

Sun Fire X4500 as a Symantec Netbackup 6.5 Media Server

A Quick Start Guide
Sun Microsystems
December 11, 2007

2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 USA

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD appliances, licensed from the University of California.

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

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun's Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-1987), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). 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 HELD TO BE LEGALLY INVALID.

Table of Contents

Introduction.....	3
Configuration	4
Overview.....	4
Solaris Configuration.....	5
Zpool Configuration.....	5
ZFS Configuration.....	6
Netbackup Configuration.....	7
Frequently Asked Questions.....	16
Summary.....	18
References.....	19

Introduction


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 Operating System with the data integrity and simplified administration of ZFS, the X4500 becomes an ideal candidate for streamlining and improving backup/restore operations.

This paper discusses how to quickly configure the X4500 as a Media Server for Symantec Veritas NetBackup.

Configuration

Overview

The below diagram shows an example NBU configuration with the X4500 acting as a Media Server. The default NIC configuration on the X4500 allows up to four physical GigaBit Ethernet connections to be attached. The four connections can be utilized to back up four separate subnets, or the links can be aggregated together using `dladm(1M)` for improved throughput to a single subnet.



The X4500 can also be connected to an optional Tape Library (using FibreChannel or SCSI connections) for Image Duplication or Archiving purposes. The X4500 can also be used as the primary data cache for Staging operations, with a Tape Library be utilized for longer term storage and data moved to tape after certain filesystem utilization or time-based thresholds are reached. It is also possible that the Tape Library is attached to a completely separate Media Server.

Solaris Configuration

The X4500 server used while testing for this paper was installed with Solaris 10 8/07. This is the recommended version to use with Netbackup as various ZFS improvements were introduced in this build.

The X4500 has 48 500GB SATA drives. The two boot drives (c5t0d0 and c5t4d0) are mirrored using Solaris Volume Manager. The rest of the drives are available for backup data.

Zpool Configuration

Zpool configuration is a decision that needs to be made based on a number of factors.

- What is the required performance needed? (IOPS/Throughput)
- What is the required data protection level needed? (RAIDZ/RAIDZ2/Mirror)
- What is the required capacity?

After gathering and analyzing the requirements, a configuration that balances the needs of those requirements can be chosen.

When configuring the system for network backups, the ability of the X4500 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 4 onboard Gigabit Network ports. Additional network cards would be needed to meet this requirement.

Care should be taken that a system is not over configured for disk performance, at the expense of space or data protection, when the entire system (encompassing backup clients and network infrastructure) is not able to send data at adequate levels. Likewise, a system should not be over configured for space at the expense of performance or data protection.

For the purposes of this Quick Start Guide, the X4500 was configured one zpool consisting of 8 x 5 disk RAIDZ vdevs. Additionally, 6 hotspares were configured for reliability. This gives the X4500 sufficient disk performance to handle 300-350MB/sec of backup traffic over the network.

Example Zpool creation script:

```
#!/bin/bash
zpool create -f nbupool raidz cot0d0 c1t0d0 c4t0d0 c6t0d0 c7t0d0 \
  raidz cot1d0 c1t1d0 c4t1d0 c5t1d0 c7t1d0 \
  raidz cot2d0 c1t2d0 c5t2d0 c6t2d0 c7t2d0 \
  raidz cot3d0 c1t3d0 c4t3d0 c6t3d0 c7t3d0 \
  raidz cot4d0 c1t4d0 c4t4d0 c6t4d0 c7t4d0 \
  raidz cot5d0 c4t5d0 c5t5d0 c6t5d0 c7t5d0 \
  raidz cot6d0 c1t6d0 c4t6d0 c5t6d0 c6t6d0 \
  raidz c1t7d0 c4t7d0 c5t7d0 c6t7d0 c7t7d0 \
  spare c6t1d0 c5t3d0 c4t2d0 c1t5d0 c7t6d0 cot7d0
```

```
root@thumper1 # zpool list
NAME          SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT
nbupool       18.1T 58.6G 18.1T  0%  ONLINE  -
```

The size listed in the above command is the actual physical space in the pool. This may differ from the actual space data may occupy as various RAIDZ methods and configurations are not taken into account. See the `zpool(1M)` man page for more details.

ZFS Configuration

In this example, the X4500 Media Server is configured to write to Basic Disk Storage Units. Each Storage Unit is created from a ZFS filesystem mountpoint. The number of filesystems needed will be a implementation by implementation decision but a few factors should be considered.

