

HOW TO MANAGE A CLUSTER

SIMONE ATZENI & MOHAMMED AL-MAHFOUDH

CONTENTS

1	Introduction	2
1.1	Simple example	2
1.2	What is next?	2
2	Sophisticated Clusterware	2
3	Simplified Admin/Clusterware (Webmin)	3
4	Webmin Setup and configuration walk through	4
5	Concluding remarks	9

LIST OF FIGURES

Figure 1	The webmin login page	4
Figure 2	Main page of the webmin console	4
Figure 3	Sub links of the webmin	5
Figure 4	The webmin servers index, through which you add all other cluster nodes using a single button click or one by one in case configurations differ on each node.	5
Figure 5	Adding a webmin server to the servers index	6
Figure 6	Cluster tab, of special interest “webmin servers” and “Cluster Shell Commands” sub-links	6
Figure 7	Cluster webmin servers link initially before populating it from the servers that were added to the webmin tab sublink “servers index”	7
Figure 8	After adding all webmin servers from the servers index to the cluster servers	8
Figure 9	Issuing custom commands in bulk to the whole/some of cluster nodes.	9

1 INTRODUCTION

Now that you have built your Pi cluster, now is the time to maintain it and administer its nodes. So, imagine instead of having a 3 nodes cluster that you suddenly became rich somehow, and you decided to build a 100 nodes Pi cluster. Assuming you figured out all the wiring and networking issues and you got it up and running. You decided after the whole hassle to install a package. How would you go about that? the naive would login to each node and starts issuing and re-issuing the command one node at a time. This is a nightmare situation. So in this chapter we will discuss how to effectively and efficiently administer both your Pi clusters and the industrial strength supercomputers.

1.1 Simple example

Lets for a second continue our example of installing “emacs” on all the 100 nodes in your cluster. If you spend around a minute to both issue the following command and then waiting for the package manager to both fetch and install it:

```
$ sudo apt-get install emacs
```

We would have waited at least 100 minutes (i.e. 1 hr and 40 minutes) just to install that alone!

This is apparently not the right way to do it. Imagine you have a much longer upgrade process per node... the nightmare continues to worsen.

1.2 What is next?

In a cluster, managing nodes on a node-by-node basis is not effective nor practical at all. So, we manage them in bulk using specialized software. Such software, are referred to cluster ware (a collection of resource managers, schedulers, and/or provisioning tools). We will know what each means shortly. For now, think of all of these as one package/program that does most of the tasks for you.

That software is mainly run from the main node and other nodes need to install only the helper processes to allow the main node to manage them. This management software can also be installed either in a node by node manner, or in bulk. We assume we made one disk image especially for the head node(s) ¹ and another disk image containing all software for the compute node(s) each containing the respective management software pre-installed. In our case it will be [Webmin](#).

2 SOPHISTICATED CLUSTERWARE

Webmin isn't primarily aimed to manage clusters, so there are more sophisticated software for each category of operations done in a cluster:

SCHEDULERS: Schedules jobs to run on clusters, i.e. time management of allocation of compute resources to specific jobs. As the number of users of a big cluster increases, this becomes cumbersome for a human intervention to happen and very error prone, and inefficient ².

¹ yes some clusters have more than one head node for backup in case one went down

² You know that clusters/supercomputers can consume megawatts of power and this is huge cost for the supercomputers maintainers

RESOURCE MANAGERS: Allocates exclusive/non-exclusive access to the cluster nodes for a specific duration of time, it starts, executes and monitors work (called a job), and manages the contention over specific resources. These are its roles rather than definition but you get the point.

OPERATING SYSTEM PROVISIONING: i.e. rebooting a number/percentage of nodes to a specific operating system according to tasks descriptions.

You can read a lot more about such software from the [wikipedia page](#). For our Pi clusters, we will use the previously mentioned **Webmin** due to many reasons some of which are the following:

1. Pi nodes have limited RAM on them, so we can't run management software than consumes much more than what's available on each node.
2. Sophisticated software require a lot of compute power since they run very sophisticated algorithms to figure out some problems and resolve them on the fly.
3. Sophisticated software also offers a lot of services that won't be viable to be used on our Pi clusters since, in all honesty, no more than two to three people would want to use at a time. Either way, Pi nodes aren't that powerful to support more than that.
4. Some of the top cluster ware, e.g. **Moab cluster suite** cost a lot of money for licenses, we think you better off buying more Pi nodes instead.

