



PART 2:
CONFIGURE & USE ZFS-HOSTED
iSCSI GUEST STORAGE
**WITHIN THE SUN
xVM HYPERVISOR**

> Technical Brief

About This Technical Brief

This technical brief is intended for users familiar with the Sun xVM hypervisor who have used the “Install Sun xVM Hypervisor and Use it to Configure Domains” technical brief to install domains. This brief is Part 2 of the “Configure and Use ZFS-hosted iSCSI Guest Storage within the Sun xVM Hypervisor” series. In Part 1, we completed the following:

1. Configured NFS and iSCSI via ZFS on a Solaris Community Edition-based server
2. Captured the guest data & transferred to the NFS server
3. Configured a Sun xVM hypervisor server to access both NFS and iSCSI
4. Extracted the guest data and recreated on iSCSI
5. Defined and started the two guests from the XML descriptor files

In Part 2, we will extend the steps completed in Part 1 to:

1. Configure a second Sun xVM hypervisor server to access NFS and iSCSI
2. Migrate the guests on the first Sun xVM hypervisor server to the second server using ZFS
3. Create an NFS-based HVM guest on the second server
4. Use Live Migration to migrate the NFS-based HVM guest from the second Sun xVM hypervisor server to the first server

This technical brief is based on the Sun xVM hypervisor components contained within the Solaris Express Community Edition (SXCE). For more information on Solaris Express, refer to:

<http://opensolaris.org/os/downloads/>

Contents

Server Environment	Page 1
Configuration Prerequisites and Assumptions	Page 2
NFS & iSCSI Configuration on Sun xVM Hypervisor Ultra 20 Server	Page 2
Migration of Guests from Sun xVM Hypervisor Ultra 40 to Sun xVM Hypervisor Ultra 20 Server Using ZFS	Page 5
Create an NFS-based Windows HVM Guest on the Sun xVM Hypervisor Ultra 20	Page 9
Migration of the NFS-based Windows HVM guest-w2003se-LM-HVM from Sun xVM Hypervisor Ultra 20 to Sun xVM Hypervisor Ultra 40 Server Using Live Migration	Page 13
Summary	Page 16

Sun xVM Hypervisor Technical Brief

Server Environment

The server environment contains the following components:

Sun Ultra 40 M2:

Hardware Configuration: 2 x 2.2Ghz Dual Core AMD CPUs with 8GB RAM

Role: Sun xVM hypervisor server running Solaris Community Edition build 85

Sun Ultra 20 M2:

Hardware Configuration: 1 x 2.6Ghz Dual Core AMD CPU with 8GB RAM

Role: Sun xVM hypervisor server running Solaris Community Edition build 85

Sun Ultra 24:

Hardware Configuration: 1 x 2.4Ghz Quad Core INTEL CPU with 4GB RAM

Role: Sun iSCSI & NFS server running Solaris Community Edition build 85

Note: All three workstations are on a single flat subnet.