- Number of Backup/Recover streams needed
- Number of duplication or stage operations needed
- Number of Tape Drives configured for duplication or stage operations

In this example, 4 backup filesystems are created (`/backup1`, `/backup2`, `/backup3`, `/backup4`). Testing has shown that creating a single ZFS filesystem for all backup streams is a bottleneck that limits the throughput to 200-250MB/sec. Using multiple filesystems in parallel eliminates that bottleneck.

Example ZFS filesystem creation script:

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

```
root@thumper1 # zfs list
NAME          USED  AVAIL  REFER  MOUNTPOINT
nbupool       46.8G 14.2T 39.1K  /nbupool
nbupool/backup1 14.8G 14.2T 14.8G  /backup1
nbupool/backup2 10.8G 14.2T 10.8G  /backup2
nbupool/backup3  5.38G 14.2T  5.38G  /backup3
nbupool/backup4 15.8G 14.2T 15.8G  /backup4
```

Netbackup Configuration

The X4500 will be configured as a Media Server. In this example, it is assumed that there is a separate Master NBU server.

The “Netbackup and Media Manager” package must be installed on the X4500. It also will require the appropriate server licensing.

It is recommended that the Master server's bp.conf file be updated with the new Media server information before installing the software on the media server. An example is shown below.

The supported installation method is using the install script that comes with the Netbackup Installation media.

```
root@thumper1 # /nfs/nbu-new/DVD2/install

Symantec Installation Script
Copyright 1993 - 2007 Symantec Corporation, All Rights Reserved.
  Installing NetBackup Server Software
NOTE: To install only NetBackup Client software locally on this machine
      or to load additional UNIX client software on this server, insert
      the NetBackup UNIX Clients cdrom.
Do you wish to continue? [y,n] (y) y
NetBackup installs to the running root environment, by default.
Are you installing to the running root environment? [y,n] (y) y
Processing package instance <SYMCnetbp> from </nfs/nbu-new/DVD2/solaris_x86>
```

(... rest of the installation output omitted for brevity...)

After the installation process prompts for the license key and the key is properly entered, the following questions need to be answered

```
Would you like to use "thumper1" as the configured
name of the NetBackup server? [y,n] (y) y
Is thumper1 the master server? [y,n] (y) n
What is the fully qualified name of the master server? sol-230
```

The previously installed master server name should be entered. In this example, the master server name was “sol-230”.

After the package is installed on the X4500. The Master Server EMM database must be updated with the new Media Server information.

From the Master Server command line:

```
# /usr/openv/netbackup/bin/admincmd/nbemmcmd -addhost -machinename thumper1 -machinetype media -masterserver sol-230 \
-netbackupversion 6.5 -operatingsystem solaris
NBEMMCMD, Version:6.5
Command completed successfully.
```

Where “machinename” is the name of the X4500 and “masterserver” is the name of the Master Server.

The results of this command can be verified

```
# nbemmcmd -listhost
NBEMMCMD, Version:6.5
The following hosts were found:
media      thumper1
master     sol-230
server     sol-230
Command completed successfully.
```

The bp.conf files on both Master and Media servers and any clients should be adjusted for the new Media Server.

On the Master Server

```
SERVER = sol-230
SERVER = thumper1
CLIENT_NAME = sol-230
EMMSERVER = sol-230
VXDBMS_NB_DATA = /usr/opensv/db/data
```

