

XCAT 2 on AIX

Migating to a new OS and xCAT version

2010/11/4, 08:45:54 AM

1.0 Overview	2
2.0 Back up the NIM and xCAT databases.....	2
2.1 NIM.....	3
2.2 XCAT.....	3
3.0 Shut down any diskless nodes that are being served by the management node....	3
4.0 Upgrade xCAT.....	4
4.1 Upgrade xCAT on the management node.....	4
4.1.1 Shut down xCAT	4
4.1.2 Download and install the latest xCAT dependency software.....	5
4.1.3 Download and install the latest xCAT database software.....	5
4.1.4 Install the latest xCAT packages.....	5
4.1.5 Restart and verify xCAT	6
4.2 Upgrade xCAT on service nodes.....	6
4.2.1 Copy the new rpms to the NIM lpp_source directories.....	6
4.2.2 Update the NIM installp_bundle resources.....	7
4.2.3 Check the osimage (optional).....	7
4.2.4 Run the updatenode command to install the software.....	8
4.2.5 Verify xCAT on the service nodes.....	8
5.0 Migrate the management node to a new version of AIX.....	8
5.1 Verify xCAT and NIM on the management node.....	8
6.0 Upgrade the xCAT service nodes to the new OS version	9
6.1 Create a new xCAT osimage for the service nodes.....	10
6.2 Modify the bosinst_data resource to do a migration install	10
6.3 Copy new rpms to the NIM lpp_source.....	11
6.4 Update the service node bundles.....	11
6.5 Verify the osimage.....	12
6.6 Shut down the diskless clients of the service nodes.....	12
6.7 Initialize the service nodes.....	12
6.8 Initiate a network install.....	12
6.9 Verify the service node	13
7.0 Phase two of xCAT upgrade.....	13
8.0 Upgrade the xCAT compute nodes.....	13
8.1 Installing AIX standalone nodes (using NIM rte method).....	13
8.2 Booting AIX diskless nodes (using stateless method).....	14
8.3 Cloning AIX nodes (using AIX mkysyb images).....	14
8.4 Using xCAT Service Nodes with AIX.....	14

1.0 Overview

This document describes a recommended process to follow when you wish to upgrade the versions of AIX and xCAT that you are using in your xCAT cluster.

Note: The AIX OS level is indicated by a “Version.Release.Modificatio.Fix” number. (ex. VRMF). For example, an AIX operating system version upgrade would be moving from AIX 5.3.0.0 to AIX 6.1.0.0.

If you only want to update your OS release, modification, or fix level you can simply apply the software updates to the management node using the standard AIX tools. (Ex. AIX 6.1.0.0 to 6.1.5.0 etc.) A re-install or a migration install would not be needed.

To upgrade the management node you have the option of doing a complete new overwrite install or doing a migration installation.

The advantage of an operating system migration installation compared to a new and complete overwrite is that most filesets are preserved. This includes almost all directories, such as /home, /var, /usr, the root volume group, logical volumes, system configurations and previously installed software. The only file system that will be new after the migration is /tmp. You can easily avoid losing information you have stored in this directory by copying the important information to another directory before the migration and move it back afterwards. Additionally, after the migration, you can import your user volume groups. It is probably the easiest way to upgrade your system while maintaining all customized information and configuration. Another advantage, especially if you need to minimize the downtime of your system, is that there are fewer reconfigurations tasks to do after the migration.

During a migration, the installation process determines which optional software products are installed on the existing version of the operating system. Any software that was previously installed that has new versions available will be updated.

Before you perform a migration installation, ensure you have reliable backups of your data and any customized applications or volume groups. In particular, on an xCAT management node you should back up your xCAT and NIM databases.

The process described below assumes you are familiar with AIX systems administration and NIM.

2.0 Back up the NIM and xCAT databases

When doing a migration of the operating system the NIM and xCAT databases would normally be preserved. HOWEVER, as a precaution it would be good to back them up and save them on some other system – just in case.

