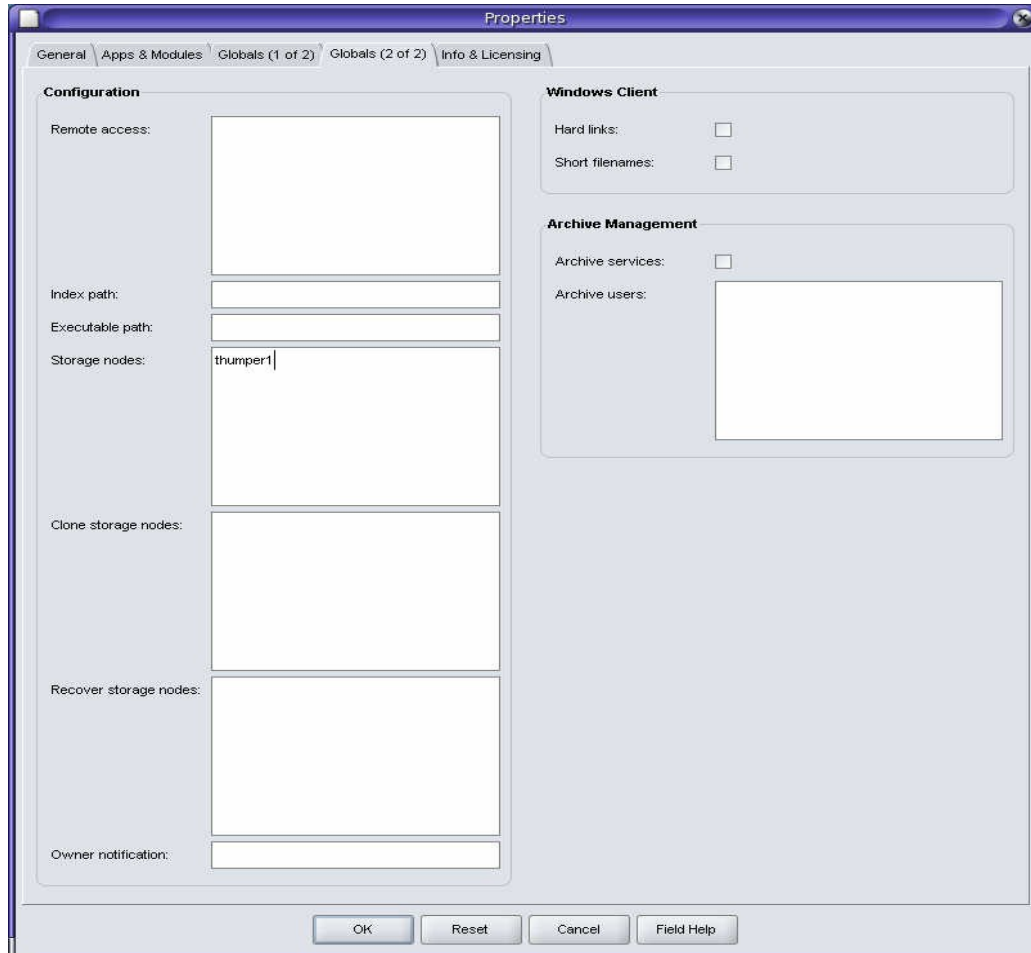
4. From the Pools list, select the new media pool, `diskpool`.
5. Accept the default name associated with volumes in the `diskpool` media pool.
6. Click OK.  
The Sun StorageTek EBS server automatically labels the volume, and then mounts the volume into the device. In addition, whenever the EBS server labels an AFTD, such as the Sun Fire X4500 server, it also creates a virtual device, with `_AF_readonly` appended. The Device navigation tree now includes:
  - The device created when adding the `adv_file` device.
  - The device created with `/_AF_readonly` appended to the name and mounted with a volume with the RO suffix. This device is used for concurrent, read-only operations.
7. Label and mount the remaining devices, such as `/default` with the correct media pools.

At this point, you have designated that certain backup groups be included in a media pool and that the media pool uses the storage node for its backups. The remaining task is to specify the EBS clients that use the storage node for backups.

## Configuring Client Backups

To configure the clients to use this storage node for backups:

1. In the server's Administration interface, click the Configuration button.
2. Select Clients, right-click the appropriate client, and select Properties.



