

xCAT 2 Cookbook for Linux on IBM System P

2010 / 9 / 21, 00:29:25

Contents

Table of Contents

1. Introduction	4
2. Install xCAT 2 on the Management node.....	5
3. Setup the management node	5
3.1. [Power 5] Workaround the atftpd issue.....	5
3.1.1. Remove atftp.....	5
3.1.2. Install the tftp server needed by xCAT, and restart it.....	5
3.1.3. Restart the tftp server.....	6
3.2. Setup common attributes for xCAT in the database.....	6
3.2.1. Add the default account for system.....	6
3.2.2. Add the default account for hmc.....	6
3.3. Define the compute nodes.....	6
3.3.1. Gather Node information using the rscan command.....	6
3.3.1.1. Define HMC as an xCAT node.....	6
3.3.1.2. Discover the LPARs managed by HMC.....	7
3.3.1.3. Define xCAT node using the stanza file.....	8
3.3.1.4. Define xCAT groups	8
3.3.2. Update the attributes of the node.....	9
3.3.2.1. Set the resource attributes of the node.....	9
3.3.2.2. Check the site.master value.....	9
3.3.2.3. Set the type attributes of the node.....	9
3.4. Set up customization scripts (optional).....	9

3.5. Add NTP setup script (optional).....	10
3.6. Setup the Services and Definition.....	10
3.6.1. Setup the networks table.....	10
3.6.2. Setup Name Resolution.....	10
3.6.2.1. Setup /etc/hosts with entries for all you nodes, hmcs, fsps.....	10
3.6.2.2. Setup the nameserver	10
3.6.2.3. Setup the DNS attributes in the Site table.....	10
3.6.2.4. Setup DNS configuration	11
3.6.2.5. Updating the DNS configuration	11
3.6.3. Configure conserver	11
3.6.3.1. Set conserver attributes for the nodes	11
3.6.3.2. Update conserver configuration.....	11
3.6.4. Check rcons(rnetboot and getmacs depend on it).....	12
3.6.5. Check hardware control setup to the nodes.....	12
3.6.6. Update the mac table with the address of the node(s).....	12
3.6.7. Setup dhcp service.....	13
3.6.7.1. Setup the dhcp listen interfaces in site table.....	13
3.6.7.2. [SLES] Check the installation of dhcp-server	13
3.6.7.3. Configure the DHCP.....	13
4. Install a Compute Node	13
4.1. Prepare the installation source.....	13
4.2. Statefull Node installation.....	13
4.2.1. Customize the install profile.....	13
4.2.1.1. Search order for installation templates.....	14
4.2.1.2. Install other specific packages.....	15
4.2.2. Set the node status to ready for installation.....	15
4.2.3. Use network boot to start the installation.....	15
4.2.4. Check the installation results.....	15
4.2.5. Install a new Kernel on the nodes.....	16
4.3. Stateless node installation.....	16
4.3.1. Generate the stateless image for compute node.....	16
4.3.1.1. Make the compute node add/exclude packaging list.....	16
4.3.1.2. Setting up postinstall files	17
4.3.1.1. Run image generation.....	18
4.3.1.2. Pack the image.....	18
4.3.2. Set the node status ready for network boot.....	18
4.3.3. Use network boot to start the installation.....	18

4.3.4. Check the installation result.....	18
4.3.5. Installing a new Kernel in the stateless image.....	19
4.3.6. Remove an image.....	19
5. Firmware upgrade.....	20
5.1. Requirements.....	20
5.1.1. Enable the HMC to allow remote ssh connections.....	20
5.1.2. Define the necessary attributes.....	20
5.1.3. Define the HMC as a node.....	20
5.1.4. Setup SSH connection to HMC.....	20
5.1.5. Get the Microcode update package and associated XML file.....	20
5.2. Perform Firmware upgrade for CEC on P5/P6.....	20
5.2.1. Define the CEC as a node on the management node.....	20
5.2.2. Setup SSH connection to HMC.....	21
5.2.3. Check firmware level.....	21
5.2.4. Update the firmware.....	21
5.3. Perform Firmware upgrades for BPA on P5/P6.....	22
5.3.1. Define the BPA as a node on the management node.....	22
5.3.2. Setup SSH connection to HMC.....	22
5.3.3. User rinv to check the firmware level (see rinv manpage).....	22
5.3.4. Update the firmware.....	22
5.4. Commit currently activated LIC update(copy T to P) for a CEC/BPA on p5/p6.....	23
5.4.1. Check firmware level.....	23
5.4.2. Commit the firmware LIC.....	23
6. Advanced features.....	23
6.1. Use the driver update disk:.....	23
7. References.....	23

1. Introduction

This cookbook introduces how to use the xCAT2 to install Linux on the IBM power system machines.