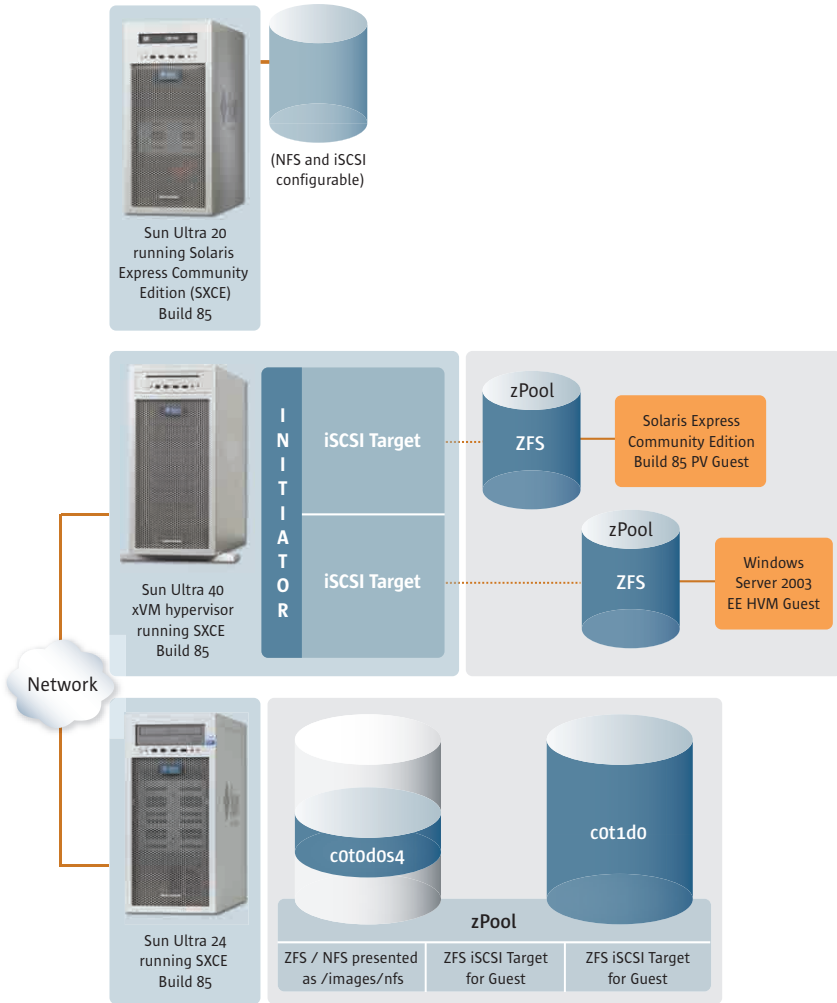


Figure 1—Details this configuration.

Configuration Prerequisites and Assumptions

This technical brief makes the following audience and system assumptions:

- Reader is familiar with the Sun xVM hypervisor and has created domains by following the “Install Sun xVM Hypervisor and Use it to Configure Domains.” technical brief.
- Reader has completed the steps in the “Configure and Use ZFS-hosted iSCSI Guest Storage Within Sun xVM Hypervisor Part 1.” technical brief.
- Reader is comfortable with basic administration of Solaris and Windows.
- Logged in as user root.
- Solaris Community Edition build 85 is installed and the Sun xVM hypervisor components are enabled on two servers.
- On Sun xVM hypervisor Ultra 40 server, two guest domains are configured under ZFS:
 - > Hardware Virtual Machine (HVM) Windows Server 2003
 - > Paravirtualized (PV) Solaris Community Edition Build 85
- Solaris Community Edition build 85 is installed with adequate spare disk space for guest domains and backup files on a single server. In this environment, the Ultra 24 Workstation has a spare slice (c0t0d0s4) on the main boot disk, and a spare internal disk (c0t1d0), which is used for a ZFS zpool (images) to store the guest and backup data.

NFS and iSCSI Configuration on Sun xVM Hypervisor Ultra 20 Server

In this section, the steps needed to configure both NFS and iSCSI on the Sun Ultra 20 server are described. This server will utilize the NFS and iSCSI resources served by the Sun Ultra 24 Server.

1. Check and enable iSCSI configuration:

```
# /usr/sbin/iscsiadm list discovery
Discovery:
    Static: disabled
    Send Targets: disabled
    iSNS: disabled
```

Enable Static Discovery:

```
# /usr/sbin/iscsiadm modify discovery -s enable
# /usr/sbin/iscsiadm list discovery
Discovery:
    Static: enabled
    Send Targets: disabled
    iSNS: disabled
```

2. Run the output from the iscsi.bash script to configure access to the iSCSI LUNs.

Preliminary tasks:

- “Configure and Use ZFS-hosted iSCSI Guest Storage Within Sun xVM Hypervisor Part 1” [Gather the iSCSI target information for initiator configuration]
- “Configure and Use ZFS-hosted iSCSI Guest Storage Within Sun xVM Hypervisor Part 1” [Appendix A iscsi.bash Script} for more information and access to the script

```
# /usr/sbin/iscsiadm add static-config iqn.1986-03.com.sun:02:6fb72f78-9b42-c080-b4fe-efa3d26f4f79,129.146.229.223
# /usr/sbin/iscsiadm add static-config iqn.1986-03.com.sun:02:5dfef158-d101-6d48-cd25-e9580e65f42c,129.146.229.223
```