On the Media Server

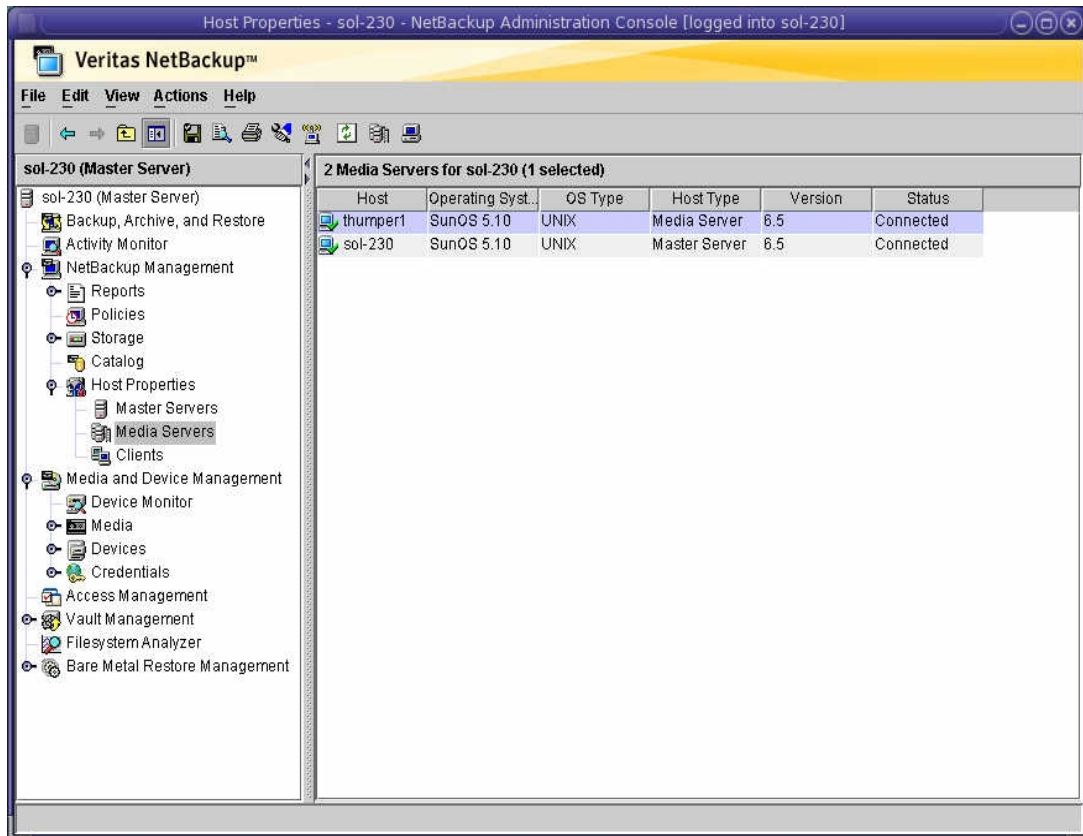
```
SERVER = sol-230
SERVER = thumper1
CLIENT_NAME = thumper1
EMMSERVER = sol-230
```

Also, any client needing access to the Media Server resources should add “SERVER=<x4500 name>” into it's bp.conf file as well.

Restart the NetBackup processes on the Master and Media servers to have the changes recognized.

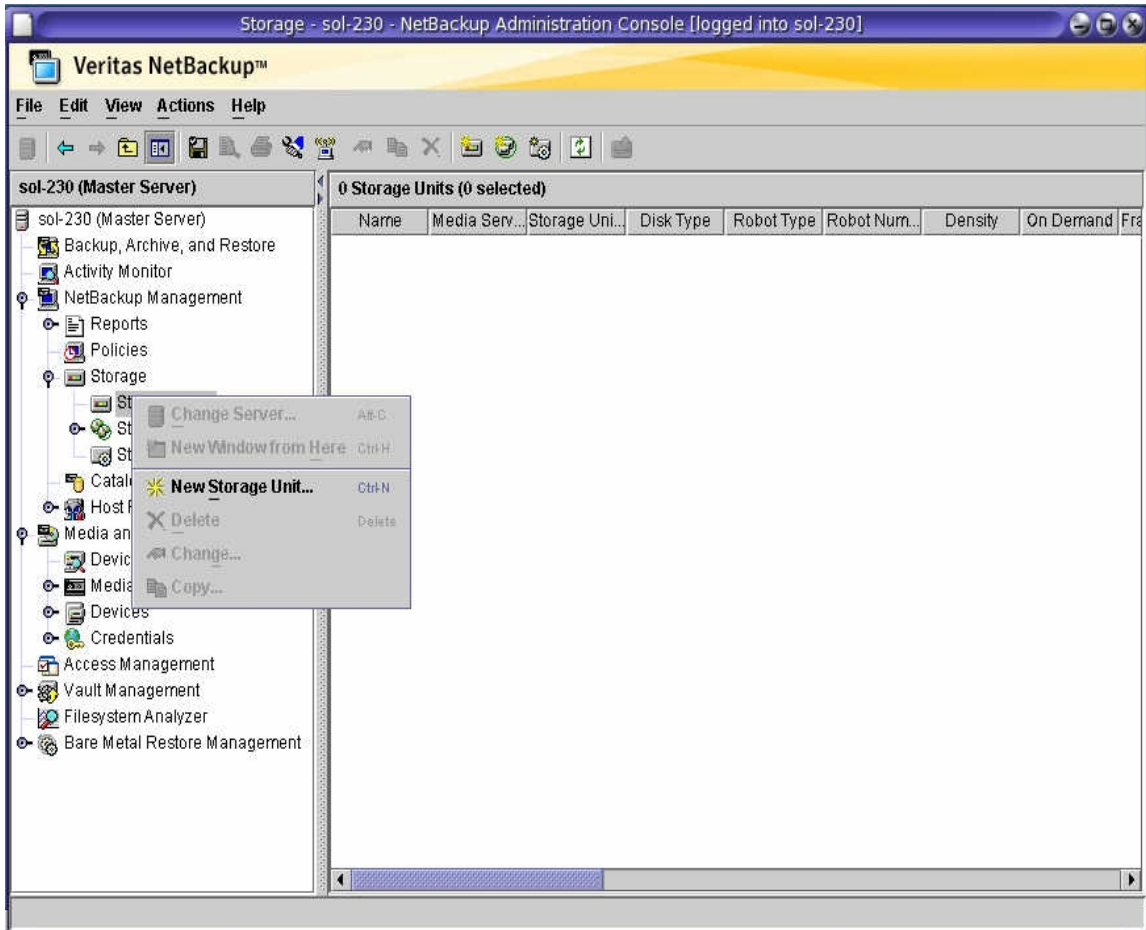
Use the NetBackup GUI for the remaining configuration