The power system machines have the following characteristics:

1. May have multiple LPARs (an LPAR will be the target machine to install an operating system image on, i.e. the LPAR will be the compute node);
2. The Ethernet card and SCSI disk can be virtual devices;
3. An HMC or IVM is used for the HCP (hardware control point)

xCAT supports two types of installations for compute nodes: Diskfull installation (Statefull) and Diskless (Stateless). xCAT also supports hierarchical management clusters where one or more service nodes are used to handle the installation and management of compute nodes. Please refer to [xCAT2SetupHierarchy.pdf](#) for hierarchical usage.

Based on the two types of installation, the following installation scenarios will be described in this document:

1. Install a stateful compute node
2. Install a stateless compute node

To provide the easier understanding of the installation steps, this cookbook provides an example to introduce the xCAT management operations:

The management node:

```
Arch: an LPAR on a p5/p6 machine
OS: Red Hat Enterprise Linux 5.2
Hostname: pmanagenode
IP: 192.168.0.1
HCP: HMC
```

The management Network:

```
Net: 192.168.0.0
NetMask: 255.255.255.0
Gateway: 192.168.0.1
Cluster-face-IF: eth1
dhcpserver: 192.168.0.1
tftpserver: 192.168.0.1
nameservers: 192.168.0.1
```

The compute nodes:

```
Arch: an LPAR on a p5/p6 machine
OS: Red Hat Enterprise Linux 5.2
HCP: HMC
```

```
Hostname: pnode1 - this node will be installed
statefull
IP: 192.168.0.10
Cluster-face-IF: eth0
```

```
Hostname: pnode2 - this node will be installed
stateless
IP: 192.168.0.20
Cluster-face-IF: eth0
```

The Hardware Control Point:

```
Name: hmcl
IP: 192.168.0.100
```

```
xCAT version:
xCAT-2.1+
```

2. Install xCAT 2 on the Management node

Before proceeding to setup your pLinux Cluster, you should first read [xCAT2top](#) for information on downloading and installing xCAT on your Management Node. Some xCAT database tables will be used in the following chapters, you can refer to [xcatdb manpage](#) for more details on xCAT database tables.

3. Setup the management node

3.1. [Power 5] Workaround the atftpd issue

The tftp client in the open firmware of p5 is only compatible with tftp-server instead of atftpd which is required by xCAT2. So we have to remove the atftpd first and then install the tftp-server. This is not required for Power6 or later.

3.1.1. Remove atftp

```
rpm -qa | grep atftp
Could find one or both of the following rpms:
atftp-xcat-*
atftp-*
```

```
service tftpd stop
rpm --nodeps -e atftp-xcat atftp
```

3.1.2. Install the tftp server needed by xCAT, and restart it

[RH]:

```
yum install tftp-server.ppc
```

[SLES]:

```
zypper install tftp
```

3.1.3. Restart the tftp server

Notes: make sure the entry "disable=no" in the /etc/xinetd.d/tftp.
service xinetd restart

3.2. Setup common attributes for xCAT in the database

The xCAT database table "passwd" contains default userids and passwords for xCAT to access cluster components. This section will describe how to set the default userids and passwords for system and hmc in xCAT database table.

3.2.1. Add the default account for system

```
chtab key=system passwd.username=root passwd.password=cluster
```

3.2.2. Add the default account for hmc

```
chtab key=hmc passwd.username=hscroot passwd.password=abc123
```

Note: The username and password for xCAT to access the HMCs can be specified through mkdef or chdef command, this is useful especially when some specific HMCs use the different username and password with the default ones. For example:

```
mkdef -t node -o hmc1 groups=hmc,all nodetype=hmc mgt=hmc username=hscroot  
password=abc1234
```

```
chdef -t node -o hmc1 username=hscroot password=abc1234
```

3.3. Define the compute nodes

The definition of a node is stored in several tables of the xCAT database.

You can use **rscan** command to discover the HCP to get the nodes that managed by this HCP. The discovered nodes can be stored into a stanza file. Then edit the stanza file to keep the nodes which you want to create and use the mkdef command to create the nodes definition.

3.3.1. Gather Node information using the rscan command

3.3.1.1. Define HMC as an xCAT node

First, define the hardware control point as a node object.