3. On the General tab, verify that the client is a member of a backup group that is included in the diskpool media pool.
4. On the Globals (2 of 2) tab, edit the Storage Nodes field by entering the host name of the Sun Fire X4500 server.
5. Click OK.

## Frequently Asked Questions

Q1: After creating my storage node and sending all my backups to media pools on it, my server bootstrap backups fail.

A1: The EBS server bootstrap backups must be stored on local storage. Create a local device and media pool and direct bootstrap backups to it. See the Sun StorageTek Enterprise Backup Software 7.4 Administration Guide for details.

Q2: What is the best zpool configuration to create to get the maximum performance?

A2: A balance of performance, data protection, and capacity must be decided based on customer requirements. Further testing is needed to quantify optimal configurations based on customer needs.

Q3: I don't want to create four separate backup subnets to utilize my four GigE ports.

A3: Use the `dladm(1M)` command to aggregate the ports under a single IP address. See the `dladm(1M)` man page and the Solaris 10 System Administrator Collection – System Administration Guide: IP Services manual for more details. Additionally, a 10GigaBit Ethernet card can be used for a backup interface.

Q4: My Sun Fire X4500 server is filling up! Help!

A4: See the Sun StorageTek Enterprise Backup Software 7.4 Administration Guide for how to set up recycling policies and staging to long-term storage.

## For More Information

Here are additional resources:

- Sun storage download page:  
<http://www.sun.com/download/index.jsp?cat=Hardware%20Drivers&tab=3&subcat=Storage>
- These and other Sun training courses at <http://www.sun.com/training/>:
  - Sun Fire X4500 Server Administration (WET-5060)
  - Sun StorEdge Enterprise Backup 7.x Administration (ES-260)
  - Designing Storage Area Networks (NWS-4940)
- Storage and Information Lifecycle Management course catalog:  
<http://www.sun.com/training/catalog/storage/index.xml>
- Support:
  - Register your Sun gear: <https://inventory.sun.com/inventory/>
  - Services: <http://www.sun.com/service/index.jsp>
  - SunSolve Online: <http://sunsolve.sun.com>
- Discussions such as the Storage Forums:  
<http://forum.java.sun.com/category.jspa?categoryID=66>
- The following documents at <http://docs.sun.com>:
  - Sun StorageTek Enterprise Backup Software Release 7.4 Documentation:  
<http://docs.sun.com/app/docs/coll/ebs7.4?l=en>
  - Sun Fire X4500 Document Collection:  
<http://docs.sun.com/app/docs/prod/f4bbfa5f-e6e5-11da-ac3d-080020a9ed93>
  - Solaris 10 System Administrator Collection:
    - Solaris ZFS Administration Guide:  
<http://docs.sun.com/app/docs/doc/819-5461>
    - System Administration Guide: IP Services:  
<http://docs.sun.com/app/docs/doc/816-4554>

- Sun wikis at <http://wikis.sun.com/dashboard.action>, for example:
  - Sun Storage Administration Wiki: <http://wikis.sun.com/display/StorageAdmin/Home>
  - Storage Developers' Wiki: <http://wikis.sun.com/display/StorageDev/Home>
  - Storage System Patterns Wiki: <http://wikis.sun.com/display/StorageSystemPatterns/Storage+System+Patterns>
  - Storage BluePrints Wiki: <http://wikis.sun.com/display/BluePrints/Storage+BluePrints>
  - BigAdmin Storage Tech Tips Wiki: <http://wikis.sun.com/display/BigAdmin/Storage+Tech+Tips>
- Related sites:
  - *ZFS Best Practices Guide* on the Solaris Internals web site: [http://www.solarisinternals.com/wiki/index.php/ZFS\\_Best\\_Practices\\_Guide](http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide)
  - BigAdmin Storage Resource Collection: <http://www.sun.com/bigadmin/collections/storage.html>
  - Sun Storage web site: <http://www.sun.com/storagetek/index.jsp>
  - Storage Stop Blog: <http://blogs.sun.com/storage>
- Events of interest to users of Sun products:
  - Worldwide developer events: <http://developers.sun.com/events/>
  - Current events: <http://www.sun.com/events/index.jsp>

## Licensing Information

Unless otherwise specified, the use of this software is authorized pursuant to the terms of the license found at [http://www.sun.com/bigadmin/common/berkeley\\_license.html](http://www.sun.com/bigadmin/common/berkeley_license.html).