2.1 NIM

Using the SMIT interface:

```
smit nim_backup_db
```

Using the command line interface:

(Assuming the data is saved in */tmp/mynim.backup*.)

```
/usr/lpp/bos.sysmgt/nim/methods/m_backup_db /tmp/mynim.backup
```

The files that are backed up include:

```
./etc/objrepos/nim_attr  
./etc/objrepos/nim_attr.vc  
./etc/objrepos/nim_object  
./etc/objrepos/nim_object.vc  
./etc/NIM.level  
./etc/niminfo
```

To restore the database you can use:

```
smit nim_restore_db
```

or

```
/usr/lpp/bos.sysmgt/nim/methods/m_restore_db /tmp/mynim.backup
```

2.2 XCAT

The **dumpxCATdb** command creates .csv (comma separated value) files for all xCAT database tables and puts them in the directory given by the -p flag. These files can be used by the **restorexCATdb** command to restore the database.

```
dumpxCATdb -p /tmp/db
```

3.0 Shut down any diskless nodes that are being served by the management node

Any diskless nodes that are clients of the management node will not be able to continue running when the management node is not available.

It is better to shut down the diskless nodes while the management node is being updated.

For example, assuming all your diskless nodes are in the xCAT group called “*dsklnodes*” then you could shut them all down by running the following command.

```
xdsh dsklnodes "shutdown -F &"
```

Any standalone nodes that are clients of the management node can be left running.

4.0 Upgrade xCAT

You may or may not need to upgrade xCAT as part of an OS migration.

The one requirement is that the version of xCAT you use must support BOTH the current OS version and the version you wish to migrate to.

In some cases you may need to upgrade xCAT in two phases. The first upgrade may be needed to migrate to a newer OS version using an xCAT version that supports both. The second phase would be to upgrade xCAT to the latest version.

Whether you need to upgrade xCAT or not, you must upgrade the versions of the xCAT dependency and database rpms to the versions that were built for the new AIX level.

The xCAT dependency and database tarballs contain versions of the rpm packages built for the different levels of AIX. So, for example, if you migrate your OS to AIX 7.1 then you must also update your dependency rpms to the versions built for AIX 7.1. The xCAT dependency and database packages contain subdirectories corresponding to the different AIX versions.

Note: If you upgrade the xCAT version on the xCAT management node you must also upgrade the xCAT version on any service nodes that are being used!

4.1 Upgrade xCAT on the management node

4.1.1 Shut down xCAT

xCAT should be shut down before trying to do any upgrades.

```
stopsrc -s xcatd
```

4.1.2 Download and install the latest xCAT dependency software

Download the latest dep-aix-*.tar.gz tar file from <http://xcat.sourceforge.net/#download> and copy it to a convenient location on your xCAT management node.

- Unwrap the tar file. For example:

```
gunzip dep-aix-2.5.tar.gz  
tar -xvf dep-aix-2.5.tar
```
- Run the **instoss** script (contained in the tar file) to install the OSS packages. Please make sure the /opt and the other file systems have enough disk space to install these OSS packages before running the **instoss** script.

4.1.3 Download and install the latest xCAT database software

Some of the optional databases that you may be using with xCAT will require upgrades to the dependency software.

XCAT currently supports SQLite, MySQL, PostgreSQL, and DB2. As a convenience, the xCAT site provides downloads for MySQL and PostgreSQL software.

([xcat-postgresql-snap201007150920.tar.gz](http://xcat.sourceforge.net/#download) and [xcat-mysql-201005260807.tar.gz](http://xcat.sourceforge.net/#download))

See the following xCAT documents for instructions on how to upgrade these databases.

[Setup MySQL as the xCAT Database](#)

[Setup PostgreSQL as the xCAT Database](#)

[Setup DB2 as the xCAT Database](#)

The database tar files that are available on the xCAT web site may contain multiple versions of RPMs – one for each AIX operating system level. You must install the versions of the rpms that correspond to your OS version.

4.1.4 Install the latest xCAT packages

Download the latest xCAT for AIX tar file from <http://xcat.sourceforge.net/#download> and copy it to a convenient location on your xCAT management node.

- Unwrap the xCAT tar file. For example,

```
gunzip core-aix-2.5.tar.gz  
tar -xvf core-aix-2.5.tar
```
- Run the **instxcat** script (contained in the tar file) to install the xCAT

software. (The **instxcat** script and all the RPMs are located in the `xcatcore` subdirectory.)

