

Node Provisioning: How do I (re)install an operating system on a compute node disk, without losing data on the other disks?

How do I (re)install an operating system on a regular node (compute node) disk, without losing data on the other disks of that node?

Sometimes an administrator may want to install the operating system onto one of the drives of a regular (compute) node, while preserving the data on the other drives of that node.

An example of this kind of situation is when the cluster has nodes with Ceph OSD, and the administrator wants to just replace the disk where the operating system is installed, while keeping the data on the other disks intact.

The procedure described here works for Bright Cluster Manager version 7.3 or higher. An example case is used to simplify this article. The actual cluster an administrator will configure is very likely to differ from the example, so appropriate changes to the procedure should be carried out, instead of simply following the exact procedure as described here. If something is not clear, then further guidance is available from Bright Computing Support.

Here is what is assumed for the example:

- The node in which the operating system will be reinstalled is node001.
- `/dev/sda` has the operating system. The `/` (root) directory is mounted on the `/dev/sda1` partition within the device.
- `/dev/sdb` has the important data that has to be preserved, so no FULL install should be performed on it, ever.
- `/dev/sdc` is another device already installed in the server. The administrator wants to do the (re)install of the operating system onto this device (the `/` (root) directory will be mounted in the `/dev/sdc1` partition)

The preceding assumptions can be verified by looking at the disk setup. The XML specification of the disk setup can be viewed from cmsh, by running the following commands on the head node:

```
# cmsh
```

```
% device use node001
```

```
% get disksetup
```

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<diskSetup>
```

Node Provisioning: How do I (re)install an operating system on a compute node disk, without losing data on the other disks?

```
<device>
```

```
<blockdev>/dev/sda</blockdev>
```

```
<partition id="a1">
```

```
<size>max</size>
```

```
<type>linux</type>
```

```
<filesystem>xf</filesystem>
```

```
<mountPoint>/</mountPoint>
```

```
<mountOptions>defaults,noatime,nodiratime</mountOptions>
```

```
</partition>
```

```
</device>
```

```
<device>
```

```
<blockdev>/dev/sdb</blockdev>
```

```
<partition id="b1">
```

```
<size>max</size>
```

```
<type>linux</type>
```

```
<filesystem>xf</filesystem>
```

```
<mountPoint>/data</mountPoint>
```

```
<mountOptions>defaults,noatime,nodiratime</mountOptions>
```

```
</partition>
```

```
</device>
```

```
</diskSetup>
```

Node Provisioning: How do I (re)install an operating system on a compute node disk, without losing data on the other disks?

The XML file shows:

```
the / directory is mounted on /dev/sda1
```

and

```
the /data directory (data that must be preserved) is mounted on /dev/sdb1.
```

Now the block devices recognized by the operating system can be verified with GNU parted:

```
# ssh node001
```

```
[root@node001 ~]# parted
```

```
GNU Parted 3.1
```

```
Using /dev/sda
```

```
Welcome to GNU Parted! Type 'help' to view a list of commands.
```

```
(parted) print devices
```

```
/dev/sda (21,5GB)
```

```
/dev/sdb (1074MB)
```

```
/dev/sdc (21,5GB)
```

```
(parted) quit
```

It is very important that the administrator knows the correct device name of the device in which the operating system will be installed, because in the steps that follow, the data on that device will be overwritten.

The disk layout XML file of the node can be backed up by running this command on the head node:

```
# cmsh -c "device use node01; get disksetup" > disksetup-node01.xml
```

Node Provisioning: How do I (re)install an operating system on a compute node disk, without losing data on the other disks?

The `datanode` property must be set to `yes` for `node001`. This ensures that the administrator will always be asked for a confirmation before a `FULL` install. Its value can be checked with:

```
# cmsh -c "device use node02; get datanode"
```

If the property is set to `no`, or if it is not necessary to set it for this node, then it is a good idea that the administrator sets it to `yes` just for this procedure, and then sets it to `no` again after the procedure has been done. In that case, if a mistake is made in the procedure, a `FULL` install will only be carried out with an explicit confirmation.

The `disksetup` property of `node001` is modified to include all the disks (the old and the new). The administrator must make sure that the definition of the partitions is exactly the same as it was before. The only thing the administrator should change, is that the `/` directory will be mounted on the new disk instead of on the old one. The administrator has to refer to the device by its device name; it is not possible to specify the UUID. For this example, the layout could be edited like this:

```
# cmsh
```

```
# device use node001
```

```
# set disksetup
```

```
<?xml version="1.0" encoding="UTF-8"?>
```

```
<diskSetup>
```

```
<device>
```

```
<blockdev>/dev/sda</blockdev>
```

```
<partition id="a1">
```

```
<size>max</size>
```

```
<type>linux</type>
```

```
<filesystem>xfs</filesystem>
```

Node Provisioning: How do I (re)install an operating system on a compute node disk, without losing data on the other disks?

```
<mountPoint>/old</mountPoint>
```

```
<mountOptions>defaults,noatime,nodiratime</mountOptions>
```

```
</partition>
```

```
</device>
```

```
<device>
```

```
<blockdev>/dev/sdb</blockdev>
```

```
<partition id="b1">
```

```
<size>max</size>
```

```
<type>linux</type>
```

```
<filesystem>xfs</filesystem>
```

```
<mountPoint>/data</mountPoint>
```

```
<mountOptions>defaults,noatime,nodiratime</mountOptions>
```

```
</partition>
```

```
</device>
```

```
<device>
```

```
<blockdev>/dev/sdc</blockdev>
```

```
<partition id="c1">
```

```
<size>max</size>
```

```
<type>linux</type>
```

```
<filesystem>xfs</filesystem>
```

```
<mountPoint>/</mountPoint>
```

Node Provisioning: How do I (re)install an operating system on a compute node disk, without losing data on the other disks?

```
<mountOptions>defaults,noatime,nodiratime</mountOptions>
```

```
</partition>
```

```
</device>
```

```
</diskSetup>
```

The changes must be committed after editing the layout.

The `blockdevicesclearedonnextboot` property of `node001` must be set to the device of the new disk. In this example it would be set like this:

```
# cmsg
```

```
% device use node001
```

```
% set blockdevicesclearedonnextboot /dev/sdc
```

```
% commit
```

A reboot of `node001` is then done. If everything was done correctly, Bright will wipe out and repartition the new disk, and do a `SYNC` install that will copy the contents of the software image to the disk.

Unique solution ID: #1414

Author: Frank Furter

Last update: 2018-03-12 15:31