The following command will create an xCAT node definition for an HMC with a host name of "hmc1". The *groups*, *nodetype*, *mgt*, *username*, and *password* attributes will be set.

```
mkdef -t node -o hmc1 groups=hmc,all nodetype=hmc mgt=hmc  
username=hscroot password=abc123
```

to change and add new groups:

```
chdef -t node -o hmc1 groups=hmc,rack1,all
```

to verify your data:

```
lsdef -l hmc1
```

If xCAT Management Node is in the same service network with HMC, you will be able to discover the HMC and create an xCAT node definition for the HMC automatically.

```
lsslp -w -s HMC
```

To check for the hmc name added to the nodelist:

```
tabdump nodelist
```

The above xCAT command `lsslp` discovers and writes the HMCs into xCAT database, but we still need to set HMCs' username and password.

```
chdef -t node -o <hmcname from lsslp> username=hscroot password=abc123
```

For more details with hardware discovery feature in xCAT, please refer to document:

<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2pHWManagement.pdf>

3.3.1.2. Discover the LPARs managed by HMC

Run the `rscan` command to gather the LPAR information. This command can be used to display the LPAR information in several formats and can also write the LPAR information directly to the xCAT database. In this example we will use the “-z” option to create a stanza file that contains the information gathered by `rscan` as well as some default values that could be used for the node definitions.

To write the stanza format output of `rscan` to a file called “node.stanza” run the following command. We are assuming, for our example, that the hmc name returned from `lsslp` was `hmc1`.

```
rscan -z hmc1 > node.stanza
```

This file can then be checked and modified as needed. For example you may need to add a different name for the node definition or add additional attributes and values.

Note: The stanza file will contain stanzas for things other than the LPARs. This information must also be defined in the xCAT database. The stanza will repeat the same bpa information for multiple fsp(s). It is not necessary to modify the non-LPAR stanzas in any way.

The stanza file will look something like the following.

```
Server-9117-MMA-SN10F6F3D:
  objtype=node
  nodetype=fsp
  id=5
  model=9118-575
  serial=02013EB
  hcp=hmc01
  pprofile=
  parent=Server-9458-10099201WM_A
```

```

groups=fsp,all
mgt=hmc

pnode1:
  objtype=node
  nodetype=lpar,osi
  id=9
  hcp=hmc1
  pprofile=lpar9
  parent=Server-9117-MMA-SN10F6F3D
  groups=lpar,all
  mgt=hmc
  cons=hmc

pnode2:
  objtype=node
  nodetype=lpar,osi
  id=7
  hcp=hmc1
  pprofile=lpar6
  parent=Server-9117-MMA-SN10F6F3D
  groups=lpar,all
  mgt=hmc
  cons=hmc

```

Note: The **rscan** command supports an option to automatically create node definitions in the xCAT database. To do this the LPAR name gathered by **rscan** is used as the node name and the command sets several default values. If you use the “-w” option, make sure the LPAR name you defined will be the name you want used as your node name.

For a node which was defined correctly before, you can use the “lsdef -z [nodename]> node.stanza” command to export the definition into the node.stanza, and use command “cat node.stanza | chdef -z” to update the node.stanza according to your need.

3.3.1.3. Define xCAT node using the stanza file

The information gathered by the **rscan** command can be used to create xCAT node definitions by running the following command:

```
cat node.stanza | mkdef -z
```

Verify the data:

```
lsdef -t node -l all
```

3.3.1.4. Define xCAT groups

See the section “xCAT node group support” in [xCAT2top](#) for more details on how to define xCAT groups. For the example below add the compute group to the nodes.

```
chdef -t node -o pnode1,pnode2 -p groups=compute
```

3.3.2. Update the attributes of the node

3.3.2.1. Set the resource attributes of the node

```
chdef -t node -o pnode1 netboot=yaboot tftpserver=192.168.0.1
      nfsserver=192.168.0.1 monserver=192.168.0.1 xcatmaster=192.168.0.1
      installnic="eth0" primarynic="eth0"
```

Note: Please make sure the attributes "installnic" and "primarynic" are set up by the correct Ethernet Interface of compute node. Otherwise the compute node installation may hang on requesting information from an incorrect interface. The "installnic" and "primarynic" can also be set to mac address if you are not sure about the Ethernet interface name, the mac address can be got through getmacs command. The installnic" and "primarynic" can also be set to keyword "mac", which means that the network interface specified by the mac address in the mac table will be used.

Make sure that the address used above (192.168.0.1) is the address of the Management Node as known by the node. Also make sure site.master has this address.

3.3.2.2. Check the site.master value

Make sure site.master is the address or name known by the node

To change site.master to this address:

```
chtab key=master site.value=192.168.0.1
```

3.3.2.3. Set the type attributes of the node

```
chdef -t node -o pnode1 os=<os> arch=ppc64 profile=compute
```

Note: The <os> can be rh, centos*, fedora*, sles*. (where * is the version #) For example, the <os> can be rhels5.2 or sles11.*

3.4. Set up customization scripts (optional)

xCAT supports the running of customization scripts on the nodes when they are installed. You can see what scripts xCAT will run by default by looking at the "xcatdefaults" entry in the xCAT "postscripts" database table. The "postscripts" attribute of the node definition can be used to specify the comma separated list of the scripts that you want to be executed on the nodes. The order of the scripts in the list determines the order in which they will be run.