4.1.5 Restart and verify xCAT

To restart xCAT run the following command.

```
restartxcatd
```

To verify xCAT function try running some basic xCAT commands.

- Run “`rpm -qa`” to see if the latest xCAT and dependency software is upgraded.
- Run the “`lsdef -h`” to check if the xCAT daemon is working. (If you get a correct response then you should be Ok.)
- Check the xCAT database.
 - Try running some commands such as: “**lsdef -t site -l**”
 - If necessary you can restore the database you saved earlier by running the **restorexCATdb** command.

4.2 Upgrade xCAT on service nodes

4.2.1 Copy the new rpms to the NIM `lpp_source` directories

To update the versions of xCAT and the dependency packages on the service nodes you must first copy all the new rpms to the `lpp_source` resource that is being used for the service nodes. This includes the new xCAT, dependency, and database rpms. These are all available in the xCAT tarballs that you download from the xCAT web site.

A simple way to copy software to the `lpp_source` locations is to use the “`nim -o update`” command. With this NIM command you simply need to provide the name of the NIM `lpp_source` resource and it will automatically copy your files to the correct locations.

For example, assume all your dependency software has been downloaded and unwrapped in the temporary location `/tmp/images`. (The rpms you need are in subdirectories corresponding to the OS version you are running.)

To add all the deps packages to the `lpp_source` resource named “`61image_lpp_source`” you could run the following command:

```
nim -o update -a packages=all -a source=/tmp/images/xcat-dep/6.1  
610image_lpp_source
```

The NIM command will find the correct directories and update the *lpp_source* resource.

4.2.2 Update the NIM *installp_bundle* resources

Make sure that any NIM *installp_bundle* resources you are using are updated to contain the new xCAT, dependency and database packages.

If you have an xCAT *osimage* defined for the service nodes it will include an *installp_bundle* file that contains a list of all the additional software that is needed on the service nodes. You must check this file to make sure it includes the new rpms that you just copied to your *lpp_source*.

The new sample xCAT bundle files are located either in the new tarball or installed in */opt/xcat/share/xcat/installp_bundles*.

4.2.3 Check the *osimage* (optional)

To avoid potential problems when installing additional software it is advisable to verify that all the software that you wish to install has been copied to the appropriate NIM *lpp_source* directory and that there are no duplicate packages.

Any software that is specified in the “otherpkgs” or the “*installp_bundle*” attributes of the *osimage* definition must be available in the corresponding *lpp_source* directories.

Also, since the sample bundle files now include wildcards (*) instead of specific versions you must make sure there are not multiple version of the same rpm in the *lpp_source* directories. If there are multiple version of the same rpm in the *lpp_source* directory then NIM will produce an error when trying to install them.

The older versions of the rpms must be removed from the *lpp_source* directory.

This can be done manually by going to the *lpp_source* location and then to the “RPMS/ppc” subdirectory and then removing the rpms.

However, starting in xCAT 2.4.3 you can use the **chkosimage** command to check the *lpp_source* directories and remove the older versions.

To check an *osimage* called *610image*:

```
chkosimage -V 610image
```

To check an *osimage* and remove duplicate rpms:

```
chkosimage -V -c 610image
```

See the **chkosimage** man page for details.

4.2.4 Run the updatenode command to install the software

You can update the versions of the xCAT, dependency, and database software on the service nodes by using the **updatenode** command.

By default the command will use the “*installp_bundle*” and “*otherpkgs*” attributes of the *osimage* definition for the node.

In the following example assume that all the service nodes have been added to the xCAT group called “service” and that the specified flags should be used when installing the packages.

```
updatenode service -V -S rpm_flags="-Uvh --replacepkgs --nodeps"  
installp_flags="-agQXY"
```

4.2.5 Verify xCAT on the service nodes

1. Run “*rpm -qa*” on service node to see if the xCAT and dependency software is upgraded.
2. Try running some xCAT commands such as “*lsdef -a*” to see if the *xcatd* daemon is running and if data from xCAT database is set properly.
3. Check the NIM configuration on the service nodes
4. Run *xdsh* to the service nodes (ex. “*xdsh service date*”)

5.0 Migrate the management node to a new version of AIX

There are several techniques provided by AIX that can be used to migrate a standalone system. They are described in the AIX documentation listed below and will not be repeated here.