3. Check that the iSCSI targets are present:

```
# /usr/sbin/iscsiadm list target
Target: iqn.1986-03.com.sun:02:5dfef158-d101-6d48-cd25-e9580e65f42c
    Alias: images/iscsi_luns/guest2
    TPGT: 1
    ISID: 4000002a0000
    Connections: 1
Target: iqn.1986-03.com.sun:02:6fb72f78-9b42-c080-b4fe-efa3d26f4f79
    Alias: images/iscsi_luns/guest1
    TPGT: 1
    ISID: 4000002a0000
    Connections: 1
```

4. Check that the iSCSI targets are presented as disks via format:

```
# /usr/sbin/format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c1t0d0 <DEFAULT cyl 9726 alt 2 hd 255 sec 63>
    /pci@0,0/pci108e,534d@5/disk@0,0
  1. c2t010000144F5857B200002A0047F12CF6d0 <SUN-SOLARIS-1-10.00GB>
    /scsi_vhci/disk@g010000144f5857b200002a0047f12cf6
  2. c2t010000144F5857B200002A0047F12CF7d0 <SUN-SOLARIS-1-10.00GB>
    /scsi_vhci/disk@g010000144f5857b200002a0047f12cf7
Specify disk (enter its number):
```

Note: ^D (CONTROL key and D key) to exit format

5. Mount the NFS share from the Ultra24 with directIO:

```
# /usr/bin/mkdir -p /nfs
# /usr/sbin/mount -o forcedirectio 129.146.229.223:/images/nfs /nfs
```

Note: The Sun Ultra24 server has IP 129.146.229.223. For performance reasons, the *forcedirectio* option is used to disable any buffering in the client.

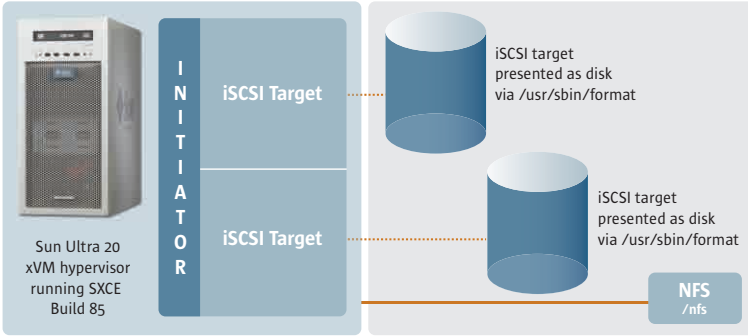


Figure 2—Details this configuration.

Migration of Guests from Sun xVM Hypervisor Ultra 40 to Sun xVM Hypervisor Ultra 20 Server Using ZFS

This section describes the procedure for migrating guests from the Sun Ultra 40 to the Sun Ultra 20 server. To achieve this, we will use the zpool export features of ZFS to export the guest zpool from the Sun Ultra 40 to the Sun Ultra 20 server.

1. On the Sun Ultra 40 server currently hosting the guests, list and shut down both guests:

```
# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 3 guest-nvb85-PV                       blocked
 7 guest-w2003se-HVM                   blocked

# /usr/bin/virsh shutdown guest-nvb85-PV
Domain guest-nvb85-PV is being shutdown

# /usr/bin/virsh shutdown guest-w2003se-HVM
Domain guest-w2003se-HVM is being shutdown

# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 - guest-nvb85-PV                       shut off
 - guest-w2003se-HVM                   shut off
```

2. Export the guest zpools:

```
# /usr/sbin/zpool list
NAME      SIZE    USED    AVAIL    CAP    HEALTH    ALTROOT
guest1    9.94G   6.11G   3.83G    61%    ONLINE    -
guest2    9.94G   2.03G   7.91G    20%    ONLINE    -

# /usr/sbin/zpool export guest1
# /usr/sbin/zpool export guest2
# /usr/sbin/zpool list
no pools available
```

3. On the Sun Ultra 20 server, import the guest ZFS zpools:

```
# /usr/sbin/zfs list
no datasets available
```

List the zpools available for import:

```
# /usr/sbin/zpool import
pool:      guest2
id:        16575101397390218891
state:     ONLINE
action:    The pool can be imported using its name or numeric identifier.
config:

    guest2                                     ONLINE
    c2t010000144F5857B200002A0047F12CF7d0    ONLINE

pool:      guest1
id:        3560710408724234808
state:     ONLINE
action:    The pool can be imported using its name or numeric identifier.
config:

    guest1                                     ONLINE
    c2t010000144F5857B200002A0047F12CF6d0    ONLINE
```