Our main reason to use webmin, however, is that it is super lightweight, runs perfectly and does exactly what we need for FREE! and without all the hassle to setup a gazillion services and configurations to run.

3 SIMPLIFIED ADMIN/CLUSTERWARE (WEBMIN)

In this section, we are explaining how to install webmin on all of the nodes following the same exact step on each (or of course setting it up before using the "dd" command to make the master image of the SD cards).

First, visit [webmin official downloads page](#), and download the .deb installation file. Copy to each node you need to install it on and execute the following command:

```
1 $ sudo dpkg -i path_to_webmin.deb
```

In case the above failed not being able to resolve dependencies, worry not. Just execute the following command and it will complete the installation:

```
1 $ sudo apt-get install -f
```

Now that we have it setup, it should be up and running as an HTTPS service on port 10000, as a super light web page. So, to access each node's webmin web interface you should enter the following URL in the web browser:

```
1 https://192.168.0.81:10000
```

The above indicates we are accessing Pi1 from the Pi cluster. Do the above per node (all nodes should have their own webmin installed and running). The good news is that it will run from boot time so no more configuration is needed to restart on reboots.

4 WEBMIN SETUP AND CONFIGURATION WALK THROUGH

After webmin has been setup and is running. In addition, all nodes of the cluster are visible to each other (assuming you configured the network correctly for them to be in the same subnet). We now add those peer webmin servers in the head node's webmin server index. To do this, point your browser to the following URL:

```
https://192.168.0.80:10000
```

For the username and password, they have to be in the sudo group. We will use the default username "pi" and password "raspberrypi" to login to the webmin console.

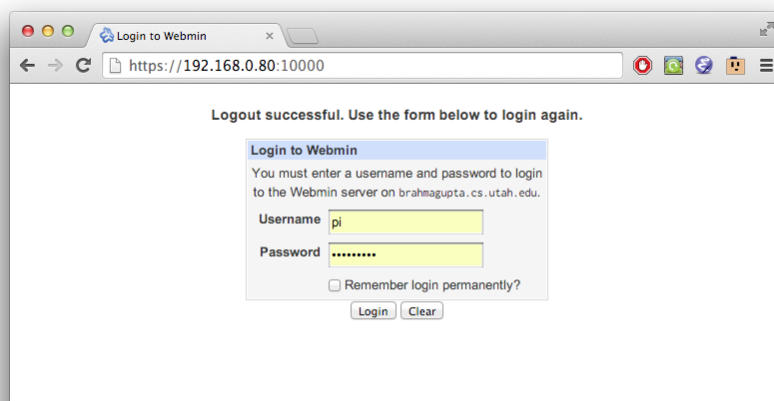


Figure 1: The webmin login page

After logging into the webmin console, you will be presented with the following web page.

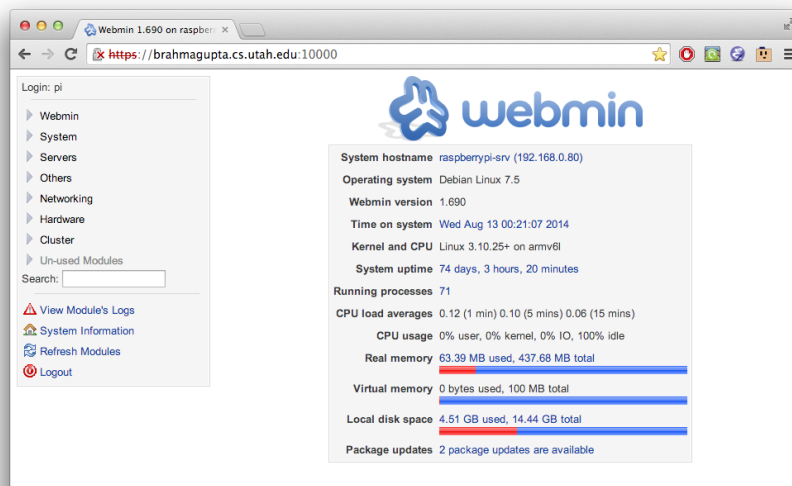


Figure 2: Main page of the webmin console

You can explore around and click links carefully to see what each link can do or where it can take you.