For example, if you want to have your two scripts called "foo" and "bar" run on node "node01" you could use the **chdef** command as follows.

```
chdef -t node -o node01 -p postscripts=foo,bar
```

(The "-p" means to add these to whatever is already set.)

See the [TopDoc Appendix D](#) for more information.

3.5. Add NTP setup script (optional)

To enable the NTP services on the cluster, first configure NTP on the management node and start **ntpd**.

Next set the “*ntp servers*” attribute in the site table. Whatever time servers are listed in this attribute will be used by all the nodes that boot directly from the management node.

```
chdef -t site ntpservers= myMN
```

To have xCAT automatically set up ntp on the cluster nodes you must add the **setupntp** script to the list of postscripts that are run on the nodes.

To do this you can either modify the “postscripts” attribute for each node individually or you can just modify the definition of a group that all the nodes belong to.

For example, if all your nodes belong to the group “compute”, see Error: Reference source not found, then you could add **setupntp** to the group definition by running the following command.

```
chdef -p -t group -o compute postscripts=setupntp
```

3.6. Setup the Services and Definition

A basic networks table was created for you during the xCAT install. Review that table and add additional networks based on your hardware configuration.

3.6.1. Setup the networks table

Create the networks that used for cluster management:

```
mkdef -t network -o net1 net=192.168.0.0 mask=255.255.255.0
  gateway=192.168.0.1 mgtifname=eth1 dhcpserver=192.168.0.1
  tftpserver=192.168.0.1 nameservers=192.168.0.1
```

3.6.2. Setup Name Resolution

3.6.2.1. Setup /etc/hosts with entries for all you nodes, hmcs, fsps

```
127.0.0.1          localhost
192.168.0.1       pmanagenode
192.168.0.10     pnode1
192.168.0.20     pnode2
192.168.0.100    hmcl
```

3.6.2.2. Setup the nameserver

Add following lines into /etc/resolv.conf

```
search cluster.net
nameserver 192.168.0.1
```

3.6.2.3. Setup the DNS attributes in the Site table

Setup local machine as nameserver:

```
chdef -t site nameservers=192.168.0.1
```

Setup the external nameserver:

```
chdef -t site forwarders=9.112.4.1
```

Setup the local domain name:

```
chdef -t site domain=cluster.net
```

3.6.2.4. Setup DNS configuration

```
makedns  
service named start  
chkconfig --level 345 named on
```

3.6.2.5. Updating the DNS configuration

If you add nodes or update the networks table at a later time, then rerun makedns:

```
makedns  
service named restart
```

3.6.3. Configure conserver

The xCAT rcons command uses the conserver package to provide support for multiple read-only consoles on a single node and the console logging. For example, if a user has a read-write console session open on node node1, other users could also log in to that console session on node1 as read-only users. This allows sharing a console server session between multiple users for diagnostic or other collaborative purposes. The console logging function will log the console output and activities for any node with remote console attributes set to file /var/log/consoles/<node_name>, the console logging files can be replayed for debugging or any other purpose.

3.6.3.1. Set conserver attributes for the nodes

Check to see if the cons attribute was set by rscan by running:

```
lsdef pnode1
```

If the cons attribute for the nodes is not set, then do the following:

```
chdef -t node -o pnode1,pnode2 cons=hmc conserver=<management node>
```

Note: conserver=<management node> is the default, so it is optional to set.

3.6.3.2. Update conserver configuration

Each xCAT node with remote console attributes set should be added into the conserver configuration file to make the rcons work. The xCAT command makeconservercf will put all the nodes into conserver configuration file /etc/conserver.cf. The makeconservercf command must be run when there is any node definition changes that will affect the conserver, such as adding new nodes, removing nodes or changing the nodes' remote console settings.

```
makeconservercf
service conserver stop
service conserver start
```

3.6.4. Check rcons(rnetboot and getmacs depend on it)

```
rcons pnode1
```

If it works ok, you will get into the console interface of the pnode1. If it does not work, review your node setup from above.

3.6.5. Check hardware control setup to the nodes

See if your setup is correct at this point, run rpower to check node status:

```
rpower pnode1 stat
```

3.6.6. Update the mac table with the address of the node(s)

Before run getmacs, make sure the node is off. If not, please force the lpar shutdown with “rpower pnode1 off” command. The reason is that HMC has one issue that cannot shutdown linux nodes which is in running state.