Import guest1 and guest2 zpools:

```
# /usr/sbin/zpool import guest1
# /usr/sbin/zpool import guest2
# /usr/sbin/zpool list
NAME      SIZE    USED    AVAIL    CAP    HEALTH    ALTROOT
guest1    9.94G   6.11G   3.83G    61%    ONLINE    -
guest2    9.94G   2.03G   7.91G    20%    ONLINE    -
```

4. On the Sun Ultra 20 server, define the guests from the XML backup files:

```
# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running

# /usr/bin/virsh define /nfs/dumps/guest-nvb85-PV.xml
Domain guest-nvb85-PV defined from /nfs/dumps/guest-nvb85-PV.xml

# /usr/bin/virsh define /nfs/dumps/guest-w2003se-HVM.xml
Domain guest-w2003se-HVM defined from /nfs/dumps/guest-w2003se-HVM.xml

# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
- guest-nvb85-PV                         shut off
- guest-w2003se-HVM                       shut off
```

5. On the Sun Ultra 20 server, start the guests:

```
# /usr/bin/virsh start guest-nvb85-PV
Domain guest-nvb85-PV started

# /usr/bin/virsh start guest-w2003se-HVM
Domain guest-w2003se-HVM started

# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 20 guest-nvb85-PV                       running
 21 guest-w2003se-HVM                     running
```

6. Access the consoles of the guests:

guest-nvb85-PV:

```
# /usr/bin/virsh console guest-nvb85-PV
v3.1.2-xvm chgset 'Mon Mar 03 23:05:33 2008 -0800 15630:ala16c966b70'
SunOS Release 5.11 Version snv_85 64-bit
Copyright 1983-2008 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.

Hostname: dhcp-umpk17-228-48
Reading ZFS config: done.

dhcp-umpk17-228-48 console login:
```

Note: To break out of the console, use the CTRL] keys together.

guest-w2003se-HVM:

Get the vnc details:

```
# /usr/bin/virsh vncdisplay guest-w2003se-HVM
:0
```

Using an xterm where the correct DISPLAY variable is set run:

```
# /usr/bin/vncviewer :0 &
```