Each tab on the navigation bar on the left means something. For example, the tab “webmin” has to do with the webmin configuration and other webmin servers (as we will see and use it shortly). The “system” tab concerns operations on the same node you are logged in to through the webmin. Webmin, although simple, is also huge in terms of feature set so we will focus our guide on cluster support while leaving the remaining features to be explored by the curious. We should mention, though, that the most interesting tabs for clustering are two: “webmin” and “Cluster” tabs.

Once you click on the webmin tab, you will be presented with many sub links as shown below:

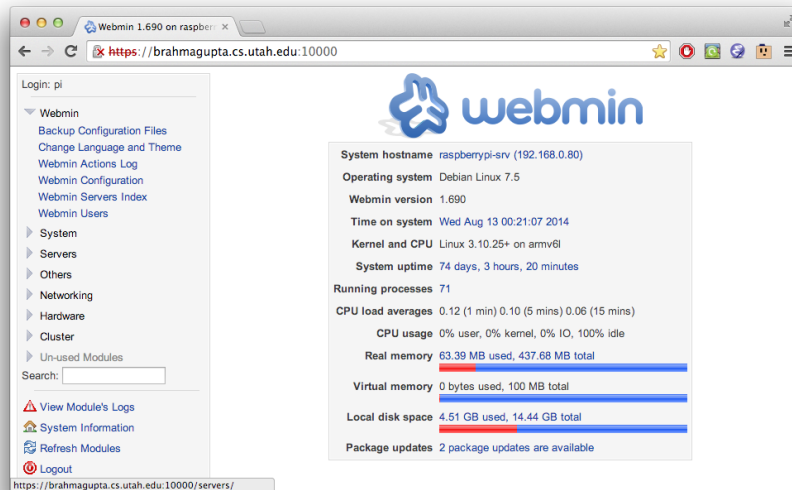


Figure 3: Sub links of the webmin

What we care most about is the link “Webmin Servers Index”. Click on it to be presented with the following screen, except you won’t see any servers added yet since you will do so now:



Figure 4: The webmin servers index, through which you add all other cluster nodes using a single button click or one by one in case configurations differ on each node.

We will assume that you have a different username and password pairs on each of the nodes and do everything manually. Otherwise a single button click should add all cluster nodes in one shot. To manually add each node, we click on the “Register a new server” link and fill all the details as shown in the screen shot below (assuming you have a node with IP address 192.168.0.81):

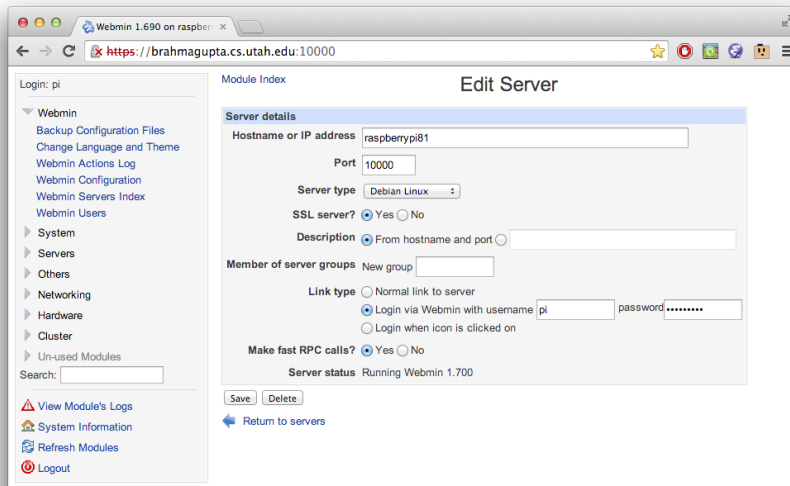


Figure 5: Adding a webmin server to the servers index

Do the above step for each of the nodes, including the head node. After adding all nodes to the index, now we can switch to the “Cluster” tab in webmin console:

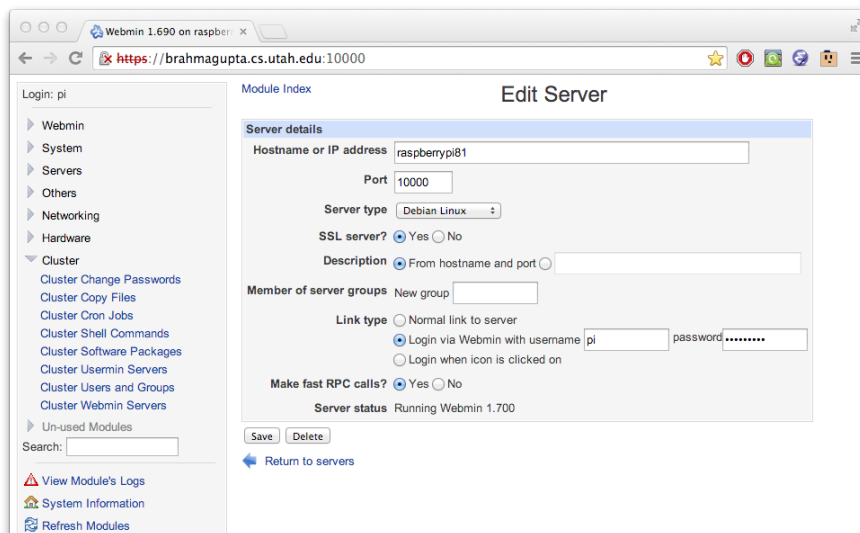


Figure 6: Cluster tab, of special interest “webmin servers” and “Cluster Shell Commands” sub-links

As indicated by Figure 6, we are most interested in “webmin servers” and “Cluster Shell Commands” sub-links of the cluster tab. Others are, of course, interesting but they are not tricky to figure out nor overly beneficial beyond what their names imply. So we will leave it for the curious explorer(s). For now, we click on the “Cluster Webmin Servers” to see the following screen.

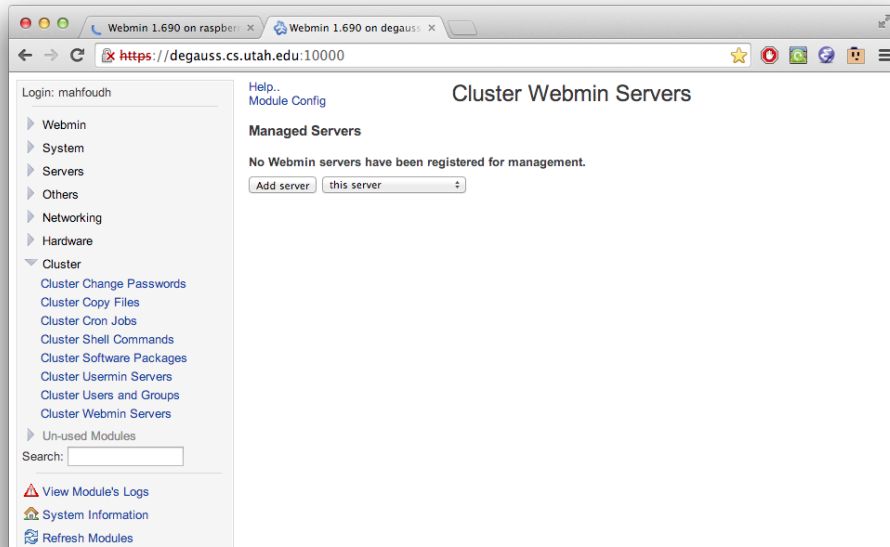


Figure 7: Cluster webmin servers link initially before populating it from the servers that were added to the webmin tab sublink “servers index”

Keep clicking on the “add server” button shown in the screen. Each time you click, the server that was added will be removed from this list till all are added then you will end up with a screen that looks like the following screen.

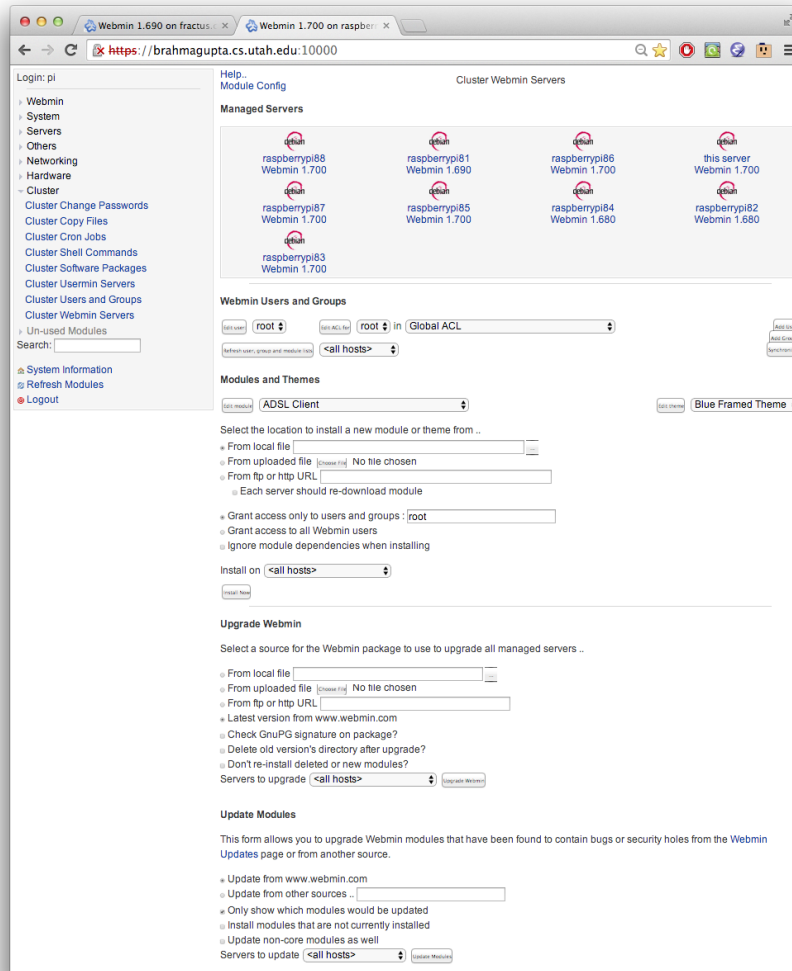


Figure 8: After adding all webmin servers from the servers index to the cluster servers

You may already have noticed that you can do the following from this very same screen: add users to all nodes all at once, synchronize users/groups, upgrade all webmin software on all of them at once in a single button click and so forth. While all of that of course is cool, some times you want to upgrade the software on each node. That is the software you installed and/or want to install more software. So the only way is to revert to command line. However, the good news is that you are not going to execute that command on each of the nodes your self, this time you will click on the “Cluster Shell Commands” and issue the command to all of them in bulk. See the following screen to know how to upgrade all software on all servers at once, note the “-y” to the package manager that is absolutely crucial not to abort the operation since webmin isn’t interactive with the command you issue. You can see in the screen shot below, I didn’t have to re-enter a previously entered command. Instead, I chose it from the drop down menu and all I have to do is to select which hosts are affected (or just choose <all hosts>).

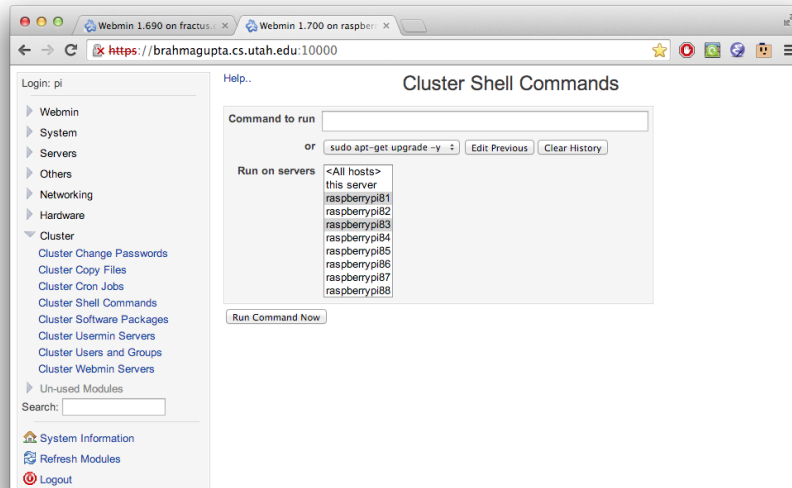


Figure 9: Issuing custom commands in bulk to the whole/some of cluster nodes.

5 CONCLUDING REMARKS

In this chapter we only gave the starting points for each of webmin's important features to the cluster. Webmin offers a lot more than that. The truth is that webmin replaces all GUI functionality for administering and maintaining your servers remotely and without the need to type a gazillion command line command/scripts all with a point and click manner, and some times fill up the gap. So, we hope that this is just the start for you to explore the power of webmin and similar resource managers/admin tools to expedite your job of cluster management and to do more in less time even in other contexts.

At the end, we wanted to give a nice hint based on our experience with webmin. If some thing breaks, just make sure you fix the head node first and then reboot if needed. After that, if webmin servers on other nodes are still running, broadcast your commands. Otherwise maybe a reboot of the whole cluster through either webmin or manual reboot should fix at least webmin, and then from there you can broadcast commands to these webmin nodes again.