Verify that the Media Server is recognized via “Host Properties -> Media Servers” Click on the X4500 hostname to ensure the Media Server is “Connected” The Master Server will already be listed.



To create a Storage Unit using the ZFS filesystems previously configured, follow the below procedures

“Storage -> Storage Units” Right-Click to add new Storage Unit



Storage Unit Configuration

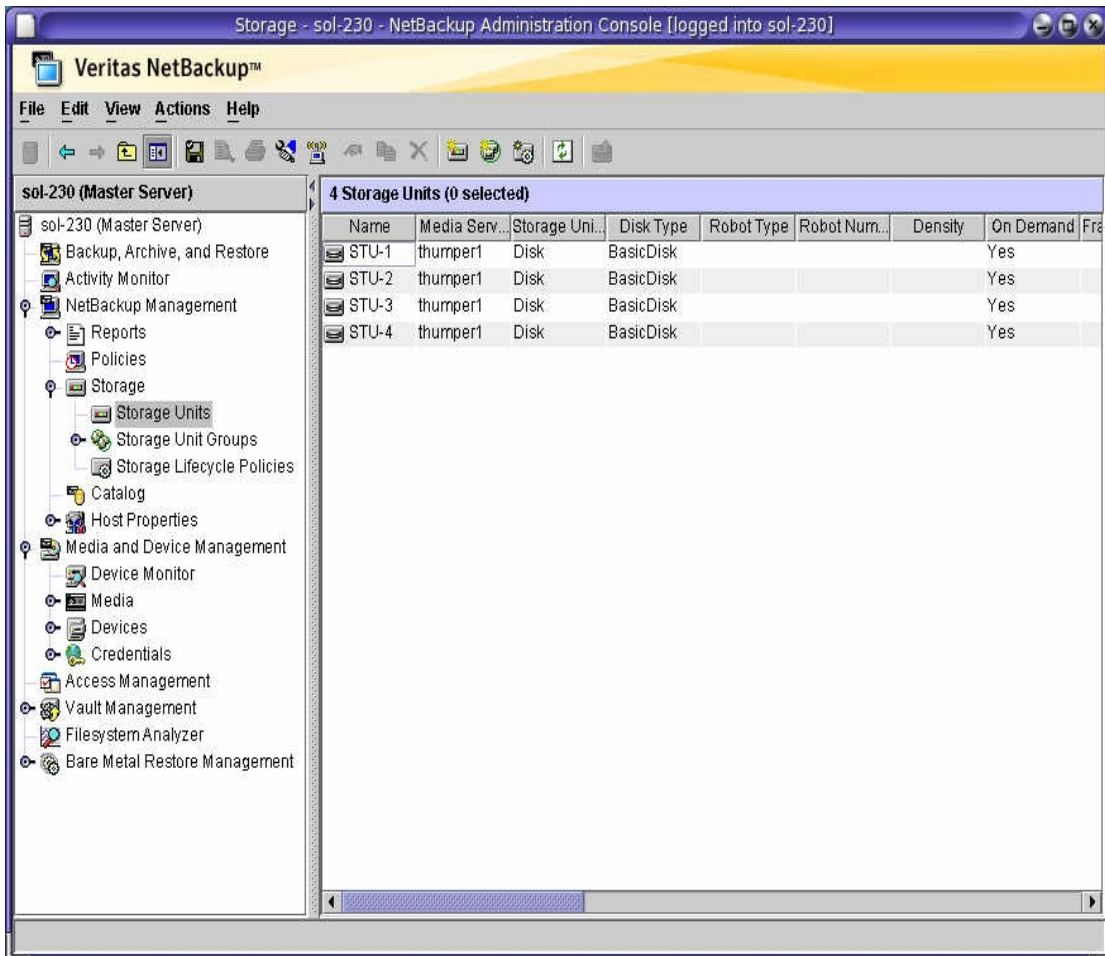
- Use Basic Disk as the Type
- Ensure the Media Server is the X4500
- Set the Absolute pathname to one of the ZFS filesystems