Listing of AIX documents.

<http://publib.boulder.ibm.com/infocenter/aix/v7r1/index.jsp>

AIX “Installation and Migration Guide”:

http://publib.boulder.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.install/doc/insgdrf/insgdrf_pdf.pdf

5.1 Verify xCAT and NIM on the management node

Check that xCAT is still functioning properly on the management node. (ex. run some xCAT commands, check the database etc.)

If xCAT isn't running or some dependency packages may not be functioning.

1. Run the **instoss** script (contained in the tar file) to re-install the OSS packages corresponding to the OS version you are running, or run “rpm -Uvh” or “rpm -Uvh --force” to upgrade the single package.
2. Run 'restartxcatd' to restart xCAT.

Depending on the method you used to do the AIX migration, the NIM master on the xCAT management node may or may not be functioning.

If it is not:

1. Check that the master fileset is installed and that it has been updated to the latest level. (*bos.sysmgt.nim.master*). This should have been in your *lpp_source* directory and also in the service node bundle file.

2. Save the current */etc/niminfo* file

```
mv /etc/niminfo /etc/niminfo.orig
```

3. Reconfigure NIM.

You could use the SMIT interface or the **nimconfig** command to re-configure NIM. (Note: Make sure you specify the appropriate interface to use as the NIM primary interface. Typically the interface that will be used to install the nodes is used as the primary interface.) The **nimconfig** command would be similar to the following:

```
nimconfig -a pif_name=en1 -a netname=clstr_net -r
```

4. It should not be required, but you can restore the previously saved NIM database backup by using SMIT:

```
smit nim_restore_db
```

or by using the following NIM method.

```
/usr/lpp/bos.sysmgt/nim/methods/m_restore_db /tmp/mynim.backup
```

5. Validate existing NIM resources.

- Make sure the NIM definitions are correct, that the files and directories all exist etc.
- Run the “nim -o check ...” operation on NIM SPOT and lpp_source resources.

6.0 Upgrade the xCAT service nodes to the new OS version

If you migrate your xCAT management node to a new OS version then any service nodes you are using should also be upgraded.

Since the AIX service nodes are standalone systems they can remain running while you update the management node. However if you leave the service nodes at back levels it could cause problems with the xCAT support.

For more details on installing a service node refer to the following xCAT document.

<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2onAIXServiceNodes.pdf>

6.1 Create a new xCAT osimage for the service nodes

Use the xCAT **mknimimage** command to create a new *osimage* definition using the new AIX OS version.

For example, assuming all the AIX software has been copied to */aix61source* you could run:

```
mknimimage -s /aix61source AIX61image
```

This will create the required NIM resources as well as the new xCAT *osimage* definition.

Make sure you update your *osimage* definition with any other resources or attributes you wish to include. (Such as the service node *installp_bundle* resource etc.) Refer to the “xCAT2onAIXServiceNodes” document mentioned above for further details.

6.2 Modify the *bosinst_data* resource to do a migration install

By doing a migration install, among other things, you will preserve the NIM database and NIM resources that were created on the service nodes.

A default NIM *bosinst_data* resource was created automatically when you created the new *osimage*.

A **bosinst_data** resource represents a file that contains information for the BOS installation program. The resource is typically created by starting with the template that is shipped with AIX. See the `/usr/lpp/bosinst/bosinst.template.README` file for more information on *bosinst_data* file contents.

To perform a migration install you must make one change to the file.

You must change the “INSTALL_METHOD” value to “migrate”.

This can be changed in the file either before or after you define the resource but it must be done before initiating a network boot of the system.

6.3 Copy new rpms to the NIM lpp_source

There are new versions of the xCAT, dependency, and database rpms that must be copied to the new *lpp_source* resource created in the previous step.

These rpms were all contained in the xCAT tarballs you downloaded when you upgraded the management node.

