

Installing BCM: How do I install Bright Cluster Manager on an SGI ICE X/XA system?

Installing Bright on an SGI ICE X/XA system involves first installing the "Admin" node using a Bright Cluster Manager ISO. The ISO could be burned to a DVD, could be loaded from a USB drive, or could be mounted using the BMC's virtual media option.

Full details on all steps of the head node installation procedure can be found in the Installation Manual (which includes a Quickstart Installation Guide). When SGI is selected as the hardware vendor during the head node installer, an SGI specific package will be included in the head node installation. The package (cluster-tools-sgi) could also be installed from the Bright YUM repository at a later stage if necessary.

Normally on an SGI ICE X/XA system, the first NIC is configured on the external network and the second NIC is configured on the internal management network.

Once the admin node has been installed as a Bright head node and is fully booted, a number of additional steps will have to be performed to allow the admin node to access the rest of the system.

To make Bright Cluster Manager fully functional a license needs to be installed. The procedure to request and install a license can be found in the Installation Manual.

Assuming the cluster-tools-sgi package has been installed on the head node, there will be an environment module called cluster-tools-sgi that can be used to get easy access to the commands that are provided by this package. To load the environment module, one can use:

```
module load cluster-tools-sgi
```

Every SGI ICE X/XA is delivered with a manufacturing configuration file, which is normally called mfgconfigfile. This file must be copied onto the admin node (e.g. using scp). The Bright cluster management infrastructure will have to be initialized using the information that is provided in the manufacturing configuration file. The following command can be used to accomplish this:

```
sgi-basecfg2cmd --configfile=/root/mfgconfigfile
```

Once this is done, the rack leaders and service nodes can be started. If for some reason the rack leaders and service nodes are already powered on, power them off first. The following

Installing BCM: How do I install Bright Cluster Manager on an SGI ICE X/XA system?

command will use the MAC address that is associated with each device in the manufacturing configuration to find the IP address that has been assigned to the BMC of each rack leader and service node. The command will then issue a power on operation to the BMC of each of the rack leaders and service nodes.

```
sgi-poweronmainnodes
```

It will take a couple of minutes for the rack leaders and service nodes to be started. Progress can be checked by using the device status command in cmesh. Once all the rack leaders and service are in the UP state, it is time to discover the chassis management controllers (CMCs). To do this, we must run a discovery command on each of the rack leaders. The command will listen for DHCP requests from the CMCs to discover their MAC addresses. The following command must be executed for each rack leader. For example, on 3 rack system, the commands would be:

```
sgi-run-cmcdhcpmonitor --rack=1  
sgi-run-cmcdhcpmonitor --rack=2  
sgi-run-cmcdhcpmonitor --rack=3
```

The CMC information that was collected using the sgi-cmcdhcpmonitor script needs to be added to the Bright cluster management infrastructure. The following command be used to do this for each rack. For our 3 rack system, the commands would be:

```
sgi-cmccfg2cmd --configfile=/root/cmc-dhcp-monitor-r01.log  
sgi-cmccfg2cmd --configfile=/root/cmc-dhcp-monitor-r02.log  
sgi-cmccfg2cmd --configfile=/root/cmc-dhcp-monitor-r03.log
```

After the CMCs have been added to the cluster management infrastructure, it is time to enable network booting on the internal rack networks. This will cause a DHCP and TFTP server to be

Installing BCM: How do I install Bright Cluster Manager on an SGI ICE X/XA system?

started on each rack leader. Initially this will just cause the CMCs to be assigned an IP address using DHCP. On our 3 rack system, the commands would be:

```
sgi-enablerlnetworkboot --rack=1  
sgi-enablerlnetworkboot --rack=2  
sgi-enablerlnetworkboot --rack=3
```

Now that the CMCs have been assigned IP addresses, we can query the CMCs for compute blade and IB switch information. The following command can be used to query a number of CMCs for compute blade and IB switch information. For our 3 rack system, we could use the following commands:

```
sgi-cmcinfo2cmd --rack=1  
sgi-cmcinfo2cmd --rack=2  
sgi-cmcinfo2cmd --rack=3
```

Once all of the compute blades and IB switches have been added to the cluster management infrastructure, we can power them on. This operation could normally be performed using Bright's CSMH or CMGUI, but because the BMC IP addresses are not known yet, we have to run a special command to do the initial power-on operation. This command will determine the MAC address of the BMC of each blade, and will then use the rack leader ARP table to figure out what IP has been assigned to this BMC. A power on operation is then issued to the BMC. For our 3 rack system, we could use the following commands:

```
sgi-poweroncomputenodes --rack=1  
sgi-poweroncomputenodes --rack=2  
sgi-poweroncomputenodes --rack=3
```

Once all of the compute blades and IB switches have been powered on, it will take a couple of minutes for all of the blades to be booted. The device status command in cmsh can be used to determine which blades are UP and DOWN.

Installing BCM: How do I install Bright Cluster Manager on an SGI ICE X/XA system?

Once all of the blades are in the UP state, the system is ready for business. For more information on how to add users to the system, and how to submit jobs, please refer to the Bright Administrator Manual and Bright User Manual respectively.

In rare cases it might be possible that the CMCs do not report the information for all compute blades and/or IB switches in the enclosures. For our 3 rack system, the CMCs can be re-initialized with the following commands:

```
sgi-reinitcmcs --rack=1  
sgi-reinitcmcs --rack=2  
sgi-reinitcmcs --rack=3
```

The following command can be used again to query a number of CMCs for updated compute blade and IB switch information. For our 3 rack system, we could use the following commands again, which will add the not yet configured compute blades and/or IB switches:

```
sgi-cmcinfo2cmd --rack=1  
sgi-cmcinfo2cmd --rack=2  
sgi-cmcinfo2cmd --rack=3
```

Unique solution ID: #1305

Author: Martijn de Vries

Last update: 2016-07-05 09:51