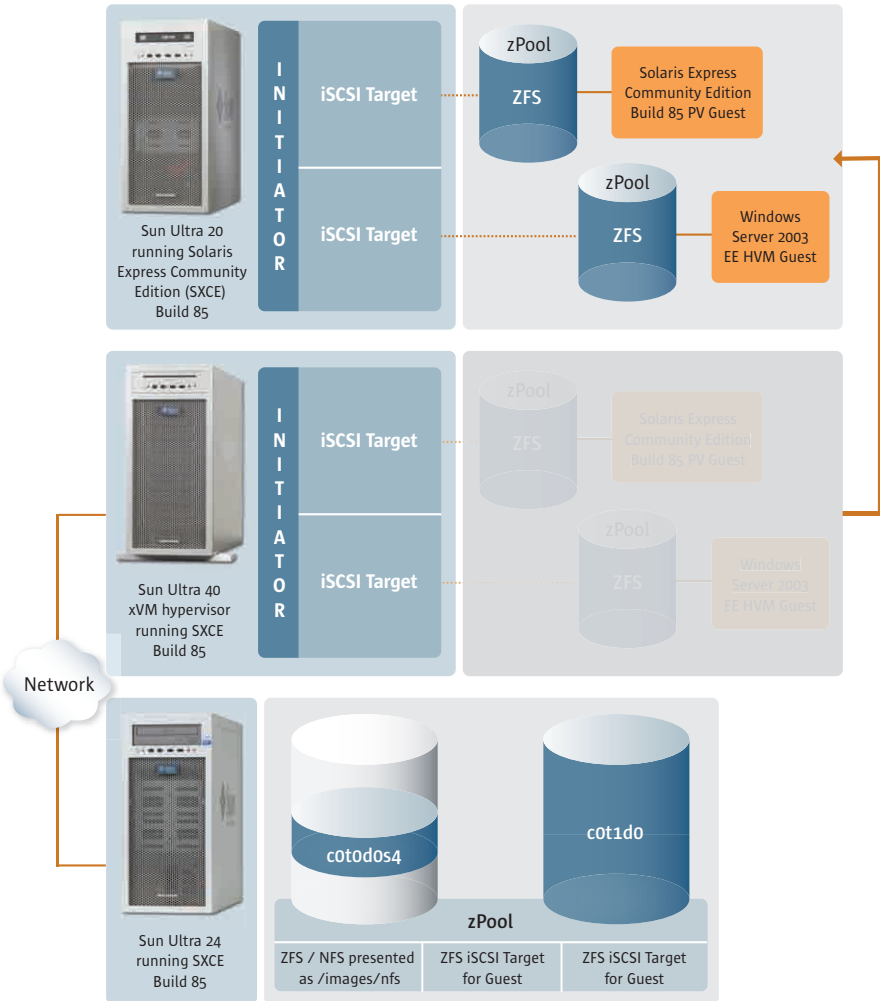


Figure 3—Details this configuration.

Create an NFS-based Windows HVM Guest on the Sun xVM Hypervisor Ultra 20

At the time of writing, Live Migration is possible only on NFS file-based guests. Therefore, we need to create an NFS file-based Windows HVM guest. For more information on creating a Windows HVM guest, go to “Install the Sun xVM Hypervisor and Use it to Configure Domains” [Example of a CLI Windows HVM Installation].

Prerequisites:

- Windows 2003 SE Server CDOM available
- Windows 2003 License Key
- Adequate space on the Sun Ultra 24 NFS/iSCSI server

1. Create a directory on the NFS Server for the NFS-based guest.

On the Sun xVM hypervisor Ultra 20 server (which has /nfs mounted from the Sun Ultra 24 NFS/iSCSI server), go to “Create an NFS-based Windows HVM Guest on the Sun xVM Hypervisor Ultra 20” [Mount the NFS share from the Ultra24 with directIO]:

```
# /usr/bin/mkdir -p /nfs/guests
```

2. Set VNC configuration.

These steps enable the graphical part of the HVM installation to take place:

```
# /usr/sbin/svccfg -s xvm/xend setprop config/vncpasswd = astring: newroot
# /usr/sbin/svcadm refresh xvm/xend; /usr/sbin/svcadm restart xvm/xend
# /usr/sbin/svccfg -s xvm/xend setprop config/vnclisten = astring: 0.0.0.0
# /usr/sbin/svcadm refresh xvm/xend; /usr/sbin/svcadm restart xvm/xend
```

3. For an HVM, install graphics are needed. If you are logged in graphically as a non-root user, perform these steps:

Run the following:

```
# /usr/openwin/bin/xhost +
# # echo $DISPLAY
:0.0
```

Then su - to root and run the following:

```
# DISPLAY=:0.0; export DISPLAY
```

Verify that an xterm can be displayed:

```
# /usr/openwin/bin/xterm
```

4. Use virt-install CLI to install a Windows HVM:

```
# /usr/bin/virt-install -n guest-w2003se-LM-HVM -r 1024 --hvm -f
/nfs/guests/guest-w2003se-LM-HVM.img -s 10 --vnc -c /dev/dsk/c0t0d0s2
```

Note: The -f path specified above automatically creates the .img file for the guest boot file using the size specified by the -s flag.

The install will start and a vncviewer session will be opened. Login using the password set in “Set vnc configuration” above.

There is a bug 6648448 with respect to the Windows 2003 guest and the CD. The install will halt with a Windows error message with regard to access to the CDROM. The workaround follows as Steps 5 through 9 of this procedure:

5. Shut down the running guest:

In a new terminal window logged in as root:

```
# /usr/bin/virsh shutdown guest-w2003se-LM-HVM
Domain guest-w2003se-LM-HVM is being shutdown
```

6. Dump the configuration to an xml file:

```
# /usr/bin/virsh dumpxml guest-w2003se-LM-HVM > /tmp/guest-w2003se-LM-HVM.xml
```

Edit the file created above and insert the following text after the first <disk> section:

```
<disk type='block' device='cdrom'>
  <driver name='phy' />
  <source dev='/dev/dsk/c0t0d0s2' />
  <target dev='hdb' />
  <readonly />
</disk>
```

Note: The <source dev/> section pointing to the Windows CD.

