

Burn: How can I flash update all my nodes (from DOS)?

How can I flash update all my nodes (from DOS)?

(For update via Linux, see

<http://kb.brightcomputing.com/faq/index.php?action=artikel&cat=11&id=105>)

The problem:

I have to perform some maintenance on the nodes of my system, that requires that I boot via PXE, only once, a special image (i.e. a DOS image loaded with ramdisk and some script to update the BIOS).

I don't want to write some script that looks into log files and changes the dhcp configuration files on the fly to do this sort of thing. Is there a handy way to ask Bright to boot with such an image only once?

Two solutions:

Bright has at least two options to do just that. Depending on what your use-case is, you'll have to choose if you want to use the burn-framework approach, or the approach where you set the 'pxelabel' property manually.

1. The Burn framework method

It's part of the burn in test framework that is installed on each Bright cluster. The burn in framework is covered in general in the Admin manual , currently in Appendix N of Version 6. The latest version of the manual is always available from <http://info.brightcomputing.com/manual-download>

To aid flashing there is a test called bios-manager. Its configuration will be covered here.

The test can be used to automatically flash different images into different mainboard models (for example).

For this example let's assume we have two different models, called "boardA" and "boardB".

First of all there need to be custom images that are supposed to do the actual flashing. These are usually DOS images. The images need to be placed in

Burn: How can I flash update all my nodes (from DOS)?

/cm/images/default-image/boot/bios

(Replace "default-image" with your software image name if needed.) So, for this example, we might have

/cm/images/default-image/boot/bios/boardA.img

and

/cm/images/default-image/boot/bios/boardB.img

In that same directory, there is already a menu.conf file:

/cm/images/default-image/boot/bios/menu.conf

This needs to be modified. It should have an entry for each flash image you would like to use. Here are two entries for our example:

```
# BEGIN SECTION boardA
LABEL boardA
  KERNEL memdisk
  APPEND initrd=bios/boardA.img
  MENU LABEL This will flash boardA
  MENU HIDE
# END SECTION boardA
```

```
# BEGIN SECTION boardB
LABEL boardB
  KERNEL memdisk
  APPEND initrd=bios/boardB.img
  MENU LABEL This will flash boardB
  MENU HIDE
# END SECTION boardB
```

Burn: How can I flash update all my nodes (from DOS)?

It is important that all occurrences of boardX are within a section, so that in the section the begin and end markers as well as the label statement, are the same. And of course they should refer to the correct flash image.

In order to distinguish between different boards that may need different flashimages, the bios-manager test will run a set of "main board detection"-scripts. These are installed in

```
/cm/shared/apps/cmburn/bios-manager
```

and should be named:

```
mainboard_<custom tag>
```

When a mainboard script detects that it is running on a specific board that it knows, it should print the corresponding section tag from the menu.conf file. Otherwise it should print "0". To continue our example, create a file called:

```
/cm/shared/apps/cmburn/bios-manager/mainboard_boardA
```

and a file called

```
/cm/shared/apps/cmburn/bios-manager/mainboard_boardB
```

Here is an example of what /cm/shared/apps/cmburn/bios-manager/mainboard_boardA might look like:

```
#!/bin/bash

if dmidecode --type 2| grep "Product Name" | grep -q boardA; then
    echo "boardA"
else
    echo "0"
fi
```

Obviously, this is just an example, these scripts can be as elaborate as required.

To prevent accidents, the bios-manager test is disabled by default. To enable it, edit

Burn: How can I flash update all my nodes (from DOS)?

/cm/shared/apps/cmburn/bios-manager/bios-manager.conf. The default settings should be suitable for most cases, but the BIOS_UPDATE_NEW_MACHINES field should be set to YES, otherwise nothing will happen.

Unfortunately, there is a bug in /cm/shared/apps/cmburn/bios-manager/getinterface. This script is used to detect

which interface on the head nodes will serve the PXE images. This script is known to sometimes return the wrong interface. For now the recommended work around is to edit the script and have it return the correct interface

hard-coded. Just replace the existing code with an "echo eth0" statement, where eth0 is the interface which the head node uses to serve PXE images, tftp and dhcp etc.