A simple way to copy software to the *lpp_source* locations is to use the “nim -o update” command. With this NIM command you simply need to provide the name of the NIM *lpp_source* resource and it will automatically copy your files to the correct locations.

For example, assume all your dependency software has been downloaded and unwrapped in the temporary location `/tmp/images`. (The rpms you need are in subdirectories corresponding to the OS version you are running.)

To add all the deps packages to the *lpp_source* resource named “*610new_lpp_source*” you could run the following command:

To copy the rpms to the correct location in the *lpp_source* you can use the “***nim -o update ..***” command.

For example:

```
nim -o update -a packages=all -a source=/tmp/images/xcat-dep/6.1  
610new_lpp_source
```

The NIM command will find the correct directories and update the *lpp_source* resource.

6.4 Update the service node bundles

XCAT ships sample NIM `installp_bundle` files. (See `/opt/xcat/share/xcat/installp_bundles`) There are bundle files to use for the service nodes and bundle files to use for the compute nodes. When you switch to a new OS version you must use the corresponding version of the bundle files.

For example, if you are now using AIX 6.1 then you need to use the `xCATaixSN61.bnd` file when installing the service node.

Copy the bundle file to an exportable location (ex. `/install/nim/installp_bundle`) and define the NIM `installp_bundle` resource using SMIT or the “`nim -o define ...`” command.

Once the resource is defined you must add it to the `osimage` definition.

```
chdef -t osimage -o AIX61image installp_bundle=xCATaixSN61
```

Check the contents of the bundle file to make sure it includes all the software you will need. There are some commented-out entries that are optional that you may need. For example, the additional rpms needed for the different databases that you may be using.

6.5 Verify the osimage

You can use the xCAT `chkosimage` command to check that all the software called out in the bundles files are included in the `lpp_source` resource.

```
chkosimage AIX61image
```

6.6 Shut down the diskless clients of the service nodes.

It is recommended that you shut down all diskless clients of the service nodes before upgrading them. Any standalone clients may be left running.

For example, to shut down the diskless compute node named “`compute02`” you could run the following command.

```
xdsh compute02 “shutdown -F &”
```

6.7 Initialize the service nodes

You can use the xCAT **nimnodeset** command to initialize the AIX standalone service nodes. This command uses information from the xCAT *osimage* definition and default values to run the appropriate NIM commands.

For example, to set up all the nodes in the group “*service*” to install using the *osimage* named “*610SNimage*” you could issue the following command.

```
nimnodeset -i 610SNimage service
```

6.8 Initiate a network install

Initiate a remote network boot request using the xCAT **rnetboot** command. For example, to initiate a network boot of all nodes in the group “*service*” you could issue the following command.

```
rnetboot service
```

Result:

The service node will be installed with the new OS and new rpms specified in the bundle files. The xCAT posts scripts will do the xCAT setup and configure the system as a service node.

6.9 Verify the service node

1. Run “rpm -qa” on service node to see if the dependency software is upgraded.
2. Try running some xCAT commands to see if xcatd daemon is running.
3. Run xdsh to the service nodes (ex. “xdsh service date”)
4. Check the NIM configuration on the service nodes

7.0 Phase two of xCAT upgrade

If you had to use a back-level version of xCAT to do your OS migration you can now upgrade xCAT to the most recent level. (For example, if you had to use xCAT 2.4 to handle the OS migration but now you wish to upgrade to xCAT 2.5.)

Use the same procedure mentioned in the “Update xCAT” section above.

8.0 Upgrade the xCAT compute nodes

The details for upgrading compute nodes is covered in several other xCAT documents.

If you have a large number of compute node to upgrade you may wish to use the xCAT rolling update support. See the **rollupdate** man page for usage details.

8.1 *Installing AIX standalone nodes (using NIM rte method)*

<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2onAIXinstall.pdf>

8.2 *Booting AIX diskless nodes (using stateless method)*

<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2onAIXDiskless.pdf>

8.3 *Cloning AIX nodes (using AIX mksysb images)*

<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2onAIXmksysb.pdf>

8.4 *Using xCAT Service Nodes with AIX*

<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2onAIXServiceNodes.pdf>