At present, the USB mouse configuration is not complete, which makes the mouse movement awkward. To rectify, change the following line from:

```
<input type='mouse' bus='ps2' />
```

to:

```
<input type='tablet' bus='usb' />
```

The full file should read:

```
<domain type='xen' id='-1'>
  <name>guest-w2003se-LM-HVM</name>
  <uuid>da4eb087-9a30-708f-b52b-009f0ea3700d</uuid>
  <os>
    <type>hvm</type>
    <loader>/usr/lib/xen/boot/hvmloder</loader>
    <boot dev='hd' />
  </os>
  <memory>1048576</memory>
  <vcpu>1</vcpu>
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <features>
    <pae/>
  </features>
```

```

<clock offset='utc' />
  <devices>
    <emulator>/usr/lib/xen/bin/qemu-dm</emulator>
    <interface type='ethernet'>
      <target dev='vif-1.0' />
      <mac address='00:16:3e:5b:08:d8' />
    </interface>
    <disk type='block' device='disk'>
      <driver name='phy' />
      <source dev='/nfs/guests/guest-w2003se-LM-HVM.img' />
      <target dev='hda' />
    </disk>
    <disk type='block' device='cdrom'>
      <driver name='phy' />
      <source dev='/dev/dsk/c0t0d0s2' />
      <target dev='hdb' />
      <readonly />
    </disk>
    <input type='tablet' bus='usb' />
    <graphics type='vnc' port='-1' />
  </devices>
</domain>

```

- Build this added information into the guest:

```

# /usr/bin/virsh define /tmp/guest-w2003se-HVM.xml
Domain guest-w2003se-LM-HVM defined from /tmp/guest-w2003se-LM-HVM.xml

```

- Start the guest:

```

# /usr/bin/virsh start guest-w2003se-LM-HVM
Domain guest-w2003se-LM-HVM started

```

- Reaccess the console to continue the installation.

First, determine which VNC port is tied to the guest:

```

# /usr/bin/virsh vncdisplay guest-w2003se-LM-HVM
:0

```

Therefore, use the `vncviewer` command with the port number to re-access the console:

```

# /usr/bin/vncviewer :0

```

Note: Remember to use the window or `xterm` that has the display variable set.

- Install Windows as normal. Windows will reboot a number of times. If the VNC session is lost, reconnect using the `vncviewer` command above.
- Log in to Windows. To log in when viewing via the `vncviewer`, a CTRL-ALT-DELETE key entry is required. Left click the mouse within the vnc window and press F8. This brings up a vnc menu where CTRL-ALT-DELETE can be selected.

12. If Windows asks for a second CD:

If the Windows Server 2003 distribution has two CDs rather than one, the procedure referenced below should be followed:

- Using “Install the Sun xVM Hypervisor and Use it to Configure Domains” Appendix A: [Detaching a CDROM Device from a Guest], remove the Windows 2003 CD1
- Then, using Appendix A: [Adding a CDROM Device to a Guest], add the Windows 2003 CD2 and finish the Windows installation

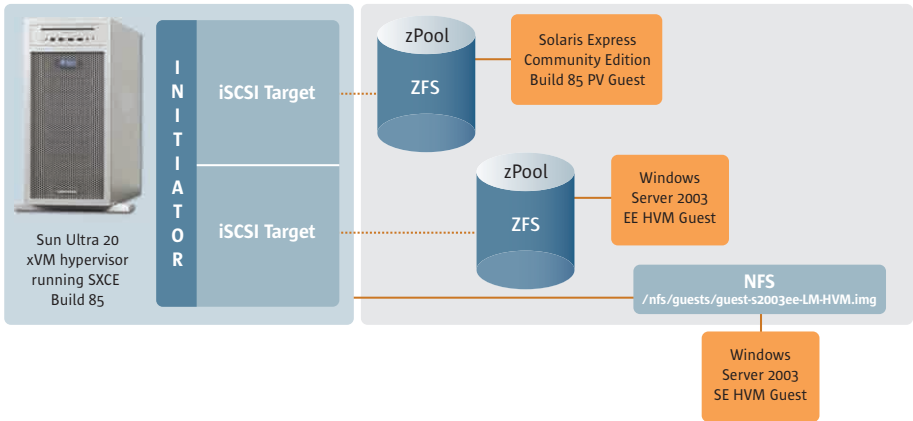


Figure 4—Details this configuration.