If there's only one Ethernet adapter on the node or you have specified the installnic or primarynic attribute of the node, using following command can get the correct mac address.

```
getmacs pnode1
```

But, if there're more than one Ethernet adapters on the node and you don't know which one has been configured for the installation process, or the lpar is just created and there is no active profile for that lpar, or the lpar is on a P5 system and there is no lhea/sea ethernet adapters, you have to specify more parameters like this for lpar to try to figure out an available interface by ping operation:

```
getmacs pnode1 -D -S 192.168.0.1 -G 192.168.0.10
```

The output looks like following:

```
pnode1:
```

Type	Location Code	MAC Address	Full Path Name	Ping Result
Device Type				

ent U9133.55A.10E093F-V4-C5-T1	f2:60:f0:00:40:05	/vdevice/1-lan@30000005	virtual	
--------------------------------	-------------------	-------------------------	---------	--

And the Mac address will be written into the xCAT mac table. Run to verify:

```
tabdump mac
```

3.6.7. Setup dhcp service

3.6.7.1. Setup the dhcp listen interfaces in site table

```
chdef -t site dhcpinterfaces='pmanagenode|eth1'
```

3.6.7.2. [SLES] Check the installation of dhcp-server

On the SLES management node, the dhcp-server rpm may not have been automatically installed. Use following command to check whether it has been installed:

```
rpm -qa | grep -E "^dhcp-server"
```

If it is not installed, installed it manually:

```
zypper install dhcp-server
```

3.6.7.3. Configure the DHCP

Add the relevant networks into the DHCP configuration:

```
makedhcp -n
```

Add the defined nodes into the DHCP configuration:

```
makedhcp -a
```

Restart the dhcp service:

```
service dhcpd restart
```

Note: Please make sure there is only one dhcpd server can serv these compute nodes.

4. Install a Compute Node

4.1. Prepare the installation source

You can use the iso file of the installed OS to extract the installation files. For example, you have a iso file /iso/RHEL5.2-Server-20080430.0-ppc-DVD.iso

```
copycds /iso/RHEL5.2-Server-20080430.0-ppc-DVD.iso
```

Note: If you encounter the issue that the iso cannot be mounted by the copycds command. Make sure the SELinux is disabled.

4.2. Statefull Node installation

4.2.1. Customize the install profile

xCAT uses KickStart or AutoYaST installation templates and related installation scripts to complete the installation and configuration of the compute node.

You can find sample templates for common profiles in following directory:

```
/opt/xcat/share/xcat/install/<os>/
```

If you customize a template then you should copy it to the /install/custom/install/<os> directory.

4.2.1.1. Search order for installation templates

The profile, os and architecture of the node was setup in Set the type attributes of the node above.

To check the setting of your node's profile, os, architecture run:

```
lsdef pnode1

Object name: pnode1
.
.
.
  arch=ppc64
  os=rhels5.5
  profile=compute
```

For this example, the search order for the template file is as follows:

```
compute.rhels5.5.ppc64.tpl
compute.rhels5.ppc64.tpl
compute.rhels.ppc64.tpl
compute.rhels5.5.tpl
compute.rhels5.tpl
compute.rhels.tpl
compute.ppc64.tpl
compute.tpl
```

The directory */install/custom/<install>/<os>* will be searched first and then search */opt/xcat/share/xcat/<install>/<os>*.

Check the directories to see which template will match.

If you want to customize a template for node , you can copy the template to the */install/custom/install/<os>/* directory and make your modifications. You need to copy to the “custom” directory, so the next install of xCAT will not wipe out your modifications. Keep in mind the above search order to make sure it picks up your template.

Note: Sometimes the directory */opt/xcat/share/xcat/install/scripts* also needs to be copied to */install/custom/install/* to make the customized profile work, because the customized profiles will need to include the files in scripts directory as the prescripts and postscripts. For example, you need to put the *.otherpkgs.pkglist* file into the */install/custom/install/<os>/* directory, if you need to install other packages.

4.2.1.2. Install other specific packages

If you want to install the specific package like `specific.rpm` onto the compute node, copy the `specific.rpm` into the following directory:

```
/install/post/otherpkgs/<os>/<arch>
```

Another thing you HAVE TO do is to create `repodata` for this directory. You can use the “`createrepo`” command to create `repodata`.

On RHEL5.x, the “`createrepo`” rpm package can be found in the install ISO; on SLES11, it can be found in SLE-11-SDK-DVD Media 1 ISO.

After “`createrepo`” is installed, you need to create one text file which contains the complete list of files to include in the repository. For example, the name of the text file is “`rpms.list`” in “`/install/post/otherpkgs/<os>/<arch>`” directory, please follow this command to include the list into “`rpms.list`”:

```
cd /install/post/otherpkgs/<os>/<arch>
ls *.rpm >rpms.list
```

Or, if you create one sub-directory to contain the rpm packages, for example, one directory named “`other`” in “`/install/post/otherpkgs/<os>/<arch>`”. Please run the following commands to add the rpm packages under “`other`” directory to this repo:

```
cd /install/post/otherpkgs/<os>/<arch>
ls other/*.rpm >>rpms.list
```

Finally, run the following command to create `repodata` for the directory:

```
createrepo -i rpms.list /install/post/otherpkgs/<os>/<arch>
```

The “`createrepo`” command with “`-i rpms.list`” option will create the repository for the rpm packages listed in the “`rpms.list`” file. It won't destroy and affect the rpm packages which is in the same directory but have been included into another repository.