The image shows a Windows-style dialog box titled "New Storage Unit". The dialog is configured with the following settings:

- Storage unit name:** STU-1
- Storage unit type:** Disk (selected in a dropdown menu). A checkbox for "On demand only" is checked.
- Disk type:** BasicDisk (selected in a dropdown menu).
- Properties and Server Selection:**
 - Media server:** thumper1 (selected in a dropdown menu).
 - Absolute pathname to directory:** /backup1 (text input field). A "Properties" button is located to the right of this field.
 - This directory can exist on the root file system or system disk.
- Maximum concurrent jobs:** 4 (spin box).
- Reduce fragment size to: 524288 (spin box) Megabytes.
- High water mark:** 98 (spin box) %.
- Low water mark:** 80 (spin box) %.
- Enable Temporary Staging Area. Copy the data to its final destination according to its staging schedule. A "Staging Schedule..." button is located below this checkbox.

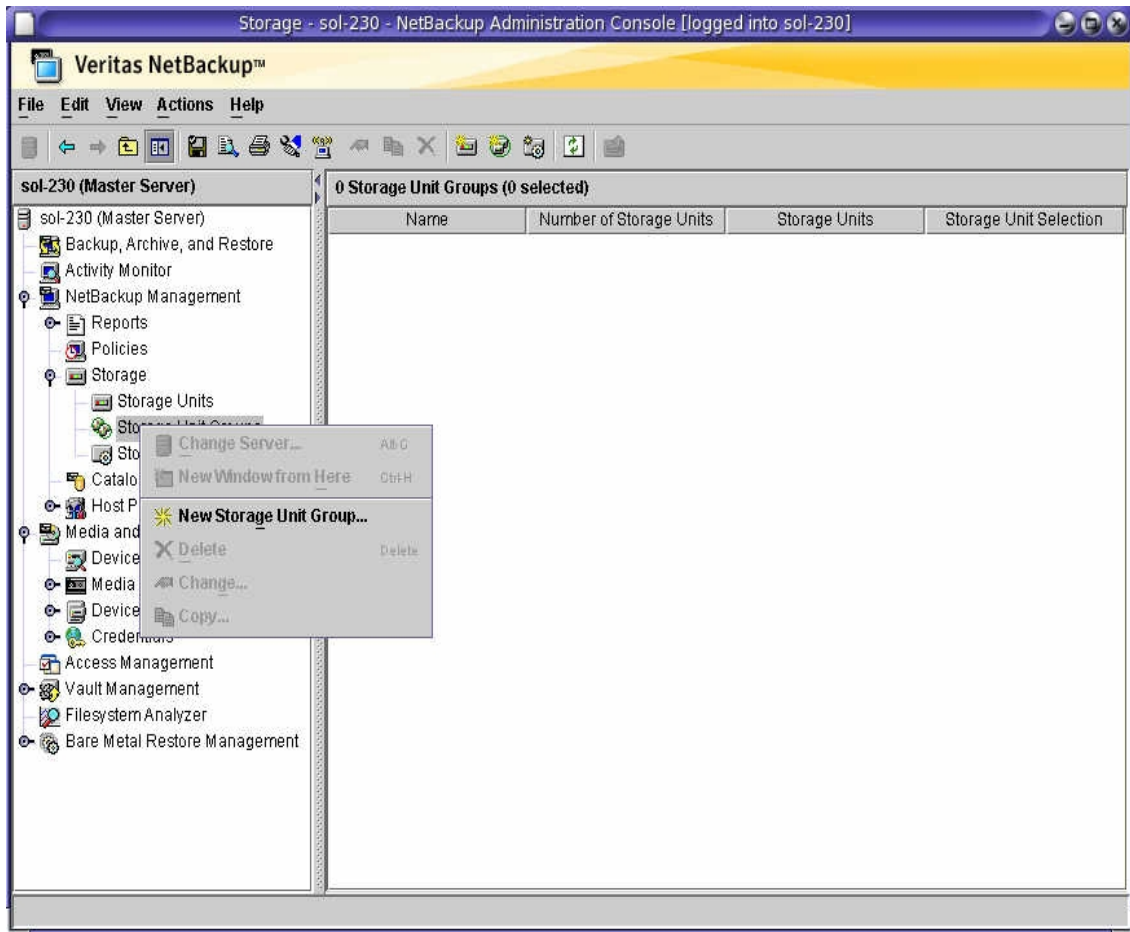
At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