Migration using Live Migration of the NFS-based Windows HVM guest-w2003se-LM-HVM from Sun xVM Hypervisor Ultra 20 to Sun xVM Hypervisor Ultra 40 Server

In this section, the steps needed to enable Live Migration within Sun xVM hypervisor and for the Live Migration of an NFS-presented Windows guest from Sun Ultra 20 xVM hypervisor server to the Sun Ultra 40 xVM hypervisor server are described.

1. Before configuring Live Migration for the Sun xVM hypervisor server, shut down any running guests.

Within this environment, the following shutdown commands are used:

```
# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 4 guest-w2003se-LM-HVM                 blocked
 5 guest-nvb85-PV                       running
 6 guest-w2003se-HVM                    no state

# /usr/bin/virsh shutdown guest-nvb85-PV
Domain guest-nvb85-PV is being shutdown

# /usr/bin/virsh shutdown guest-w2003se-HVM
Domain guest-w2003se-HVM is being shutdown

# /usr/bin/virsh shutdown guest-w2003se-LM-HVM
Domain guest-w2003se-LM-HVM is being shutdown

# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 - guest-nvb85-PV                       shut off
 - guest-w2003se-HVM                   shut off
 - guest-w2003se-LM-HVM                shut off
```

2. On hypervisors on both servers:

The following commands must be run on both Sun xVM hypervisor servers to enable Live Migration:

```
# /usr/sbin/svccfg -s xend setprop config/xend-relocation-address = ""
```

The xend-relocation-address field is the address which xend listens on for relocation requests. If blank or not present, all interfaces are used. In the case above, it is set blank, and therefore all interfaces are used. In an enterprise configuration, it would be expected that a separate network would be required to separate Live Migration traffic from the guest traffic.

```
# /usr/sbin/svccfg -s xend setprop config/xend-relocation-hosts-allow =
\"^u20$ ^localhost$\"
```

Note: There must be a space between the hosts above, i.e, “^u20\$ ^localhost\$”

The `xend-relocation-hosts` field is a space-separated list of regular expressions. If the host name of a system matches any one of the given regular expressions, it is allowed to connect and interact with the relocation server if the server has been enabled by the `xend-relocation-server` property. The default is `^localhost$`. The hostnames contained within the command above must be known to the host, which means contained within the `/etc/inet/hosts` file or via DNS / LDAP.

Output of the `xend-relocation` fields for the Sun xVM Ultra 40 server:

```
# /usr/sbin/svccfg -s xend listprop config/xend*
config/xend-relocation-server          boolean          true
config/xend-unix-server                 boolean          true
config/xend-relocation-address          astring
config/xend-relocation-hosts-allow     astring          "^u20$ ^localhost$"
```

Output of the `xend-relocation` fields for the Sun xVM Ultra 20 server:

```
# /usr/sbin/svccfg -s xend listprop config/xend*
config/xend-relocation-server          boolean          true
config/xend-unix-server                 boolean          true
config/xend-relocation-address          astring
config/xend-relocation-hosts-allow     astring          "^u40$ ^localhost$"
```

If any of the `xend-relocation` fields have been changed, a restart of `xend` is required using the following commands:

```
# /usr/sbin/svccadm refresh xend && /usr/sbin/svccadm restart xend
```

3. On the Sun xVM hypervisor Ultra 20 server, restart the guests:

```
# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
- guest-nvb85-PV                         shut off
- guest-w2003se-HVM                       shut off
- guest-w2003se-LM-HVM                    shut off

# /usr/bin/virsh start guest-nvb85-PV
Domain guest-nvb85-PV started

# /usr/bin/virsh start guest-w2003se-HVM
Domain guest-w2003se-HVM started

# /usr/bin/virsh start guest-w2003se-LM-HVM
Domain guest-w2003se-LM-HVM started

# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 4 guest-w2003se-LM-HVM                  LM-HVM blocked
 5 guest-nvb85-PV                         blocked
 6 guest-w2003se-HVM                     blocked
```

- On the Sun xVM hypervisor Ultra 20 server, Live Migrate guest-w2003se-HVM to the Sun xVM hypervisor Ultra 40 server:

```
# /usr/sbin/xm migrate --live guest-w2003se-LM-HVM u40
# /usr/bin/virsh list --all
[root@dhcp-umpk17-229-162:/var/log/xen] virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
 6 guest-w2003se-HVM                    blocked
 8 guest-nvb85-PV                       blocked
- guest-w2003se-LM-HVM                  shut off
```

Note: After the migration command completes, the migrated guest may show as shut off. This entry will be removed in time.