4.2.2. Set the node status to ready for installation

```
nodeset pnode1 install
```

4.2.3. Use network boot to start the installation

```
rnetboot pnode1
```

4.2.4. Check the installation results

1. After the node installation is completed successfully, the node's status will be changed to “`booted`”, the following command to check the node's status:

```
lsdef pnode1 -i status
```

2. When the node's status is changed to “`booted`”, you can also check ssh service on the node is working and you can login without password.

Note: Do not run ssh or xds against the node until the node installation is completed successfully. Running ssh or xds against the node before the node installation completed may result in ssh hostkeys issues.

3. If ssh is working but cannot login without password, force exchange the ssh key to the compute node using xdsh:

```
xdsh pnode1 -K
```

After exchanging ssh key, following command should work.

```
xdsh pnode1 date
```

4.2.5. Install a new Kernel on the nodes

Using a postinstall script (you could also use the updatenode method):

```
mkdir /install/postscripts/data  
cp <kernel> /install/postscripts/data
```

Create the postscript updatekernel:

```
vi /install/postscripts/updatekernel  
#!/bin/bash  
rpm -Uivh data/kernel-*rpm  
chmod 755 /install/postscripts/updatekernel
```

Add the script to the postscripts table

```
chdef -p -t group -o compute postscripts=updatekernel
```

```
rinstall compute
```

4.3. Stateless node installation

4.3.1. Generate the stateless image for compute node

Typically, you can build your stateless compute node image on the Management Node, if it has the same OS and architecture as the node. If you need another OS image or architecture than the OS installed on the Management Node, you will need a machine that meets the OS and architecture you want for the image and create the image on that node.

4.3.1.1. Make the compute node add/exclude packaging list

The default list of rpms to added or exclude to the diskless images is shipped in /opt/xcat/share/xcat/netboot/<os> directory. If you want to modify the current defaults for *.pkglist or *.exlist or *.postinstall, copy the shipped default lists to the /install/custom/netboot/<os> directory, so your modifications will not be removed on the next xCAT rpm update. xCAT will first look in the custom directory for the files before going to the share directory.

If you want to exclude more packages, add them into the following exlist file:

/install/custom/netboot/<os>/<profile>.exlist

Add more packages names that need to be installed on the stateless node into the pkglist file

/install/custom/netboot/<os>/<profile>.pkglist

4.3.1.2. Setting up postinstall files

There are rules (release 2.4 or later) for which * postinstall files will be selected to be used by genimage.

Copy the appropriate /opt/xcat/share/xcat/netboot/<OS>/*.postinstall file to the /install/custom/netboot/<os> directory, if you are going to make modifications.

Use these basic rules to edit the correct file in the “/install/custom/netboot/<OS>” directory. The rule allows you to customize your image down to the profile, os and architecture level, if needed.

You will find *.postinstall files of the following forms:

- 1) <profile>.<os>.<arch>.postinstall
- 2) <profile>.<arch>.postinstall
- 3) <profile>.<os>.postinstall
- 4) <profile>.postinstall

genimage will select the *.postinstall file from the list , if it exist, in the order 1-4.

This means, if “<profile>.<os>.<arch>.postinstall” is there, it will be used .

If there is no such a file, then the “<profile>.<arch>.postinstall” file will be used.

If there's no such a file , then the “<profile>.<os>.postinstall” file will be used.

If there is no such file, then it will use “<profile>.postinstall”.

Make sure you have the basic postinstall script setup in the directory to run for your genimage. The one shipped will setup fstab and rcons to work properly and is required. You can add more postinstall process ,if you want. The basic postinstall script (2.4) will be named <profile>.<arch>.postinstall (e.g. compute.ppc64.postinstall). You can create one for a specific os by copying the shipped one to , for example, compute.rhels5.4.ppc64.postinstall

Note: you can use the sample here: /opt/xcat/share/xcat/netboot/<os>/

[RH]:

Add following packages name into the <profile>.pkglist

```
bash
nfs-utils
stunnel
dhclient
kernel
openssh-server
openssh-clients
busybox-anaconda
```

```
wget
vim-minimal
ntp
```

You can add any other packages that you want to install on your compute node. For example, if you want to have users with passwords you should add the following:

```
cracklib
libuser
passwd
```

[SLES11]:

Add following packages name into the <profile>.pkglist

```
aaa_base
bash
nfs-utils
dhcpcd
kernel
openssh
psmisc
wget
sysconfig
syslog-ng
klogd
vim
```

4.3.1.1. Run image generation

[RHEL]:

```
cd /opt/xcat/share/xcat/netboot/rh
./genimage -i eth0 -n ibmveth -o rhels5.2 -p compute
```