Follow the same procedures for the rest of the ZFS filesystems you have configured.



The individual Storage Units could then be used as backup destinations, however, configuring Storage Groups can prove beneficial.

Add a New Storage Group



Add the Storage Units that have been created into a Storage Group. The Storage Unit Selection will be “Round-Robin”. This will ensure that backup streams are allocated across all units.

New Storage Unit Group

Storage unit group name:
SG-1

Change priority:
Move Up Move Down

Storage units in the group:

- STU-1
- STU-2
- STU-3
- STU-4

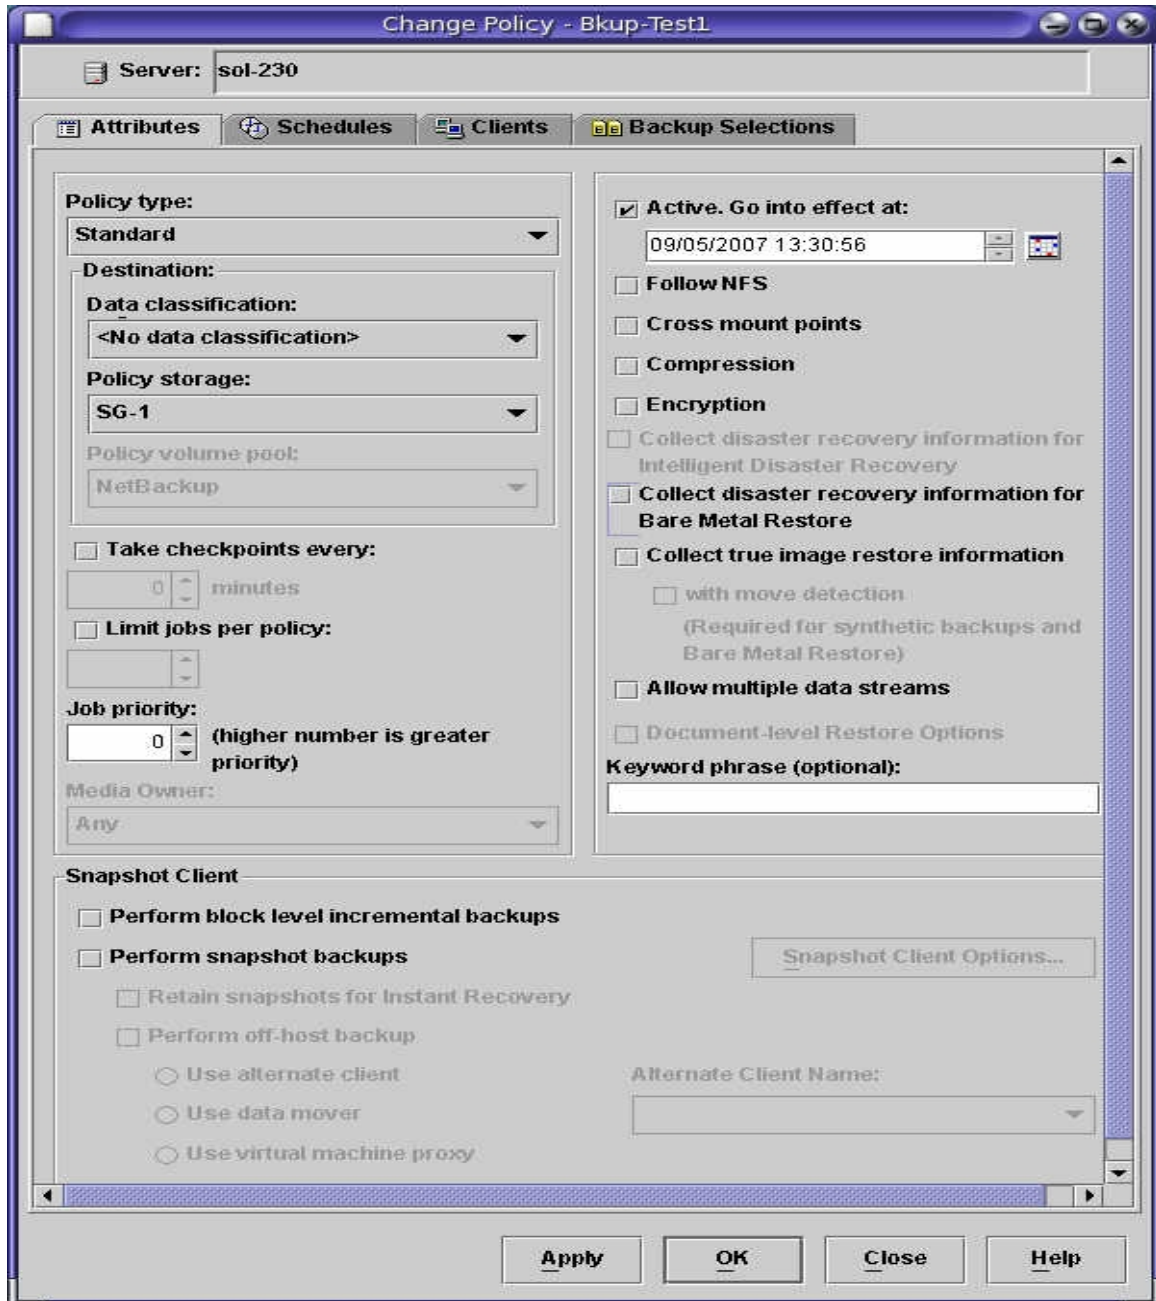
Storage units not in the group:
Add Remove