After all of the above is set up, you can start an actual burn cycle to test it. To do so, setup a burn configuration. Using cmsg, the following steps could be used:

```
[temp60-centos5pbs]% partition use base
[temp60-centos5pbs->partition[base]]% burnconfigs
[temp60-centos5pbs->partition[base]->burnconfigs]% add flash
[temp60-centos5pbs->partition*[base*]->burnconfigs*[flash*]]% set xml
```

Enter into the editor and save-quit:

```
<burnconfig>
<pre-install>
  <phase name="flash">
    <test name="bios-manager"/>
  </phase>
</pre-install>
<post-install>
</post-install>
</burnconfig>
[temp60-centos5pbs->partition*[base*]->burnconfigs*[flash*]]% set description "flash the bios"
[temp60-centos5pbs->partition*[base*]->burnconfigs*[flash*]]% commit
[temp60-centos5pbs->partition[base]->burnconfigs[flash]]%
```

And then to test it on node001:

```
[temp60-centos5pbs->device[node001]]% burn start -o flash
Power reset nodes
[temp60-centos5pbs->device[node001]]%
custom ..... [ RESET ] node001
[temp60-centos5pbs->device[node001]]%
Tue Feb 5 12:56:04 2013 [notice] temp60-centos5pbs: node001 [ DOWN ] health check failed
[temp60-centos5pbs->device[node001]]%
Tue Feb 5 12:56:39 2013 [alert] node001: Check 'DevicesUp' is in state FAIL on node001
[temp60-centos5pbs->device[node001]]%
Tue Feb 5 12:56:55 2013 [notice] temp60-centos5pbs: node001 [ INSTALLING ] (node
```

Burn: How can I flash update all my nodes (from DOS)?

installer started)

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:56:55 2013 [notice] temp60-centos5pbs: node001 [ INSTALLING ] (running burn in tests)
```

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:57:28 2013 [notice] temp60-centos5pbs: node001
```

```
[INSTALLER_UNREACHABLE ] (running burn in tests)
```

<this is normal, the node is now rebooting into the dos image>

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:58:39 2013 [notice] temp60-centos5pbs: node001 [ INSTALLING ] (running burn in tests)
```

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:58:50 2013 [notice] temp60-centos5pbs: node001 [ INSTALLING ] (node installer started)
```

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:58:50 2013 [notice] temp60-centos5pbs: node001 [ INSTALLING ] (running burn in tests)
```

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:59:03 2013 [notice] temp60-centos5pbs: node001 [ INSTALLING ] (burn in test completed successfully)
```

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 12:59:36 2013 [notice] temp60-centos5pbs: node001 [INSTALLER_CALLINGINIT ] (switching to local root)
```

```
[temp60-centos5pbs->device[node001]]%
```

```
Tue Feb 5 13:00:15 2013 [notice] temp60-centos5pbs: node001 [ UP ]
```

It is strongly recommended to first test all of the above on a single node, and with a DOS image which does not actually do anything, before running the actual flash on a large number of nodes. Also, it is assumed that the flash images do a reset of the machine after flashing. And it is assumed that flashing will take at least ten seconds. So as an extra safety, you should add a sleep command to the DOS image. Two binaries installed on the cluster might come in handy here:

```
/cm/shared/apps/cmbios/nodebios/reset.com
```

and

```
/cm/shared/apps/cmbios/nodebios/sleep.exe
```

Burn: How can I flash update all my nodes (from DOS)?

2. The PXE MENULABEL method

It is also possible to control which PXE menu option a node defaults to by setting the 'pxelabel' property of a node. This way you could let a node, or all nodes, PXE boot into your DOS image. If you clear the 'pxelabel' property for your node(s), you can have them boot into the node-installer again.

The details of PXE MENU configuration are given in the BCM6 admin manual.

The configuration can be done like this:

A LABEL statement is built in `/cm/images/<default-image>/boot/pxelinux.cfg/default` file on the head node.

For example:

```
LABEL DOSIMAGE
MENU LABEL D^OSIMAGEFLASH
KERNEL dosimage
```

where `dosimage` is placed in `/cm/images/default-image/boot/`

Then, to enable a default boot to the `dosimage`, you can set this

```
cmsh -c "device; foreach -c default (set pxelabel DOSIMAGEFLASH); commit"
```

On reboot, the `dosimage` will load. To clear the `pxelabel` setting for the subsequent boot:

```
cmsh -c "device; foreach -c default (clear pxelabel); commit"
```

Burn: How can I flash update all my nodes (from DOS)?

Unique solution ID: #1093

Author: Frank Furter

Last update: 2013-03-22 12:49