[SLES11]:

```
cd /opt/xcat/share/xcat/netboot/sles
./genimage -i eth0 -n ibmveth -o sles11 -p compute
```

4.3.1.2. Pack the image

[RHEL]:

```
packimage -o rhels5.2 -p compute -a ppc64
```

[SLES]:

```
packimage -o sles11 -p compute -a ppc64
```

4.3.2. Set the node status ready for network boot

```
nodeset pnode2 netboot
```

4.3.3. Use network boot to start the installation

```
rnetboot pnode2
```

4.3.4. Check the installation result

1. After the node installation is completed successfully, the node's status will be changed to “booted”, the following command to check the node's status:

```
lsdef pnode1 -i status
```

2. When the node's status is changed to “booted”, you can also check ssh service on the node is working and you can login without password.

Note: Do not run ssh or xdsh against the node until the node installation is completed successfully. Running ssh or xdsh against the node before the node installation completed may result in ssh hostkeys issues.

3. If ssh is working but cannot login without password, force exchange the ssh key to the compute node using xdsh:

```
xdsh pnode1 -K
```

After exchanging ssh key, following command should work.

```
xdsh pnode1 date
```

4.3.5. Installing a new Kernel in the stateless image

Obtain you new kernel and kernel modules on the MN.

Copy the kernel into /boot :

```
cp vmlinux-2.6.32.10-0.5-ppc64 /boot
```

Copy the kernel modules into /lib/modules/<new kernel directory>

```
xcatlinuxmn:/lib/modules # ls -al
total 16
drwxr-xr-x  4 root root 4096 Apr 19 10:39 .
drwxr-xr-x 17 root root 4096 Apr 13 08:39 ..
drwxr-xr-x  3 root root 4096 Apr 13 08:51 2.6.32.10-0.4-ppc64
drwxr-xr-x  4 root root 4096 Apr 19 10:12 2.6.32.10-0.5-ppc64
```

Run genimage to update the image with the new kernel

```
genimage -i eth0 -n ibmveth -o sles11.1 -p compute -k 2.6.32.10-0.5-ppc64
```

```
packimage -o sles11 -p compute -a ppc64
```

Run nodeset command and netboot.

```
uname -a shows the new kernel.
```

4.3.6. Remove an image

If you want to remove an image, there's one command named “rmimage”, which is used to remove the Linux stateless or statelite image from the file system.

You can specify the <os>, <arch> and <profile> value to the “rmimage” command:

```
rmimage -o <os> -a <arch> -p <profile>
```

Or, you can specify one imagename to the command:
`rmimage <imagename>`

5. Firmware upgrade

5.1. Requirements

POWER5 and POWER6 Licensed Internal Code updates must meet the following prerequisites:

5.1.1. Enable the HMC to allow remote ssh connections.

[AIX]

Ensure that ssh is installed on the AIX xCAT management node. If you are using an AIX management node, make sure the value of "useSSHonAIX" is "yes" in the site table.

```
chtab key="useSSHonAIX" site.value=yes
```

5.1.2. Define the necessary attributes

The Lpar , CEC, or BPA has been defined in the nodelist, nodehm, nodetype, vpd, ppc tables.

5.1.3. Define the HMC as a node

Define the HMC as a node on the management node. For example,

```
nodeadd hmc01.clusters.com groups=hmc
```

5.1.4. Setup SSH connection to HMC

Run the rspconfig command to set up and generate the ssh keys on the xCAT management node and transfer the public key to the HMC. You must also manually configure the HMC to allow remote ssh connections. For example:

```
rspconfig hmc01.clusters.com sshcfg=enable
```

5.1.5. Get the Microcode update package and associated XML file

Download the Microcode update package and associated XML file from the IBM Web site: <http://www14.software.ibm.com/webapp/set2/firmware/gjsn>.

5.2. Perform Firmware upgrade for CEC on P5/P6

5.2.1. Define the CEC as a node on the management node

Update the xCAT required xCAT tables:

Modify the nodelist table

```
nodeadd Server-m_tmp-SNs_tmp groups=hmc,all
```

Modify the table nodehm

```
chtab node="Server-m_tmp-SNs_tmp" nodehm.mgt="hmc"
```

Modify the table nodetype:

```
chtab node="Server-m_tmp-SNs_tmp" nodetype.nodetype="fsp"
```

Modify the table ppc:

```
chtab node="Server-m_tmp-SNs_tmp" ppc.hcp=  
hmc01.clusters.com
```

Modify the tab vpd:

```
chtab node=Server-m_tmp-SNs_tmp vpd.serial=s_tmp  
vpd.mtm=m_tmp
```

Set the account of the HMC(Modify the ppchcp):

```
chtab hcp=hmc01.clusters.com ppchcp.username=hscroot  
ppchcp.password=abc123
```

5.2.2. Setup SSH connection to HMC

Generate the ssh keys on the xCAT management node and transfer the public key to the HMC to configure the HMC to allow remote ssh connections.

```
rspconfig hmc01.clusters.com sshcfg=enable
```

5.2.3. Check firmware level

```
rinv Server-m_tmp-SNs_tmp firm
```

5.2.4. Update the firmware

Download the Microcode update package and associated XML file from the IBM Web site: <http://www14.software.ibm.com/webapp/set2/firmware/gjsn>. Create the /tmp/fw directory, if necessary, and copy the downloaded files to the /tmp/fw directory.