Storage unit selection

- Prioritized (Choose the first storage unit in the list that is not busy, down, or out of media.)
- Failover (Choose the first storage unit in the list that is not down or out of media.)
- Round Robin (Choose the least recently selected storage unit in the list.)
- Load Balance

OK Cancel Help

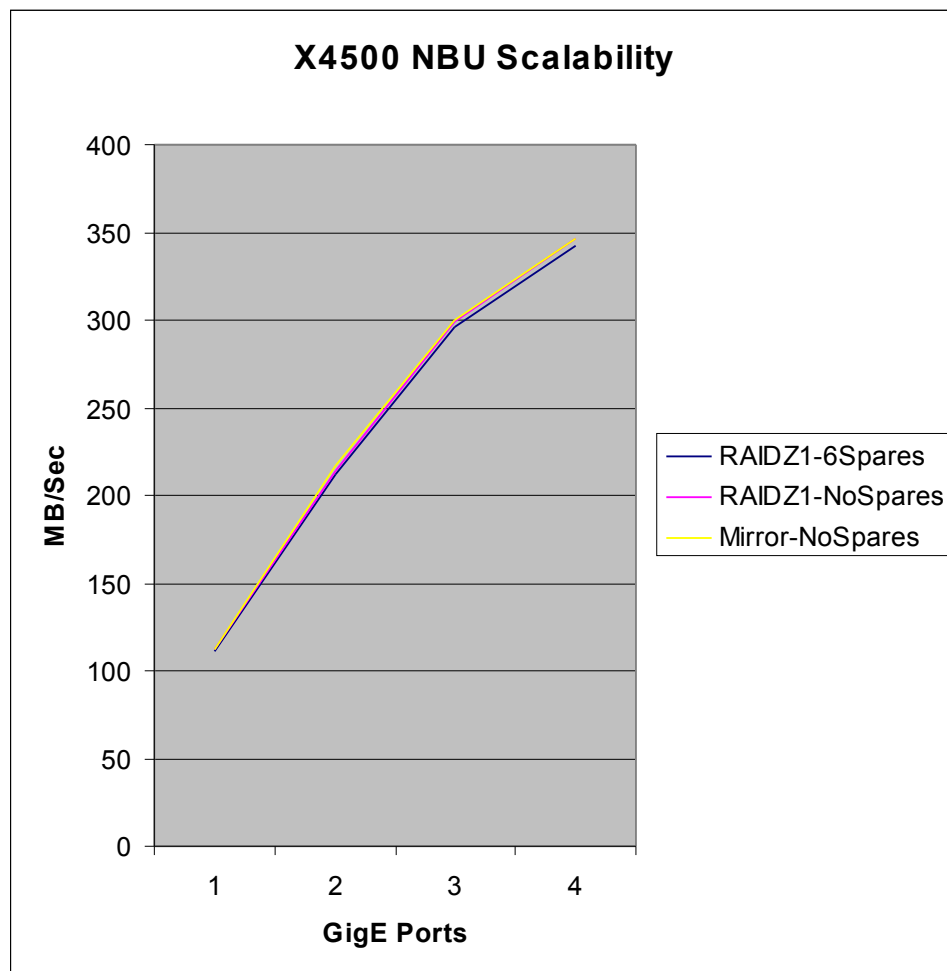
Once the Storage Group is configured, the Policy Attributes for the backup jobs can be configured to use the Storage Group as the “Policy Storage”.