- On the Sun xVM hypervisor Ultra 40 server, check and access guest-w2003se-HVM:

```
# /usr/bin/virsh list --all
Id Name                               State
-----
 0 Domain-0                             running
12 guest-w2003se-LM-HVM                blocked
- guest-nvb85-PV                       shut off
- guest-w2003se-HVM                    shut off
```

Get the vnc details:

```
# /usr/bin/virsh vncdisplay guest-w2003se-LM-HVM
:0
```

Use an xterm where the correct DISPLAY variable is set run:

```
# /usr/bin/vncviewer :0 &
```

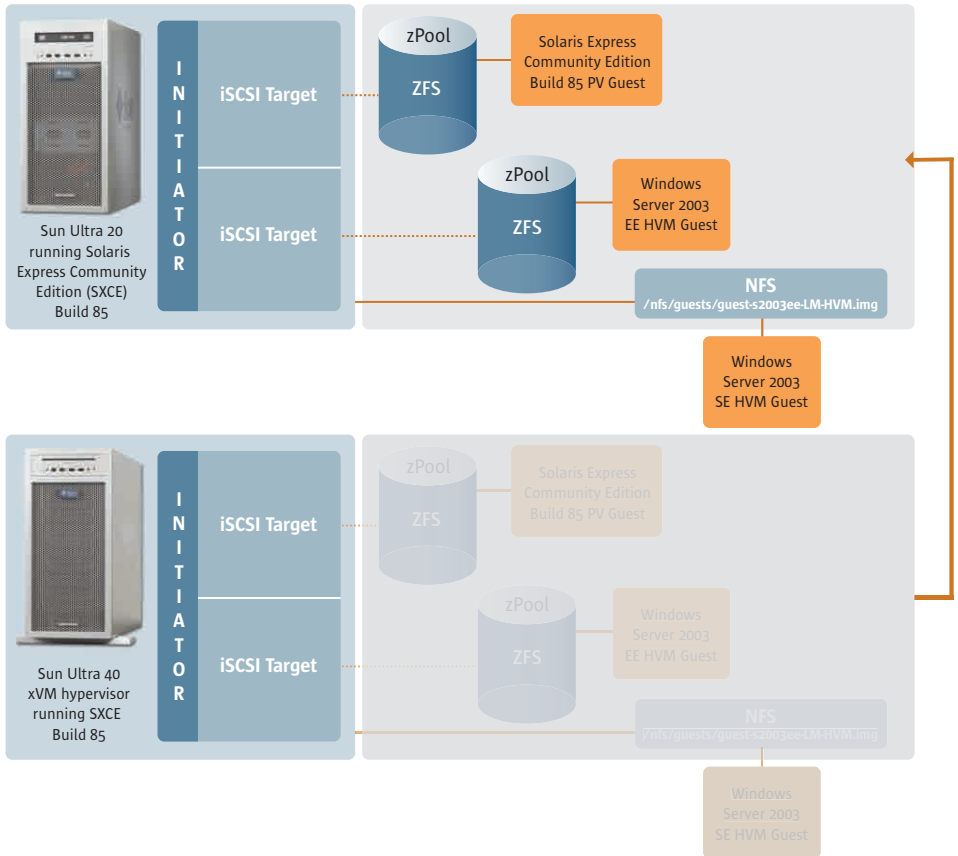


Figure 5—Details this configuration.

Summary

In summary, we have built and enabled the following:

Step 1: Configured a second Sun xVM hypervisor server to access NFS and iSCSI

Step 2: Migrated the guests on the first Sun xVM hypervisor server to the second server using ZFS

Step 3: Created an NFS-based HVM guest on the second server

Step 4: Migrated the NFS-based guest from the second Sun xVM hypervisor server to the first server using Live Migration

developers.sun.com

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN Web sun.com



©2008 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.