Run the rflash command with the --activate flag to specify the update mode to perform the updates. (Please see the “rflash” manpage for more information)

```
rflash Server-m_tmp-SNs_tmp -p /tmp/fw --activate  
disruptive
```

NOTE: You Need check your update is concurrent or disruptive here!! other commands sample:

```
rflash Server-m_tmp-SNs_tmp -p /tmp/fw --activate  
concurrent
```

Notes:

- 1)If the noderange is the group lpar, the upgrade steps are the same as the CEC's.*
- 2)System p5 and p6 updates can require time to complete and there is no visual indication that the command is proceeding.*

5.3. Perform Firmware upgrades for BPA on P5/P6

5.3.1. Define the BPA as a node on the management node

Update the xCAT tables:

Modify the nodelist table. Define the BPA as a node

```
nodeadd Server-m_tmps_tmp groups=hmc,all
```

Modify the table nodehm

```
chtab node="Server-m_tmps_tmp" nodehm.mgt="hmc"
```

Modify the table nodetype:

```
chtab node="Server-m_tmps_tmp" nodetype.nodetype="fsp"
```

Modify the table ppc:

```
chtab node="Server-m_tmps_tmp" ppc.hcp= hmc01.clusters.com  
ppc.id=x
```

Modify the tab vpd:

```
chtab node=Server-m_tmps_tmp vpd.serial=s_tmp vpd.mtm=m_tmp
```

Set the account of the HMC(Modify the ppchcp):

```
chtab hcp=hmc01.clusters.com ppchcp.username=hscroot  
ppchcp.password=abc123
```

5.3.2. Setup SSH connection to HMC

Generate the ssh keys on the xCAT management node and transfer the public key to the HMC to configure the HMC to allow remote ssh connections.

```
rspconfig hmc01.clusters.com sshcfg=enable
```

5.3.3. User rinv to check the firmware level (see rinv manpage)

```
rinv Server-m_tmps_tmp firm
```

5.3.4. Update the firmware

Download the Microcode update package and associated XML file from the IBM Web site:

<http://www14.software.ibm.com/webapp/set2/firmware/gjsn>

Create the /tmp/fw directory, if necessary, and copy the downloaded files to the /tmp/fw directory.

Run the rflash command with the --activate flag to specify the update mode to perform the updates.

```
rflash Server-m_tmps_tmp -p /tmp/fw --activate disruptive
```

NOTE: You Need check your update is concurrent or disruptive here!! other commands sample:

```
rflash Server-m_tmps_tmp -p /tmp/fw --activate concurrent
```

5.4. Commit currently activated LIC update(copy T to P) for a CEC/BPA on p5/p6

5.4.1. Check firmware level

Refer to the environment setup in the section 'Firmware upgrade for CEC on P5/P6' to make sure the firmware version is correct.

5.4.2. Commit the firmware LIC

Run the rflash command with the `-commit` flag.

```
rflash Server-m_tmp-SNs_tmp --commit
```

Notes:

(1) If the noderange is Lpar, the commit steps are the same as the CEC's.

(2) When the `-commit` or `--recover` two flags is used, the noderange cannot be BPA . It only can be CEC or LPAR ,and will take effect for both managed systems and power subsystems.

6. Advanced features

6.1. Use the driver update disk:

Refer to the “Use the driver update disk” section in the document [xCAT-iDpx.pdf](#).

7. References

- xCAT web site: <http://xcat.sf.net/>
- xCAT man pages: <http://xcat.sf.net/man1/xcat.1.html>
- xCAT DB table descriptions: <http://xcat.sf.net/man5/xcatdb.5.html>
- Monitoring Your Cluster with xCAT:
<http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2-Monitoring.pdf>
- xCAT on AIX Cookbook: <http://xcat.svn.sourceforge.net/viewvc/xcat/xcat-core/trunk/xCAT-client/share/doc/xCAT2onAIX.pdf>
- xCAT wiki: <http://xcat.wiki.sourceforge.net/>

- xCAT mailing list: <http://xcat.org/mailman/listinfo/xcat-user>
- xCAT bugs: https://sourceforge.net/tracker/?group_id=208749&atid=1006945
- xCAT feature requests: https://sourceforge.net/tracker/?group_id=208749&atid=1006948