Frequently Asked Questions

Q1: How does the X4500 scale as the workload increases?

A1: Testing was done to measure scalability by saturating each onboard GigE port in turn. Three Zpool configurations were tested. The RAIDZ1-6Spares (14.3TB of usable space) configuration is the configuration listed in the example earlier in this guide. The RAIDZ1-NoSpares (16.9TB of usable space) configuration is based on the default Zpool configuration pre-installed at the factory. The Mirror-NoSpares (10.3TB of usable space) configuration is a Single Zpool consisting of 22 mirrored pairs. Sufficient client load was generated to ensure each GigE interface could be fully saturated at 1 Gigabit per sec. The test data consisted of 32 GB of large files on each client. As shown in the graph below, the X4500 can deliver close to 350 MB/sec using the 4 onboard GigE ports. Also, the testing shows that at these data rates Zpool configuration is not a driving factor for performance. The network/CPU load becomes the bottleneck before the Zpool.



Q2: Has the configuration been tested with 10GB Ethernet?

A2: Yes. A Proof-of-Concept was done for a large customer and 10GB Ethernet was part of the testing. Using a Neterion Xframe II and the xge driver, 8 clients capable of delivering 100MB/sec each were configured to backup to a X4500. Jumbo Frames were enabled on all components. Throughput rates of 500-550 MB/sec were achieved using the Zpool/ZFS configurations mentioned above. Image duplication between two X4500 Media Servers was measured at 350 MB/sec.

Q3. How can I achieve even higher aggregate backup speeds?

A4. By configuring additional X4500 Media Servers and placing all the Disk Storage Units from all Media Servers into a single Storage Unit Group, Netbackup can effectively round-robin across multiple X4500s for even higher aggregate throughput rates.

Summary

This guide was a Quick Start introduction into configuring the Sun Fire X4500 as a disk-cache Media Server for Veritas Netbackup 6.5. It gave examples on configuring the zpool and zfs filesystems on Solaris and how to configuration the Netbackup application to utilize the X4500. Using this guide, backup throughput rates of over 1TB per hour are achievable.

References

- i. Veritas Netbackup 6.5 Documentation
 1. http://seer.entsupport.symantec.com/search_forms/doc_list.asp?DocType=Comp&LSUB=NBUESVRSUB11682&MSUB=NBUESVRSUB11686&SearchArea=ALL&el=%2C+1&ddProduct=NBUESVR&pid=&width=&SearchTerm=&subject=NBUESVRSUB11686&ProductVersion=NBUESVRPVER24910&language=English
- ii. Sun Fire X4500 Document Collection
 1. http://www.sun.com/products-n-solutions/hardware/docs/Servers/x64_servers/x4500/index.html
- iii. Solaris 10 System Administrator Collection
 1. ZFS Admin Guide - <http://docs.sun.com/app/docs/doc/819-5461>
 2. IP Services Guide - <http://docs.sun.com/app/docs/doc/816-4554>
- iv. Solaris Internals Website
 1. ZFS Best Practices Guide - http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide