

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

Next-generation High Availability Linux Clustering

Thomas Cameron, RHCA, RHCSS, RHCDS, RHCVA, RHCX
Chief Architect (Central US), Red Hat

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

AGENDA

Agenda

- What is clustering?
 - HA
 - Computational
- What we're going to discuss today
- Host installation
- Set up the iSCSI target (server)
- Set up the iSCSI initiators (clients)

Agenda

- Install the clustering software
- Start the cluster
- Configure STONITH
- Set up an active/passive Apache cluster
- Add a clustered filesystem
- Test

WHAT IS CLUSTERING?

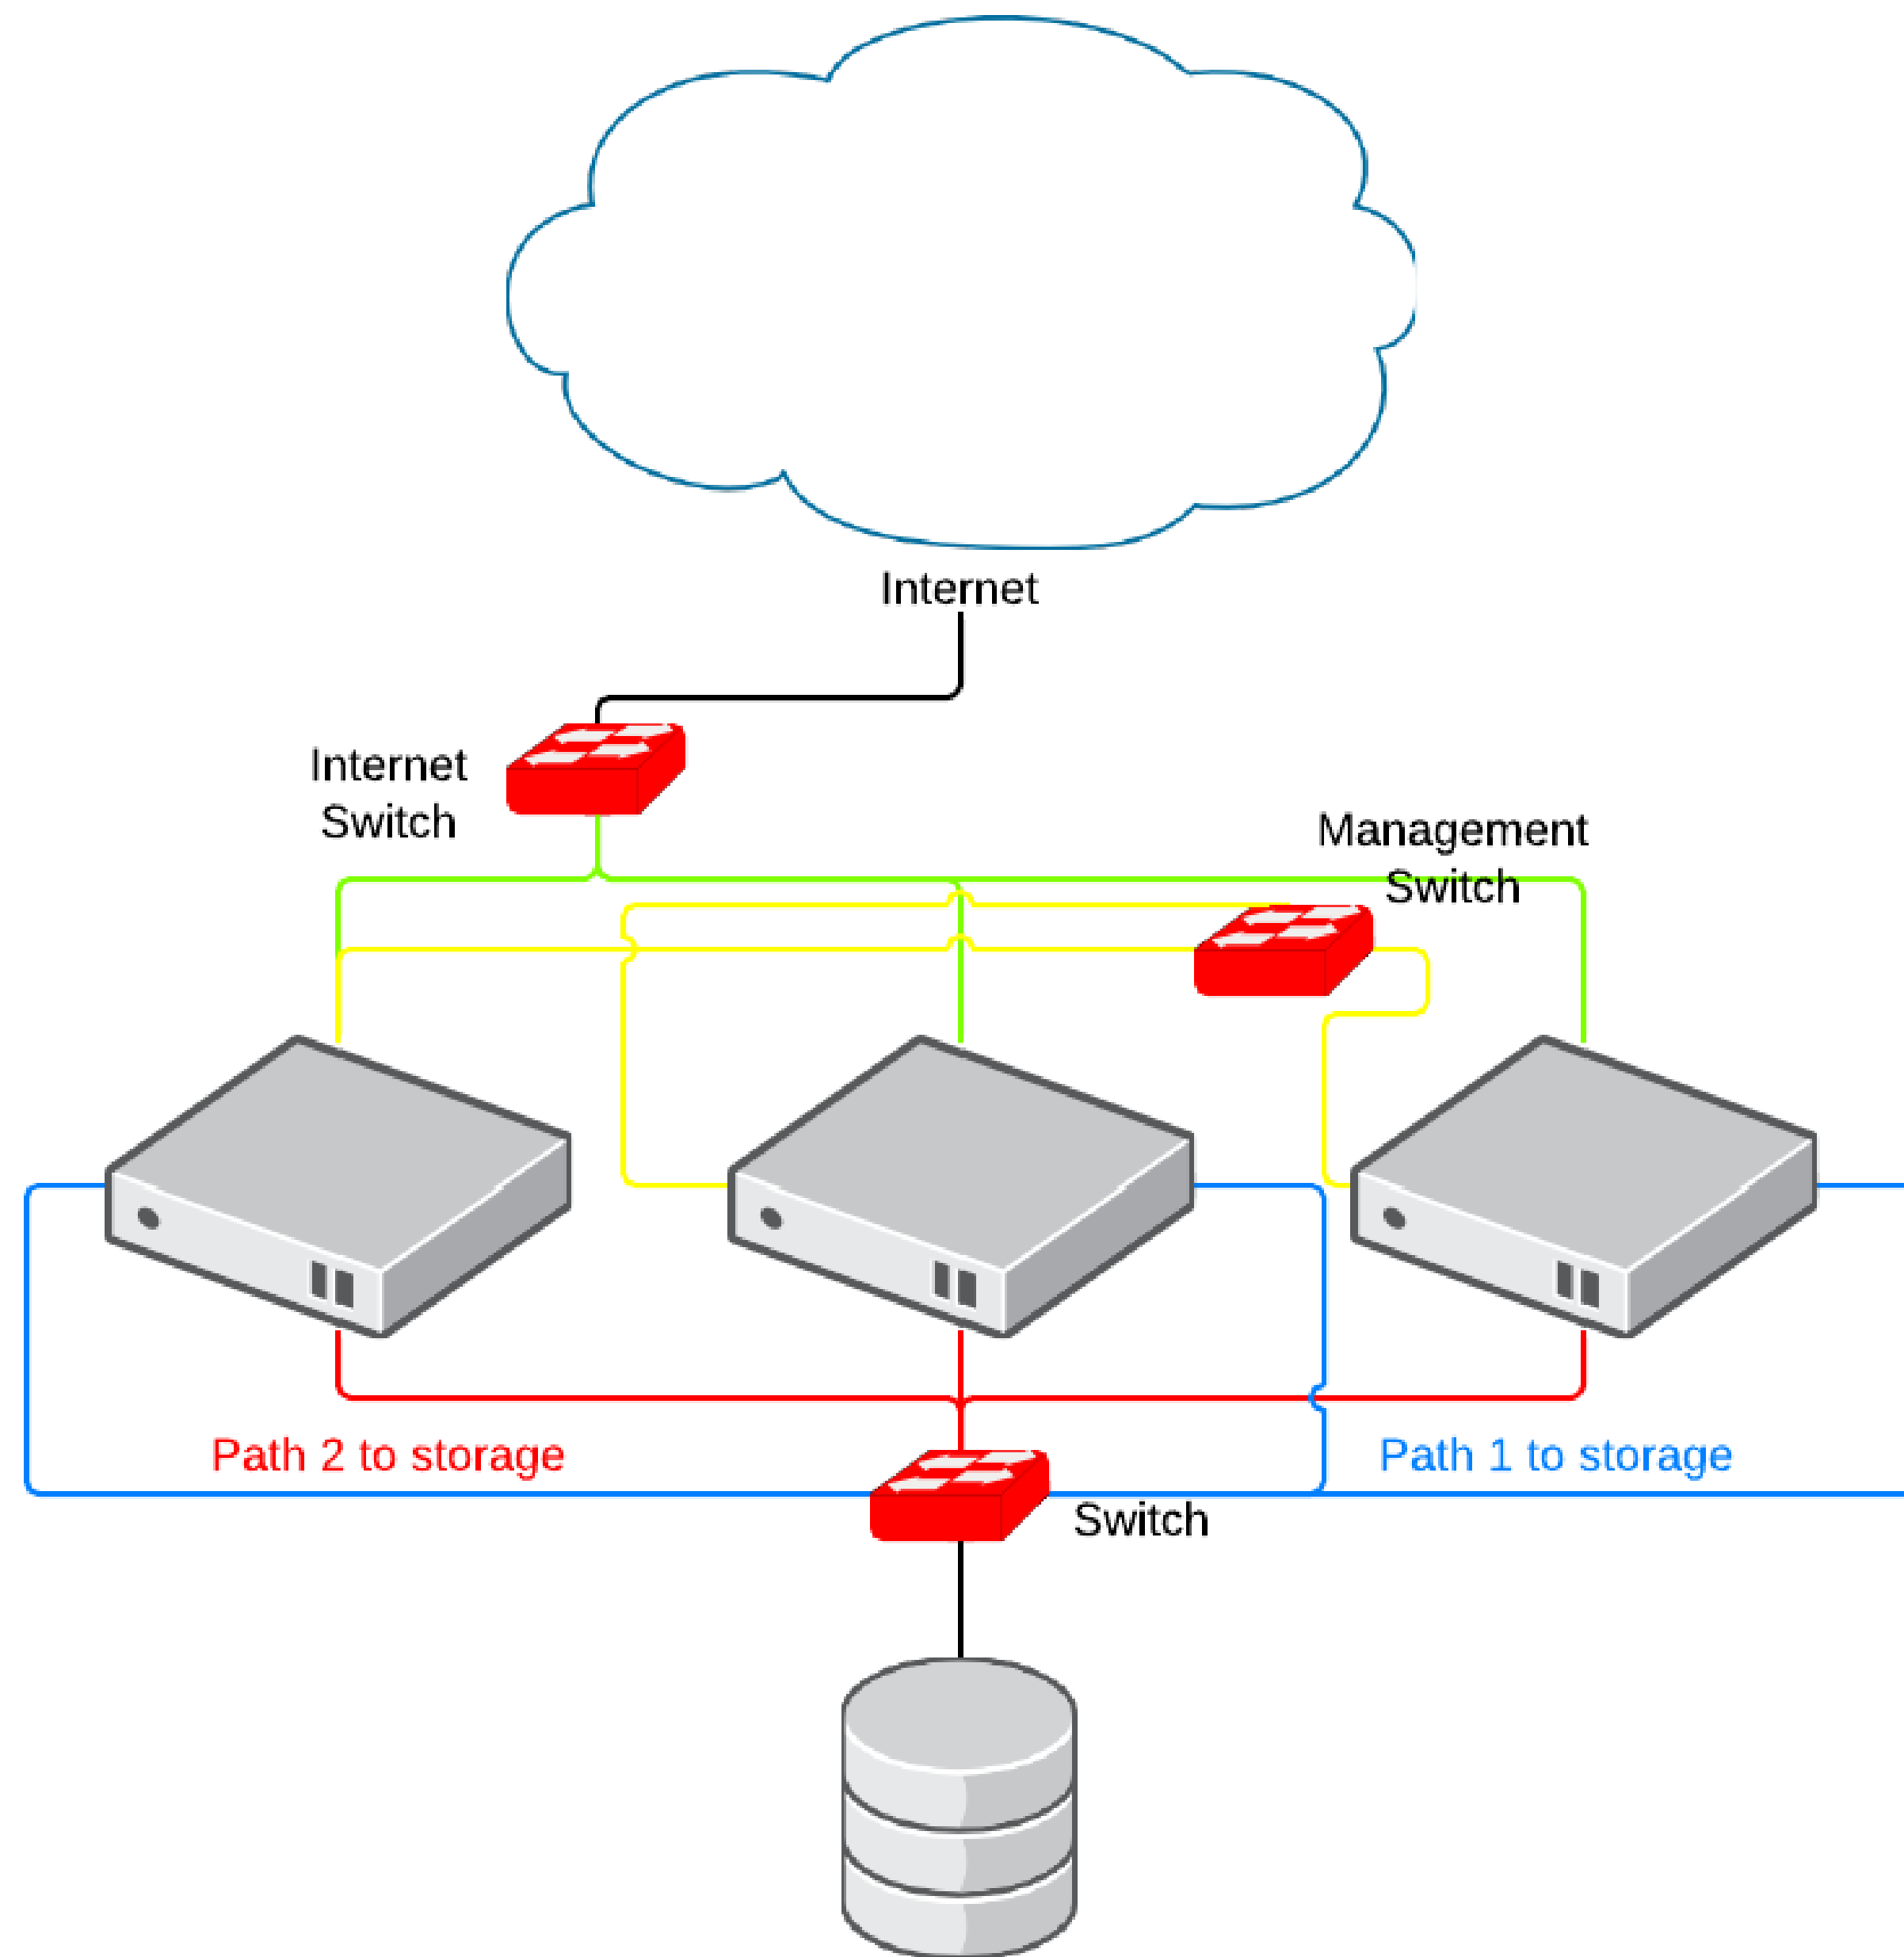
Types of clusters

Types of Clusters

- Generally speaking, there are two common types of clusters
 - High availability
 - Computational



High Availability Clusters

- Multiple nodes serving the same workload. Primary design goal is that if one node goes down, the application continues with minimal interruption.
- Can use shared storage
 - Clustered DB
 - Clustered web apps
 - Clustered file servers



Computational Clusters

- Multiple nodes working on the same or similar datasets. Design goal is to throw massive resources at a complex problem to solve it quickly.
- Typically uses local or network storage
 - Monte Carlo simulations for financial services
 - Oil field reservoir simulation and modeling
 - Chip design and simulation and modeling



NFS Storage


RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

WHAT WE'RE GOING TO DISCUSS TODAY

High Availability Clustering

- We're going to do a **very** basic cluster. We'd never use this design for production, but it is great for self-study and attaining a basic understanding of clustering.
 - Single ethernet network
 - iSCSI storage
 - Three nodes



RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

HOST INSTALLATION AND REQUIREMENTS

Kickstart

- Super basic kickstart – just @core and @base

Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://molly.tc.... /pub/rhel7ks.cfg

molly.tc.redhat.com/pub/rhel7ks.cfg

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home

```
#version=RHEL7
# System authorization information
auth --enablesshadow --passalgo=sha512

text
reboot

# Use network installation
url --url="http://molly.tc.redhat.com/pub/rhel7/beta/"
# Run the Setup Agent on first boot
firstboot --enable
#ignoredisk --only-use=vda
# Keyboard layouts
keyboard --vckeymap=us --xlayouts='us'
# System language
lang en_US.UTF-8
repo --name=updates --baseurl=http://molly.tc.redhat.com/pub/rhel7/Server/

# Network information
network --hostname=localhost.localdomain
# Root password
rootpw --iscrypted $6$U2JEK41IJh3H9hC2$xD28cZkW7U4KdFle8W4Cw/i7lBEJPrkFtEgA1SVUgdbGG3rJY8NP90Mi8bw8z90.XcHzEwh2mW0Kz03FZLNKE1
# System timezone
timezone America/Chicago --isUtc --ntpervers=time.nist.gov
# System bootloader configuration
bootloader --location=mbr

# Partition clearing information
clearpart --all --initlabel
zerombr

# Disk partitioning information
part /boot --fstype="xfs" --size=500
part swap --fstype="swap" --size=2100
part / --fstype="xfs" --size=10000

%packages
@core
@base

%end
```


Disable the Firewall

- We would not do this in production – this is only for lab purposes
 - `systemctl disable firewalld`
 - `systemctl stop firewalld`
 - `iptables -L`

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# systemctl disable firewalld  
rm '/etc/systemd/system/basic.target.wants/firewalld.service'  
rm '/etc/systemd/system/dbus-org.fedoraproject.FirewallD1.service'  
[root@neuromancer ~]#
```

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# systemctl stop firewalld  
[root@neuromancer ~]# iptables -L  
Chain INPUT (policy ACCEPT)  
target      prot opt source                destination  
  
Chain FORWARD (policy ACCEPT)  
target      prot opt source                destination  
  
Chain OUTPUT (policy ACCEPT)  
target      prot opt source                destination  
[root@neuromancer ~]#
```

Set SELinux to enforcing

- THIS IS A TERRIBLE, TERRIBLE PRACTICE.
 - The **only** reason I'm doing this here is because this is still beta software and the SELinux policies have not been completed.
- `perl -pi.orig -e 's/SELINUX=enforcing/SELINUX=permissive/g' /etc/selinux/config`
- `setenforce 0`

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# perl -pi.orig -e 's/SELINUX=enforcing/SELINUX=permissive/'  
/etc/selinux/config  
[root@neuromancer ~]# setenforce 0  
[root@neuromancer ~]# getenforce  
Permissive  
[root@neuromancer ~]#
```


Set up time

- Run “timedatectl status” to make sure that the systemd time and date service is active and using NTP

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# timedatectl status  
    Local time: Sat 2014-03-22 15:48:47 CDT  
    Universal time: Sat 2014-03-22 20:48:47 UTC  
    RTC time: Sat 2014-03-22 20:48:48  
    Timezone: America/Chicago (CDT, -0500)  
    NTP enabled: yes  
    NTP synchronized: yes  
    RTC in local TZ: no  
    DST active: yes  
    Last DST change: DST began at  
                     Sun 2014-03-09 01:59:59 CST  
                     Sun 2014-03-09 03:00:00 CDT  
    Next DST change: DST ends (the clock jumps one hour backwards) at  
                     Sun 2014-11-02 01:59:59 CDT  
                     Sun 2014-11-02 01:00:00 CST  
[root@neuromancer ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# timedatectl status  
    Local time: Sat 2014-03-22 15:48:51 CDT  
    Universal time: Sat 2014-03-22 20:48:51 UTC  
    RTC time: Sat 2014-03-22 20:48:51  
    Timezone: America/Chicago (CDT, -0500)  
    NTP enabled: yes  
    NTP synchronized: yes  
    RTC in local TZ: no  
    DST active: yes  
    Last DST change: DST began at  
                     Sun 2014-03-09 01:59:59 CST  
                     Sun 2014-03-09 03:00:00 CDT  
    Next DST change: DST ends (the clock jumps one hour backwards) at  
                     Sun 2014-11-02 01:59:59 CDT  
                     Sun 2014-11-02 01:00:00 CST  
[root@hideo ~]#
```

```
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# timedatectl status  
    Local time: Sat 2014-03-22 15:48:52 CDT  
    Universal time: Sat 2014-03-22 20:48:52 UTC  
    RTC time: Sat 2014-03-22 20:48:51  
    Timezone: America/Chicago (CDT, -0500)  
    NTP enabled: yes  
    NTP synchronized: yes  
    RTC in local TZ: no  
    DST active: yes  
    Last DST change: DST began at  
                     Sun 2014-03-09 01:59:59 CST  
                     Sun 2014-03-09 03:00:00 CDT  
    Next DST change: DST ends (the clock jumps one hour backwards) at  
                     Sun 2014-11-02 01:59:59 CDT  
                     Sun 2014-11-02 01:00:00 CST  
[root@lady3jane ~]#
```

```
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# timedatectl status  
    Local time: Sat 2014-03-22 15:48:51 CDT  
    Universal time: Sat 2014-03-22 20:48:51 UTC  
    RTC time: Sat 2014-03-22 20:48:51  
    Timezone: America/Chicago (CDT, -0500)  
    NTP enabled: yes  
    NTP synchronized: yes  
    RTC in local TZ: no  
    DST active: yes  
    Last DST change: DST began at  
                     Sun 2014-03-09 01:59:59 CST  
                     Sun 2014-03-09 03:00:00 CDT  
    Next DST change: DST ends (the clock jumps one hour backwards) at  
                     Sun 2014-11-02 01:59:59 CDT  
                     Sun 2014-11-02 01:00:00 CST  
[root@maelcum ~]#
```

Set up passwordless ssh between hosts

- On each node, run ssh-keygen


```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# ssh-keygen  
Generating public/private rsa key pair.  
Enter file in which to save the key (/root/.ssh/id_rsa):  
Enter passphrase (empty for no passphrase):  
Enter same passphrase again:  
Your identification has been saved in /root/.ssh/id_rsa.  
Your public key has been saved in /root/.ssh/id_rsa.pub.  
The key fingerprint is:  
e9:63:2a:8a:e1:95:c8:f8:93:4f:ce:e0:d2:49:fc:cc root@neuromancer.tc.redhat.com  
The key's randomart image is:  
+--[ RSA 2048 ]-----+  
|  
|  
|. . S  
|o + . .  
|o=.0. +  
|o=*E o .  
|.O+++..  
+-----+  
[root@neuromancer ~]#
```

```

root@lady3jane:~
File Edit View Search Terminal Help
[root@lady3jane ~]# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
b9:b5:d7:6b:cb:f9:1d:22:04:21:86:de:1e:95:9e:c3 root@lady3jane.tc.redhat.com
The key's randomart image is:
+---[ RSA 2048]-----+
|      .o  ..          |
|      .. .o.         |
|      . . +..        |
|      . o E.         |
|      . S o.         |
|      . o.. .        |
|      . ....         |
|      ...ooo         |
|      . =oo         |
+-----+
[root@lady3jane ~]# 
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# ssh-keygen  
Generating public/private rsa key pair.  
  
Enter file in which to save the key (/root/.ssh/id_rsa): Enter passphrase (empty  
for no passphrase):  
Enter same passphrase again:  
Your identification has been saved in /root/.ssh/id_rsa.  
Your public key has been saved in /root/.ssh/id_rsa.pub.  
The key fingerprint is:  
3b:7e:bb:68:07:ec:e8:fa:64:66:6a:08:76:38:66:c4 root@hideo.tc.redhat.com  
The key's randomart image is:  
+--[ RSA 2048 ]-----+  
|  
|  
|.E  
|..S  
|*..o.  
|+.o.=oo.  
|. *..ooo  
|.o+o.o0000  
+-----+  
[root@hideo ~]#
```

```

root@maelcum:~
File Edit View Search Terminal Help

[root@maelcum ~]# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
b6:83:bc:20:0e:04:f2:c3:95:92:41:be:fb:f3:e3:72 root@maelcum.tc.redhat.com
The key's randomart image is:
+--[ RSA 2048 ]-----+
| .o                    |
| . o .                 |
| o + o                 |
| oo +                  |
| . =          S        |
| . o . o .             |
| . o . o o             |
| o oo.E. .             |
| . .*+.                |
+-----+
[root@maelcum ~]# 

```


Distribute keys

- Use ssh-copy-id to distribute the keys to every host in the cluster (including itself)
 - Something like “for i in host1 host2 host3 host4; do ssh-copy-id \$i; done”

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# for i in neuromancer lady3jane hideo maelcum; do ssh-copy-  
id $i; done  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@neuromancer's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'neuromancer'"  
and check to make sure that only the key(s) you wanted were added.  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@lady3jane's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'lady3jane'"  
and check to make sure that only the key(s) you wanted were added.
```

root@lady3jane:~

File Edit View Search Terminal Help

```
[root@lady3jane ~]# cat .ssh/authorized_keys
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQBAQC1cRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF
Iqwjfi6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRo1Sk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNCr8
CzifNjPipmJYLxJ/nfbgRqhjswECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS
xo07iZHsycSiT6QXJ0QPoIE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8W+R
DJ4quaY6rGtxRTHSr3j fDk7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc
.redhat.com
[root@lady3jane ~]#
```

```
root@neuromancer:~  
File Edit View Search Terminal Help  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@hideo's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'hideo'"  
and check to make sure that only the key(s) you wanted were added.  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@maelcum's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'maelcum'"  
and check to make sure that only the key(s) you wanted were added.  
  
[root@neuromancer ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
  
[root@lady3jane ~]# cat .ssh/authorized_keys  
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAClRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF  
Iqwji6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRoLSk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNCr8  
CzifNjPipmJYLxJ/nfbgRqhjswECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS  
xo07iZHsycSiT6QXJ0QP0IE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8W+R  
DJ4quaY6rGtxRTHSr3j fDk7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc  
.redhat.com  
[root@lady3jane ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
  
[root@hideo ~]# cat .ssh/authorized_keys  
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAClRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF  
Iqwji6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRoLSk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNCr8  
CzifNjPipmJYLxJ/nfbgRqhjswECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS  
xo07iZHsycSiT6QXJ0QP0IE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8W+R  
DJ4quaY6rGtxRTHSr3j fDk7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc  
.redhat.com  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
  
[root@maelcum ~]# cat .ssh/authorized_keys  
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAClRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF  
Iqwji6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRoLSk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNCr8  
CzifNjPipmJYLxJ/nfbgRqhjswECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS  
xo07iZHsycSiT6QXJ0QP0IE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8W+R  
DJ4quaY6rGtxRTHSr3j fDk7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc  
.redhat.com  
[root@maelcum ~]#
```

```
root@neuromancer:~  
File Edit View Search Terminal Help  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@hideo's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'hideo'"  
and check to make sure that only the key(s) you wanted were added.  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@maelcum's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'maelcum'"  
and check to make sure that only the key(s) you wanted were added.  
  
[root@neuromancer ~]# for i in neuromancer lady3jane hideo maelcum; do ssh-copy-  
id $i; done  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# for i in neuromancer lady3jane hideo maelcum; do ssh-copy-id  
$i; done  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# for i in neuromancer lady3jane hideo maelcum; do ssh-copy-id $i;  
done  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# for i in neuromancer lady3jane hideo maelcum; do ssh-copy-id $  
i; done
```



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@hideo's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'hideo'"  
and check to make sure that only the key(s) you wanted were added.  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@maelcum's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'maelcum'"  
and check to make sure that only the key(s) you wanted were added.  
  
[root@lady3jane ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@hideo's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'hideo'"  
and check to make sure that only the key(s) you wanted were added.  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@maelcum's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'maelcum'"  
and check to make sure that only the key(s) you wanted were added.  
  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@hideo's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'hideo'"  
and check to make sure that only the key(s) you wanted were added.  
  
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter  
out any that are already installed  
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompt  
ed now it is to install the new keys  
root@maelcum's password:  
  
Number of key(s) added: 1  
  
Now try logging into the machine, with: "ssh 'maelcum'"  
and check to make sure that only the key(s) you wanted were added.  
  
[root@maelcum ~]#
```



```
root@neuromancer:~
File Edit View Search Terminal Help
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAC1cRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF
Iqwji6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRoLSk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNcr8
CzifNjPipmJYLxJ/nfbgRqhjsWECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS
xo07iZHsycSiT6QXJ0QPoIE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8w+R
DJ4quaY6rGtxRTHS3jfdK7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc
.redhat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADQJLBzCIbz7jgUtiP0Rx6M7gaacBlQMqMF0PZc4HnN
eI8CURR0mx+W6Q+A2oypg4ExdfTs6bycUNww8dCWoHzPU0ywmw3upJP/Gb9DX7bYmiVli5MbBj femlMk
iUcMFVbMGPC27EuQq/DnZTjSrll58b7X5TWL+LYPAgo/zJqCqrsjiE1u1exGiT6sFYbaeXS7nA9ypY
Hu/WVuPUZ/ym6vf08Zh5keZ6bshUwR4SL0RZbG5S7nPXqxmSQQL/RKNKwA6TDPGJIPsg0kJSCs0rrbNY
sCdQzVt1gGZeUjeU4Jocjkr6pBsJYHwvAb7C4m0ugX8G26X4aMUuBSQ5xAV3 root@lady3jane.tc.r
edhat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADe0A6jqFmp93y1p+Tzy9IMA1j+Esvbp9o/NKiRMjW+
QN+d0kSQ2DELPMLFqpHsl1g5qVcbQsESYGSis6r3DdQMzJK34ZdYjUQ8IYGhyJA0z0f4/D8qvNu/j6Vf
Sst8HNeiYjwkNq4jWfePmVNYDHYkfZpYRF3C3Dqo0WR8SKvBEF1J7cLf/oL9qj95ZQaR0feThYzMYxf4
H9k3dAPg3DbBbVtvuic6Z6zW13rLiLW8DKK0i1gQ86cb4LZEde+6sa7NM8KdJ+aNPkn9kuh2T/muIyug
aL8eR3ZXbTGC5XwL2RhzdKE07xRSuqP0X0PhvS4rSoreNDjtt2Zk6dtiI0/v root@maelcum.tc.red
hat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACV9IDLwMCD9Xeyz1V15CrC8m/AkvvQr9Hc+QvL+mpH
79yPzE+C322rgW6nIBvYXQ6BLlC9evPhToZ4i2u96ycEuJSm9Mn1LN2RbDzbUaigDr0SA672AQya7Bjd
+fj5XpV6C/rczEFQj652m50jhDcxRyBsQnb/E3S7nC+KIiDtYjM0fPmuscPiW/LdfHt07g5eVY10rnKl
3ay5408WLGJFhacA6FhosPovl5wayV0VbikKrvwfSIEtbtBzi6+Ha3np06Y+B5KU0g3XxVNLtEdjJb0K
0pJWY4vfS+Tf5GFcKhG/54svk0j8/UzA/ORBAJCU5vKjUq+6duYbJU5AzSkn root@hideo.tc.redha
.ssh/authorized_keys

root@hideo:~
File Edit View Search Terminal Help
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAC1cRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF
Iqwji6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRoLSk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNcr8
CzifNjPipmJYLxJ/nfbgRqhjsWECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS
xo07iZHsycSiT6QXJ0QPoIE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8w+R
DJ4quaY6rGtxRTHS3jfdK7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc
.redhat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADQJLBzCIbz7jgUtiP0Rx6M7gaacBlQMqMF0PZc4HnN
eI8CURR0mx+W6Q+A2oypg4ExdfTs6bycUNww8dCWoHzPU0ywmw3upJP/Gb9DX7bYmiVli5MbBj femlMk
iUcMFVbMGPC27EuQq/DnZTjSrll58b7X5TWL+LYPAgo/zJqCqrsjiE1u1exGiT6sFYbaeXS7nA9ypY
Hu/WVuPUZ/ym6vf08Zh5keZ6bshUwR4SL0RZbG5S7nPXqxmSQQL/RKNKwA6TDPGJIPsg0kJSCs0rrbNY
sCdQzVt1gGZeUjeU4Jocjkr6pBsJYHwvAb7C4m0ugX8G26X4aMUuBSQ5xAV3 root@lady3jane.tc.r
edhat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADe0A6jqFmp93y1p+Tzy9IMA1j+Esvbp9o/NKiRMjW+
QN+d0kSQ2DELPMLFqpHsl1g5qVcbQsESYGSis6r3DdQMzJK34ZdYjUQ8IYGhyJA0z0f4/D8qvNu/j6Vf
Sst8HNeiYjwkNq4jWfePmVNYDHYkfZpYRF3C3Dqo0WR8SKvBEF1J7cLf/oL9qj95ZQaR0feThYzMYxf4
H9k3dAPg3DbBbVtvuic6Z6zW13rLiLW8DKK0i1gQ86cb4LZEde+6sa7NM8KdJ+aNPkn9kuh2T/muIyug
aL8eR3ZXbTGC5XwL2RhzdKE07xRSuqP0X0PhvS4rSoreNDjtt2Zk6dtiI0/v root@maelcum.tc.red
hat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACV9IDLwMCD9Xeyz1V15CrC8m/AkvvQr9Hc+QvL+mpH
79yPzE+C322rgW6nIBvYXQ6BLlC9evPhToZ4i2u96ycEuJSm9Mn1LN2RbDzbUaigDr0SA672AQya7Bjd
+fj5XpV6C/rczEFQj652m50jhDcxRyBsQnb/E3S7nC+KIiDtYjM0fPmuscPiW/LdfHt07g5eVY10rnKl
3ay5408WLGJFhacA6FhosPovl5wayV0VbikKrvwfSIEtbtBzi6+Ha3np06Y+B5KU0g3XxVNLtEdjJb0K
0pJWY4vfS+Tf5GFcKhG/54svk0j8/UzA/ORBAJCU5vKjUq+6duYbJU5AzSkn root@hideo.tc.redha
.ssh/authorized_keys

root@maelcum:~
File Edit View Search Terminal Help
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAC1cRXy0Abtok9jHhv6jwyNed+wSaHveBnprgRP4IGF
Iqwji6ER50/KPM71P1ZrI7GBEs4I+HeBVqzEPRoLSk/TSRq+/fyP+1fERqX+eoVpp+w8rVR/z/hNcr8
CzifNjPipmJYLxJ/nfbgRqhjsWECHi9g/Z5+iGJB9Lbgh05ZdvYiaGn5E6WbPUD4LRYFxbupYth/FexS
xo07iZHsycSiT6QXJ0QPoIE8rI5ieJz+RB4ZHhpc9wMuZGv6+HEzmYJkFTRVuQqNzBBaetYwAC0I8w+R
DJ4quaY6rGtxRTHS3jfdK7ZQd0QEf/LLf74E2jRLgJcnepw54/wQjifu00h root@neuromancer.tc
.redhat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADQJLBzCIbz7jgUtiP0Rx6M7gaacBlQMqMF0PZc4HnN
eI8CURR0mx+W6Q+A2oypg4ExdfTs6bycUNww8dCWoHzPU0ywmw3upJP/Gb9DX7bYmiVli5MbBj femlMk
iUcMFVbMGPC27EuQq/DnZTjSrll58b7X5TWL+LYPAgo/zJqCqrsjiE1u1exGiT6sFYbaeXS7nA9ypY
Hu/WVuPUZ/ym6vf08Zh5keZ6bshUwR4SL0RZbG5S7nPXqxmSQQL/RKNKwA6TDPGJIPsg0kJSCs0rrbNY
sCdQzVt1gGZeUjeU4Jocjkr6pBsJYHwvAb7C4m0ugX8G26X4aMUuBSQ5xAV3 root@lady3jane.tc.r
edhat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADe0A6jqFmp93y1p+Tzy9IMA1j+Esvbp9o/NKiRMjW+
QN+d0kSQ2DELPMLFqpHsl1g5qVcbQsESYGSis6r3DdQMzJK34ZdYjUQ8IYGhyJA0z0f4/D8qvNu/j6Vf
Sst8HNeiYjwkNq4jWfePmVNYDHYkfZpYRF3C3Dqo0WR8SKvBEF1J7cLf/oL9qj95ZQaR0feThYzMYxf4
H9k3dAPg3DbBbVtvuic6Z6zW13rLiLW8DKK0i1gQ86cb4LZEde+6sa7NM8KdJ+aNPkn9kuh2T/muIyug
aL8eR3ZXbTGC5XwL2RhzdKE07xRSuqP0X0PhvS4rSoreNDjtt2Zk6dtiI0/v root@maelcum.tc.red
hat.com
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACV9IDLwMCD9Xeyz1V15CrC8m/AkvvQr9Hc+QvL+mpH
79yPzE+C322rgW6nIBvYXQ6BLlC9evPhToZ4i2u96ycEuJSm9Mn1LN2RbDzbUaigDr0SA672AQya7Bjd
+fj5XpV6C/rczEFQj652m50jhDcxRyBsQnb/E3S7nC+KIiDtYjM0fPmuscPiW/LdfHt07g5eVY10rnKl
3ay5408WLGJFhacA6FhosPovl5wayV0VbikKrvwfSIEtbtBzi6+Ha3np06Y+B5KU0g3XxVNLtEdjJb0K
0pJWY4vfS+Tf5GFcKhG/54svk0j8/UzA/ORBAJCU5vKjUq+6duYbJU5AzSkn root@hideo.tc.redha
.ssh/authorized_keys
```


Distribute keys

- Make sure you can log into all of the machines without being prompted to accept the key

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# for i in neuromancer lady3jane hideo maelcum; do ssh $i "h  
ostname"; ssh $i.tc.redhat.com "hostname"; done  
neuromancer.tc.redhat.com  
Warning: Permanently added 'neuromancer.tc.redhat.com' (ECDSA) to the list of kn  
own hosts.  
neuromancer.tc.redhat.com  
lady3jane.tc.redhat.com  
Warning: Permanently added 'lady3jane.tc.redhat.com' (ECDSA) to the list of know  
n hosts.  
lady3jane.tc.redhat.com  
hideo.tc.redhat.com  
Warning: Permanently added 'hideo.tc.redhat.com' (ECDSA) to the list of known ho  
sts.  
hideo.tc.redhat.com  
maelcum.tc.redhat.com  
Warning: Permanently added 'maelcum.tc.redhat.com' (ECDSA) to the list of known  
hosts.  
maelcum.tc.redhat.com  
[root@neuromancer ~]#
```

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# for i in neuromancer lady3jane hideo maelcum; do ssh $i "hostname"; ssh $i.tc.redhat.com "hostname"; done  
neuromancer.tc.redhat.com  
Warning: Permanently added 'neuromancer.tc.redhat.com' (ECDSA) to the list of known hosts.  
neuromancer.tc.redhat.com  
lady3jane.tc.redhat.com  
Warning: Permanently added 'lady3jane.tc.redhat.com' (ECDSA) to the list of known hosts.  
lady3jane.tc.redhat.com  
hideo.tc.redhat.com  
Warning: Permanently added 'hideo.tc.redhat.com' (ECDSA) to the list of known hosts.  
hideo.tc.redhat.com  
maelcum.tc.redhat.com  
Warning: Permanently added 'maelcum.tc.redhat.com' (ECDSA) to the list of known hosts.  
maelcum.tc.redhat.com  
[root@neuromancer ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# for i in neuromancer lady3jane hideo maelcum; do ssh $i "hostname"; ssh $i.tc.redhat.com "hostname"; done  
neuromancer.tc.redhat.com  
Warning: Permanently added 'neuromancer.tc.redhat.com' (ECDSA) to the list of known hosts.  
neuromancer.tc.redhat.com  
lady3jane.tc.redhat.com  
Warning: Permanently added 'lady3jane.tc.redhat.com' (ECDSA) to the list of known hosts.  
lady3jane.tc.redhat.com  
hideo.tc.redhat.com  
Warning: Permanently added 'hideo.tc.redhat.com' (ECDSA) to the list of known hosts.  
hideo.tc.redhat.com  
maelcum.tc.redhat.com  
Warning: Permanently added 'maelcum.tc.redhat.com' (ECDSA) to the list of known hosts.  
maelcum.tc.redhat.com  
[root@lady3jane ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# for i in neuromancer lady3jane hideo maelcum; do ssh $i "hostname"; ssh $i.tc.redhat.com "hostname"; done  
neuromancer.tc.redhat.com  
Warning: Permanently added 'neuromancer.tc.redhat.com' (ECDSA) to the list of known hosts.  
neuromancer.tc.redhat.com  
lady3jane.tc.redhat.com  
Warning: Permanently added 'lady3jane.tc.redhat.com' (ECDSA) to the list of known hosts.  
lady3jane.tc.redhat.com  
hideo.tc.redhat.com  
Warning: Permanently added 'hideo.tc.redhat.com' (ECDSA) to the list of known hosts.  
hideo.tc.redhat.com  
maelcum.tc.redhat.com  
Warning: Permanently added 'maelcum.tc.redhat.com' (ECDSA) to the list of known hosts.  
maelcum.tc.redhat.com  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# for i in neuromancer lady3jane hideo maelcum; do ssh $i "hostname"; ssh $i.tc.redhat.com "hostname"; done  
neuromancer.tc.redhat.com  
Warning: Permanently added 'neuromancer.tc.redhat.com' (ECDSA) to the list of known hosts.  
neuromancer.tc.redhat.com  
lady3jane.tc.redhat.com  
Warning: Permanently added 'lady3jane.tc.redhat.com' (ECDSA) to the list of known hosts.  
lady3jane.tc.redhat.com  
hideo.tc.redhat.com  
Warning: Permanently added 'hideo.tc.redhat.com' (ECDSA) to the list of known hosts.  
hideo.tc.redhat.com  
maelcum.tc.redhat.com  
Warning: Permanently added 'maelcum.tc.redhat.com' (ECDSA) to the list of known hosts.  
maelcum.tc.redhat.com  
[root@maelcum ~]#
```


Distribute keys

- Make sure that you can log in via short and long hostname. Don't quit until get this result:

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# for i in neuromancer lady3jane hideo maelcum; do ssh $i "h  
ostname"; ssh $i.tc.redhat.com "hostname"; done  
neuromancer.tc.redhat.com  
neuromancer.tc.redhat.com  
lady3jane.tc.redhat.com  
lady3jane.tc.redhat.com  
hideo.tc.redhat.com  
hideo.tc.redhat.com  
maelcum.tc.redhat.com  
maelcum.tc.redhat.com  
[root@neuromancer ~]#
```

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

SET UP THE ISCSI TARGET (SERVER)

Linux-IO Target (LIO)

- Linux-IO Target (LIO™) has been the Linux SCSI target since kernel version 2.6.38.
- Linux-IO Target is based on a SCSI engine that implements the semantics of a SCSI target as described in the SCSI Architecture Model (SAM), and supports its comprehensive SPC-3/SPC-4 feature set in a fabric-agnostic way. The SCSI target core does not directly communicate with initiators and it does not directly access data on disk.

Source: <http://linux-iscsi.org/wiki/Target>

Cheat Sheet

- I love the “cheat sheet” at http://linux-iscsi.org/wiki/ISCSI#Cheat_sheet

Create the partition

- Use your favorite partitioning tool. Since this is a small disk, I used fdisk. For larger disks, parted makes sense.
- I created a new primary partition, using the rest of the disk. Again, this is just for a lab, in the real world I would probably use LVM.
- Reboot if prompted

root@neuromancer:~

File Edit View Search Terminal Help

[root@neuromancer ~]# fdisk /dev/sda

Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.

Be careful before using the write command.

Command (m for help): p

Disk /dev/sda: 500.1 GB, 500107862016 bytes, 976773168 sectors

Units = sectors of 1 * 512 = 512 bytes

Sector size (logical/physical): 512 bytes / 512 bytes

I/O size (minimum/optimal): 512 bytes / 512 bytes

Disk label type: dos

Disk identifier: 0x000ab112

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	2048	1026047	512000	83	Linux
/dev/sda2		1026048	21506047	10240000	83	Linux
/dev/sda3		21506048	25806847	2150400	82	Linux swap / Solaris

Command (m for help):


```
root@neuromancer:~  
File Edit View Search Terminal Help  
Command (m for help): n  
Partition type:  
  p   primary (3 primary, 0 extended, 1 free)  
  e   extended  
Select (default e): p  
Selected partition 4  
First sector (25806848-976773167, default 25806848):  
Using default value 25806848  
Last sector, +sectors or +size{K,M,G} (25806848-976773167, default 976773167):  
Using default value 976773167  
Partition 4 of type Linux and of size 453.5 GiB is set  
  
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
  
WARNING: Re-reading the partition table failed with error 16: Device or resource  
        busy.  
The kernel still uses the old table. The new table will be used at  
the next reboot or after you run partprobe(8) or kpartx(8)  
Syncing disks.  
[root@neuromancer ~]#
```

Install the iSCSI management software

- There are a number of ways we can create an iSCSI export. The simplest is probably targetcli
 - yum install targetcli
 - Brings in necessary python libs and utils as dependencies

root@neuromancer:~

File Edit View Search Terminal Help

```
[root@neuromancer ~]# yum install targetcli
```

Loaded plugins: langpacks, product-id, subscription-manager

This system is not registered to Red Hat Subscription Management. You can use subscription-manager to register.

Resolving Dependencies

--> Running transaction check

---> Package targetcli.noarch 0:2.1.fb34-1.el7 will be installed

--> Processing Dependency: python-rtslib >= 2.1.fb41 for package: targetcli-2.1.fb34-1.el7.noarch

--> Processing Dependency: python-configshell for package: targetcli-2.1.fb34-1.el7.noarch

--> Running transaction check

---> Package python-configshell.noarch 1:1.1.fb11-3.el7 will be installed

--> Processing Dependency: pyparsing for package: 1:python-configshell-1.1.fb11-3.el7.noarch

--> Processing Dependency: python-urwid for package: 1:python-configshell-1.1.fb11-3.el7.noarch

---> Package python-rtslib.noarch 0:2.1.fb46-1.el7 will be installed

--> Processing Dependency: python-kmod for package: python-rtslib-2.1.fb46-1.el7.noarch

--> Running transaction check

---> Package pyparsing.noarch 0:1.5.6-9.el7 will be installed

---> Package python-kmod.x86_64 0:0.9-4.el7 will be installed

---> Package python-urwid.x86_64 0:1.1.1-3.el7 will be installed

```
root@neuromancer:~
File Edit View Search Terminal Help
Running transaction
  Installing : python-kmod-0.9-4.el7.x86_64                                1/6
  Installing : python-rtslib-2.1.fb46-1.el7.noarch                        2/6
  Installing : pyparsing-1.5.6-9.el7.noarch                              3/6
  Installing : python-urwid-1.1.1-3.el7.x86_64                           4/6
  Installing : 1:python-configshell-1.1.fb11-3.el7.noarch                 5/6
  Installing : targetcli-2.1.fb34-1.el7.noarch                           6/6
  Verifying   : python-rtslib-2.1.fb46-1.el7.noarch                       1/6
  Verifying   : python-urwid-1.1.1-3.el7.x86_64                          2/6
  Verifying   : targetcli-2.1.fb34-1.el7.noarch                           3/6
  Verifying   : 1:python-configshell-1.1.fb11-3.el7.noarch               4/6
  Verifying   : pyparsing-1.5.6-9.el7.noarch                             5/6
  Verifying   : python-kmod-0.9-4.el7.x86_64                             6/6

Installed:
  targetcli.noarch 0:2.1.fb34-1.el7

Dependency Installed:
  pyparsing.noarch 0:1.5.6-9.el7      python-configshell.noarch 1:1.1.fb11-3.el7
  python-kmod.x86_64 0:0.9-4.el7      python-rtslib.noarch 0:2.1.fb46-1.el7
  python-urwid.x86_64 0:1.1.1-3.el7

Complete!
[root@neuromancer ~]#
```

Enable the “target” service

- The target service is what causes LIO to read its config at boot time.

root@neuromancer:~

File Edit View Search Terminal Help

```
[root@neuromancer ~]# systemctl enable target
ln -s '/usr/lib/systemd/system/target.service' '/etc/systemd/system/multi-user.target.wants/target.service'
[root@neuromancer ~]#
```


Using targetcli

- targetcli can be used interactively via a shell or from the command line. We'll use the shell.
- Check the current config with “ls”

root@neuromancer:~

File Edit View Search Terminal Help

```
[root@neuromancer ~]# targetcli
targetcli shell version 2.1.fb34
Copyright 2011-2013 by Datera, Inc and others.
For help on commands, type 'help'.
```

```
/> ls
```

```
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 0]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 0]
  o- loopback ..... [Targets: 0]
  o- srpt ..... [Targets: 0]
```

```
/> 
```

Define the backing store for your LUN

- You have to tell the LIO iSCSI target software about the block device you want to use
 - `cd /backstores/block` (“cd” is optional)
 - `create [lun] /dev/[device]` (I called mine lun0 on sda4)
 - `ls` to check results
- `/backstores/block create lun0 /dev/sda4`

root@neuromancer:~

File Edit View Search Terminal Help

```
[root@neuromancer ~]# targetcli
targetcli shell version 2.1.fb34
Copyright 2011-2013 by Datera, Inc and others.
For help on commands, type 'help'.

/> /backstores/block create lun0 /dev/sda4
Created block storage object lun0 using /dev/sda4.
/> 
```

root@neuromancer:~

File Edit View Search Terminal Help

```
/> /backstores/block create lun0 /dev/sda4
```

Created block storage object lun0 using /dev/sda4.

```
/> ls
```

```
o- / ..... [ ... ]
  o- backstores ..... [ ... ]
    | o- block ..... [Storage Objects: 1]
    | | o- lun0 ..... [/dev/sda4 (453.5GiB) write-thru deactivated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 0]
  o- loopback ..... [Targets: 0]
  o- srpt ..... [Targets: 0]
```

```
/> 
```

Create an iSCSI Qualified Name (IQN) record

- /iscsi
- create (you **can** specify and IQN, but it's easier to let it auto-generate it for you)


```
root@neuromancer:~  
File Edit View Search Terminal Help  
/> /iscsi create  
Created target iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.a730e011546e.  
Created TPG 1.  
/> █
```

```
/> /iscsi create
```

```
Created target iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646.
```

```
Created TPG 1.
```

```
/> ls
```

```
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 1]
    | | o- lun0 ..... [/dev/sda4 (453.5GiB) write-thru deactivated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 . [TPGs: 1]
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 0]
    |     o- luns ..... [LUNs: 0]
    |     o- portals ..... [Portals: 0]
  o- loopback ..... [Targets: 0]
  o- srpt ..... [Targets: 0]
```

```
/> □
```

Create an iSCSI portal address

- Unless you specify an address, it will listen on all addresses 0.0.0.0
 - `cd /iscsi/[your_iqn]/tpg1/portals`
 - `create`
 - `ls`

root@neuromancer:~

File Edit View Search Terminal Help

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646/tpg1/portals create
```

```
Using default IP port 3260
```

```
Binding to INADDR_ANY (0.0.0.0)
```

```
Created network portal 0.0.0.0:3260.
```

```
/> 
```

root@neuromancer:~

File Edit View Search Terminal Help

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646/tpg1/portals create
```

```
Using default IP port 3260
```

```
Binding to INADDR_ANY (0.0.0.0)
```

```
Created network portal 0.0.0.0:3260.
```

```
/> ls
```

```
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 1]
    | | o- lun0 ..... [/dev/sda4 (453.5GiB) write-thru deactivated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 . [TPGs: 1]
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 0]
    |     o- luns ..... [LUNs: 0]
    |     o- portals ..... [Portals: 1]
    |       o- 0.0.0.0:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
  o- srpt ..... [Targets: 0]
/> 
```


Export the LUN

- Now you need to make that LUN available via that portal and IQN
 - `cd /iscsi/[your_iqn]/tpg1/luns`
 - `create /backstores/block/[lun]`
 - `ls`

root@neuromancer:~

File Edit View Search Terminal Help

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646/tpg1/lun  
s create /backstores/block/lun0
```

Created LUN 0.

```
/> 
```

```
root@neuromancer:~
File Edit View Search Terminal Help
/> /iscsi/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646/tpg1/lun
s create /backstores/block/lun0
Created LUN 0.
/> ls
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 1]
    | | o- lun0 ..... [/dev/sda4 (453.5GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 . [TPGs: 1]
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 0]
    |     o- luns ..... [LUNs: 1]
    |       | o- lun0 ..... [block/lun0 (/dev/sda4)]
    |       o- portals ..... [Portals: 1]
    |         o- 0.0.0.0:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
  o- srpt ..... [Targets: 0]
/> 
```

Enable Demo Mode

- This is dangerous! Don't do this in production, you'll want to set ACLs. We're doing this in a compressed time frame for a demo setup.
- Go to your target portal group (TPG) and set the following attributes:
 - `cd /iscsi/[your_IQN]/tpg1/`
 - `set attribute authentication=0`
 - `set attribute demo_mode_write_protect=0`
 - `set attribute generate_node_acls=1`
 - `set attribute cache_dynamic_acls=1`

```
root@neuromancer:~  
File Edit View Search Terminal Help  
/> /iscsi/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.a730e011546e/tpg1/  
/iscsi/iqn.20...e011546e/tpg1> set attribute authentication=0  
Parameter authentication is now '0'.  
/iscsi/iqn.20...e011546e/tpg1> set attribute demo_mode_write_protect=0  
Parameter demo_mode_write_protect is now '0'.  
/iscsi/iqn.20...e011546e/tpg1> set attribute generate_node_acls=1  
Parameter generate_node_acls is now '1'.  
/iscsi/iqn.20...e011546e/tpg1> set attribute cache_dynamic_acls=1  
Parameter cache_dynamic_acls is now '1'.  
/iscsi/iqn.20...e011546e/tpg1>
```


Check the top level view

- `cd /`
- `ls`

```
root@neuromancer:~
File Edit View Search Terminal Help
/iscsi/iqn.20...e011546e/tpg1> /
/> ls
o- / ..... [....]
  o- backstores ..... [....]
    | o- block ..... [Storage Objects: 1]
    |   | o- lun0 ..... [/dev/sda4 (453.5GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.a730e011546e . [TPGs: 1]
    |   o- tpg1 ..... [gen-acls, no-auth]
    |     o- acls ..... [ACLs: 0]
    |     o- luns ..... [LUNs: 1]
    |       | o- lun0 ..... [block/lun0 (/dev/sda4)]
    |     o- portals ..... [Portals: 1]
    |       o- 0.0.0.0:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
  o- srpt ..... [Targets: 0]
/> 
```

Save the configuration

- saveconfig
- Exiting saves as well

```
root@neuromancer:~  
File Edit View Search Terminal Help  
/> saveconfig  
Last 10 configs saved in /etc/target/backup.  
Configuration saved to /etc/target/saveconfig.json  
/> exit  
Global pref auto_save_on_exit=true  
Last 10 configs saved in /etc/target/backup.  
Configuration saved to /etc/target/saveconfig.json  
[root@neuromancer ~]#
```

Check out `/etc/target/saveconfig.json`

- cat the file

root@neuromancer:~

File Edit View Search Terminal Help

```
[root@neuromancer ~]# cat /etc/target/saveconfig.json
```

```
{
  "fabric_modules": [],
  "storage_objects": [
    {
      "attributes": {
        "block_size": 512,
        "emulate_dpo": 0,
        "emulate_fua_read": 0,
        "emulate_fua_write": 1,
        "emulate_model_alias": 1,
        "emulate_rest_reord": 0,
        "emulate_tas": 1,
        "emulate_tpu": 0,
        "emulate_tpws": 0,
        "emulate_ua_intlck_ctrl": 0,
        "emulate_write_cache": 0,
        "enforce_pr_isids": 1,
        "fabric_max_sectors": 8192,
        "is_nonrot": 0,
        "max_unmap_block_desc_count": 0,
        "max_unmap_lba_count": 0,
        "max_write_same_len": 65535,
        "optimal_sectors": 8192,
```


root@neuromancer:~

File Edit View Search Terminal Help

```
    "queue_depth": 20480,
    "unmap_granularity": 0,
    "unmap_granularity_alignment": 0
  },
  "dev": "/dev/sda4",
  "name": "lun0",
  "plugin": "block",
  "readonly": false,
  "write_back": false,
  "wwn": "4c0669b8-b674-4e52-bcc0-cae78806818e"
},
"targets": [
  {
    "fabric": "iscsi",
    "tpgs": [
      {
        "attributes": {
          "authentication": 0,
          "cache_dynamic_acls": 1,
          "default_cmdsn_depth": 16,
          "demo_mode_write_protect": 0,
          "generate_node_acls": 1,
          "login_timeout": 15,
```

root@neuromancer:~

File Edit View Search Terminal Help

```
    "netif_timeout": 2,
    "prod_mode_write_protect": 0
  },
  "enable": true,
  "luns": [
    {
      "index": 0,
      "storage_object": "/backstores/block/lun0"
    }
  ],
  "node_acls": [],
  "parameters": {
    "AuthMethod": "CHAP,None",
    "DataDigest": "CRC32C,None",
    "DataPDUInOrder": "Yes",
    "DataSequenceInOrder": "Yes",
    "DefaultTime2Retain": "20",
    "DefaultTime2Wait": "2",
    "ErrorRecoveryLevel": "0",
    "FirstBurstLength": "65536",
    "HeaderDigest": "CRC32C,None",
    "IFMarkInt": "2048~65535",
    "IFMarker": "No",
    "ImmediateData": "Yes",
```

```
root@neuromancer:~  
File Edit View Search Terminal Help  
    "InitialR2T": "Yes",  
    "MaxBurstLength": "262144",  
    "MaxConnections": "1",  
    "MaxOutstandingR2T": "1",  
    "MaxRecvDataSegmentLength": "8192",  
    "MaxXmitDataSegmentLength": "262144",  
    "OFMarkInt": "2048~65535",  
    "OFMarker": "No",  
    "TargetAlias": "LIO Target"  
  },  
  "portals": [  
    {  
      "ip_address": "0.0.0.0",  
      "iser": false,  
      "port": 3260  
    }  
  ],  
  "tag": 1  
},  
"wwn": "iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646"  
]  
}
```



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

SET UP THE ISCSI INITIATORS (CLIENTS)

Install the software on each cluster node

- `yum install iscsi-initiator-utils`

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# yum -y install iscsi-initiator-utils  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
  
===== Size  
Installing:  
iscsi-initiator-utils          x86_64    6.2.0.873-21.el7    rhel    411 k  
Installing for dependencies:  
iscsi-initiator-utils-iscsiuio x86_64    6.2.0.873-21.el7    rhel     79 k  
Transaction Summary  
=====  
Install 1 Package (+1 Dependent package)  
  
Total download size: 490 k  
Installed size: 2.5 M  
Downloading packages:  
(1/2): iscsi-initiator-utils-6.2.0.873-21.el7.x86_64.rpm | 411 kB 00:00  
(2/2): iscsi-initiator-utils-iscsiuio-6.2.0.873-21.el7.x86_64 | 79 kB 00:00  
-----  
Total                               2.0 MB/s | 490 kB 00:00  
Running transaction check  
Running transaction test  
Transaction test succeeded  
Running transaction  
Installing : iscsi-initiator-utils-iscsiuio-6.2.0.873-21.el7.x86_64 1/2  
Installing : iscsi-initiator-utils-6.2.0.873-21.el7.x86_64 2/2  
[##  
  
root@maelcum:~  
File Edit View Search Terminal Help  
=====   
Installing:  
iscsi-initiator-utils          x86_64    6.2.0.873-21.el7    rhel    411 k  
Installing for dependencies:  
iscsi-initiator-utils-iscsiuio x86_64    6.2.0.873-21.el7    rhel     79 k  
Transaction Summary  
=====   
Install 1 Package (+1 Dependent package)  
  
Total download size: 490 k  
Installed size: 2.5 M  
Downloading packages:  
(1/2): iscsi-initiator-utils-iscsiuio-6.2.0.873-21.el7.x86_64 | 79 kB 00:00  
(2/2): iscsi-initiator-utils-6.2.0.873-21.el7.x86_64.rpm | 411 kB 00:00  
-----  
Total                               1.7 MB/s | 490 kB 00:00  
Running transaction check  
Running transaction test  
Transaction test succeeded  
Running transaction  
Installing : iscsi-initiator-utils-iscsiuio-6.2.0.873-21.el7.x86_64 1/2  
Installing : iscsi-initiator-utils-6.2.0.873-21.el7.x86_64 2/2  
[
```


Check that the iSCSI initiator service is enabled

- `systemctl status iscsi`
 - It should be enabled but dead since it has not been started yet

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# systemctl status iscsi  
iscsi.service - Login and scanning of iSCSI devices  
    Loaded: loaded (/usr/lib/systemd/system/iscsi.service; enabled)  
    Active: inactive (dead)  
        Docs: man:iscsid(8)  
              man:iscsiadm(8)  
  
[root@hideo ~]#
```

Discover the target

- Note that there are no iscsi processes running
- /var/lib/scsi is empty
- `iscsiadm --mode discoverydb --type sendtargets --portal [portal] --discover`
- Afterwards, the iscsid process is running
- /var/lib/iscsi is populated

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# ps ax | grep iscsi  
2192 pts/0 S+ 0:00 grep --color=auto iscsi  
[root@hideo ~]# find /var/lib/iscsi/  
/var/lib/iscsi/  
/var/lib/iscsi/ifaces  
/var/lib/iscsi/isns  
/var/lib/iscsi/nodes  
/var/lib/iscsi/send_targets  
/var/lib/iscsi/slp  
/var/lib/iscsi/static  
[root@hideo ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# iscsiadm --mode discoverydb --type sendtargets --portal neuroman  
cer.tc.redhat.com --discover  
172.31.100.14:3260,1 iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd6  
46  
[root@hideo ~]# ps ax | grep iscsi  
2203 ?          Ss      0:00 /usr/sbin/iscsid  
2204 ?          S<Ls    0:00 /usr/sbin/iscsid  
2207 ?          S<      0:00 [iscsi_eh]  
2218 pts/0      S+      0:00 grep --color=auto iscsi  
[root@hideo ~]# find /var/lib/iscsi/  
/var/lib/iscsi/  
/var/lib/iscsi/ifaces  
/var/lib/iscsi/isns  
/var/lib/iscsi/nodes  
/var/lib/iscsi/nodes/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd6  
46  
/var/lib/iscsi/nodes/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd6  
46/172.31.100.14,3260,1  
/var/lib/iscsi/nodes/iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd6  
46/172.31.100.14,3260,1/default  
/var/lib/iscsi/send_targets  
/var/lib/iscsi/send_targets/neuromancer.tc.redhat.com,3260  
/var/lib/iscsi/send_targets/neuromancer.tc.redhat.com,3260/st_config  
/var/lib/iscsi/send_targets/neuromancer.tc.redhat.com,3260/iqn.2003-01.org.linux  
-iscsi.neuromancer.x8664:sn.f63ec35cd646,172.31.100.14,3260,1,default  
/var/lib/iscsi/slp  
/var/lib/iscsi/static  
[root@hideo ~]#
```

Log into the target

- Note that the kernel only “sees” block devices for locally installed hard drives
- `iscsiadm --mode node --targetname [iqn] --portal [portal] --login`
- You should see a new block device

root@hideo:~

File Edit View Search Terminal Help

[root@hideo ~]# cat /proc/partitions

major	minor	#blocks	name
-------	-------	---------	------

8	0	976762584	sda
---	---	-----------	-----

8	1	512000	sda1
---	---	--------	------

8	2	10240000	sda2
---	---	----------	------

8	3	2150400	sda3
---	---	---------	------

[root@hideo ~]# iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --login

Logging in to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] (multiple)

Login to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] successful.

[root@hideo ~]# cat /proc/partitions

major	minor	#blocks	name
-------	-------	---------	------

8	0	976762584	sda
---	---	-----------	-----

8	1	512000	sda1
---	---	--------	------

8	2	10240000	sda2
---	---	----------	------

8	3	2150400	sda3
---	---	---------	------

8	16	475483160	sdb
---	----	-----------	-----

[root@hideo ~]#

Log into the target

- Lather, rinse, repeat for each initiator

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# cat /proc/partitions  
major minor #blocks name  
  
8 0 976762584 sda  
8 1 512000 sda1  
8 2 10240000 sda2  
8 3 2150400 sda3  
[root@hideo ~]# iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --login  
Logging in to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] (multiple)  
Login to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] successful.  
[root@hideo ~]# cat /proc/partitions  
major minor #blocks name  
  
8 0 976762584 sda  
8 1 512000 sda1  
8 2 10240000 sda2  
8 3 2150400 sda3  
8 16 475483160 sdb  
[root@hideo ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# man iscsiadm  
[root@lady3jane ~]# iscsiadm --mode discoverydb --type sendtargets --portal neuromancer.tc.redhat.com --discover  
172.31.100.14:3260,1 iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646  
[root@lady3jane ~]# iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --login  
Logging in to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] (multiple)  
Login to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] successful.  
[root@lady3jane ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# iscsiadm --mode discoverydb --type sendtargets --portal neuromancer.tc.redhat.com --discover  
172.31.100.14:3260,1 iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646  
[root@maelcum ~]# iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --login  
Logging in to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] (multiple)  
Login to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] successful.  
[root@maelcum ~]#
```

We'll come back to this shared storage later

- For now, let's move on to setting up the clustering software

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

COROSYNC AND PACEMAKER

About corosync and pacemaker

- Low level infrastructure corosync provides reliable messaging, membership and quorum information about the cluster (red)



About corosync and pacemaker

- Resource management Pacemaker provides the brain (illustrated in blue) that processes and reacts to events regarding the cluster. These events include:
 - nodes joining or leaving the cluster
 - resource events caused by failures, maintenance, scheduled activities
 - other administrative actions



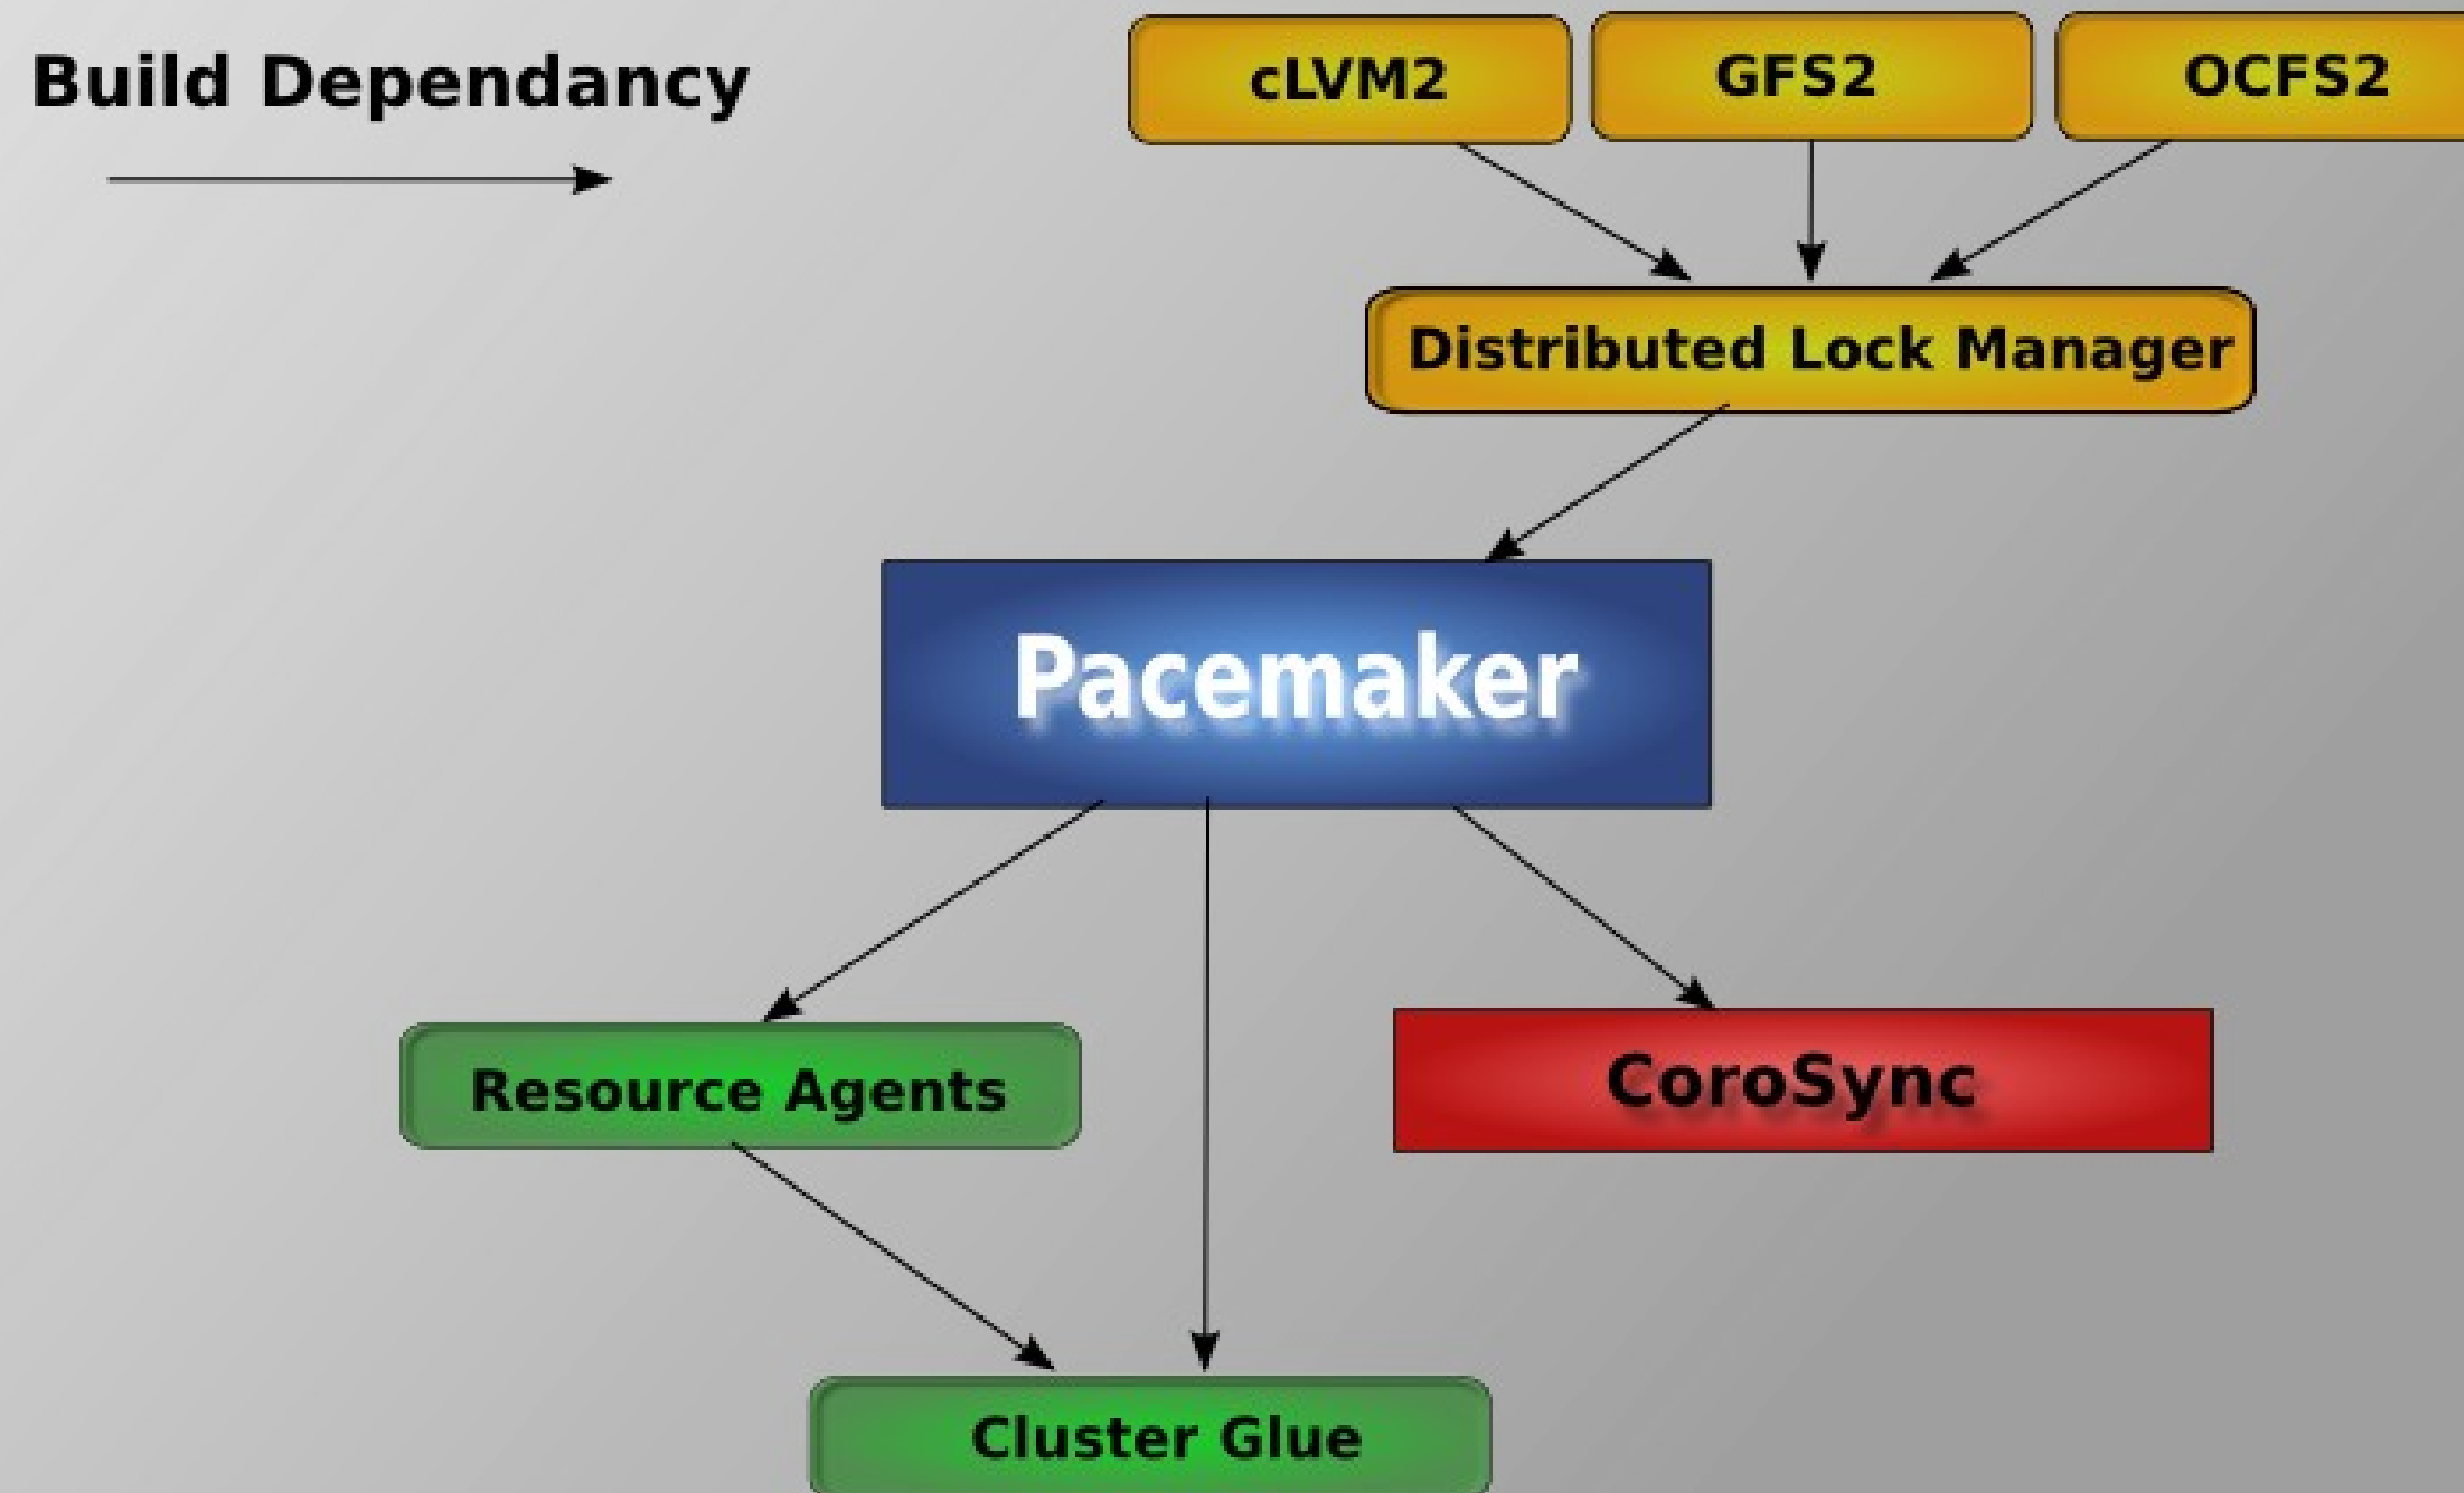
About corosync and pacemaker

- Pacemaker will compute the ideal state of the cluster and plot a path to achieve it after any of these events. This may include moving resources, stopping nodes and even forcing them offline with remote power switches.

The Pacemaker Stack

- When combined with Corosync, Pacemaker also supports popular open source cluster filesystems.
- Due to recent standardization within the cluster filesystem community, they make use of a common distributed lock manager which makes use of Corosync for its messaging capabilities and Pacemaker for its membership (which nodes are up/down) and fencing services.

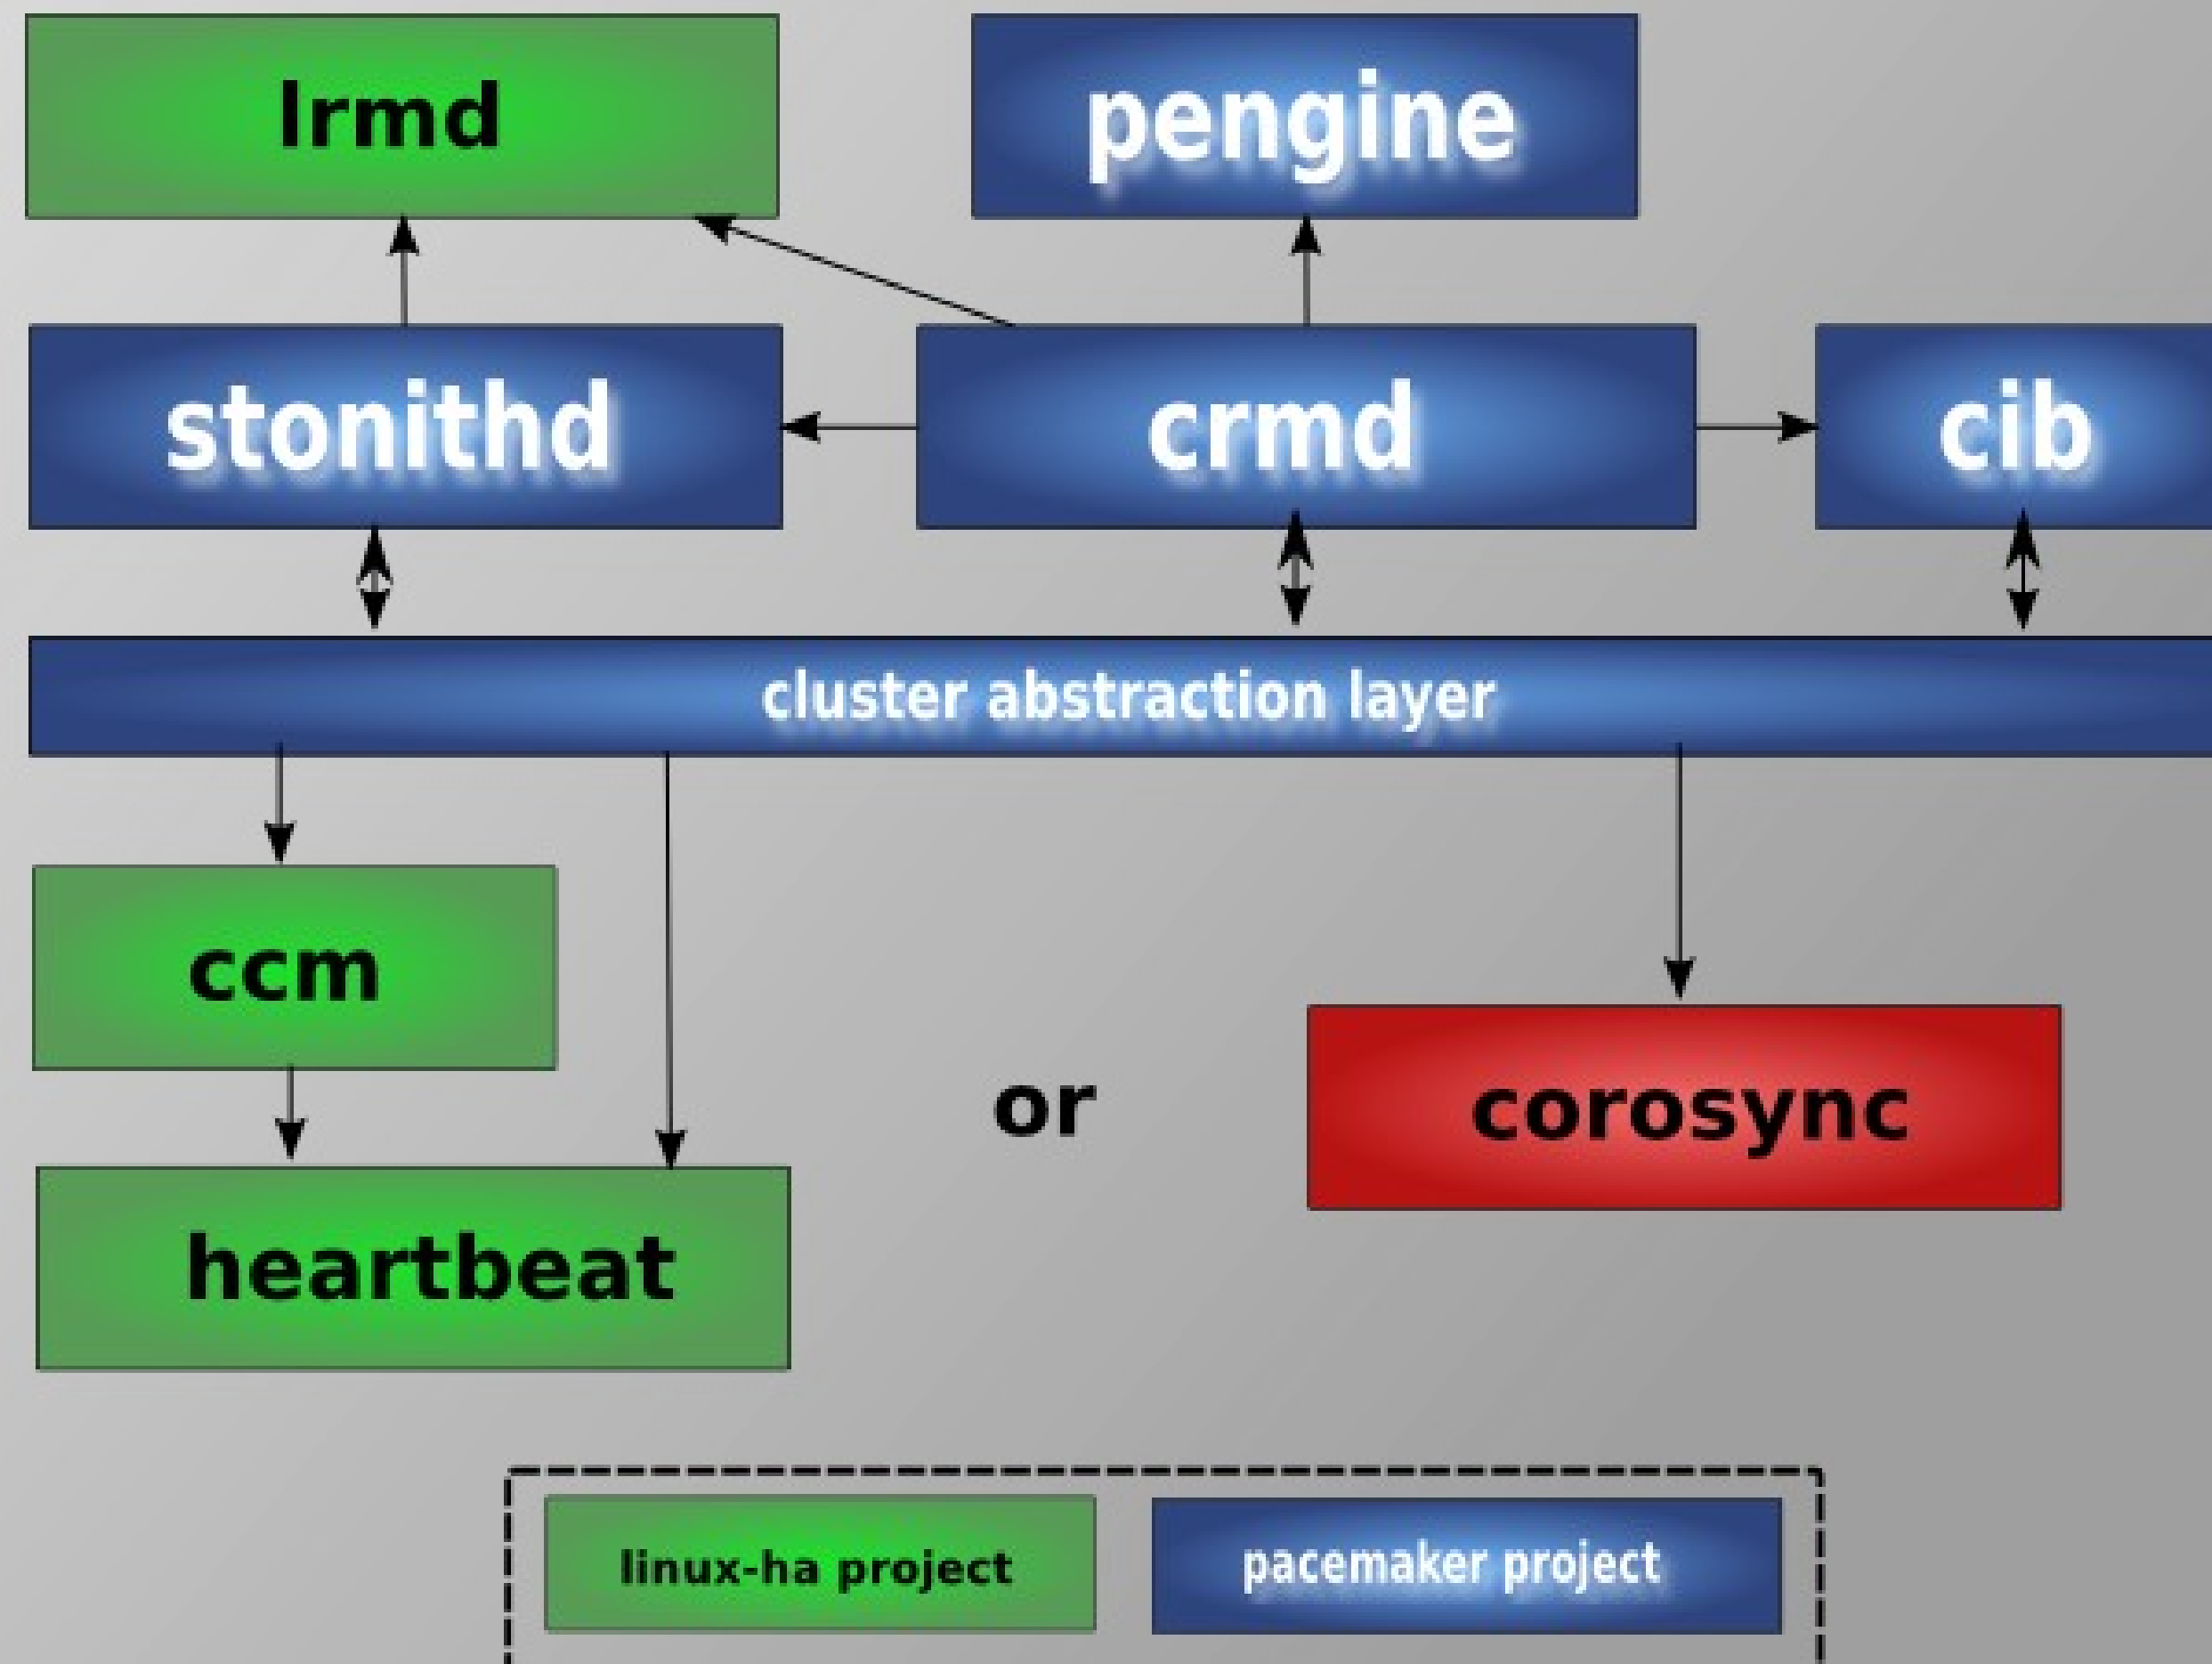
Pacemaker Stack



Internal Components

- Pacemaker itself is composed of four key components (illustrated below in the same color scheme as the previous diagram):
 - CIB (aka. Cluster Information Base)
 - CRMd (aka. Cluster Resource Management daemon)
 - PEngine (aka. PE or Policy Engine)
 - STONITHd

Pacemaker Internals



Internal Components

- The CIB uses XML to represent both the cluster's configuration and current state of all resources in the cluster. The contents of the CIB are automatically kept in sync across the entire cluster and are used by the PEngine to compute the ideal state of the cluster and how it should be achieved.
- This list of instructions is then fed to the DC (Designated Co-ordinator). Pacemaker centralizes all cluster decision making by electing one of the CRMD instances to act as a master. Should the elected CRMD process, or the node it is on, fail... a new one is quickly established.

Internal Components

- The DC carries out the PEngine's instructions in the required order by passing them to either the LRMd (Local Resource Management daemon) or CRMd peers on other nodes via the cluster messaging infrastructure (which in turn passes them on to their LRMd process).
- The peer nodes all report the results of their operations back to the DC and based on the expected and actual results, will either execute any actions that needed to wait for the previous one to complete, or abort processing and ask the PEngine to recalculate the ideal cluster state based on the unexpected results.

Internal Components

- In some cases, it may be necessary to power off nodes in order to protect shared data or complete resource recovery. For this Pacemaker comes with STONITHd. STONITH is an acronym for Shoot-The-Other-Node-In-The-Head and is usually implemented with a remote power switch. In Pacemaker, STONITH devices are modeled as resources (and configured in the CIB) to enable them to be easily monitored for failure, however STONITHd takes care of understanding the STONITH topology such that its clients simply request a node be fenced and it does the rest.

Types of Pacemaker Clusters


- Pacemaker makes no assumptions about your environment, this allows it to support practically any redundancy configuration including Active/Active, Active/Passive, N+1, N+M, N-to-1 and N-to-N.
- In this document we will focus on the setup of a highly available Apache web server with an Active/Passive configuration using GFS2.

Active / Active

Services

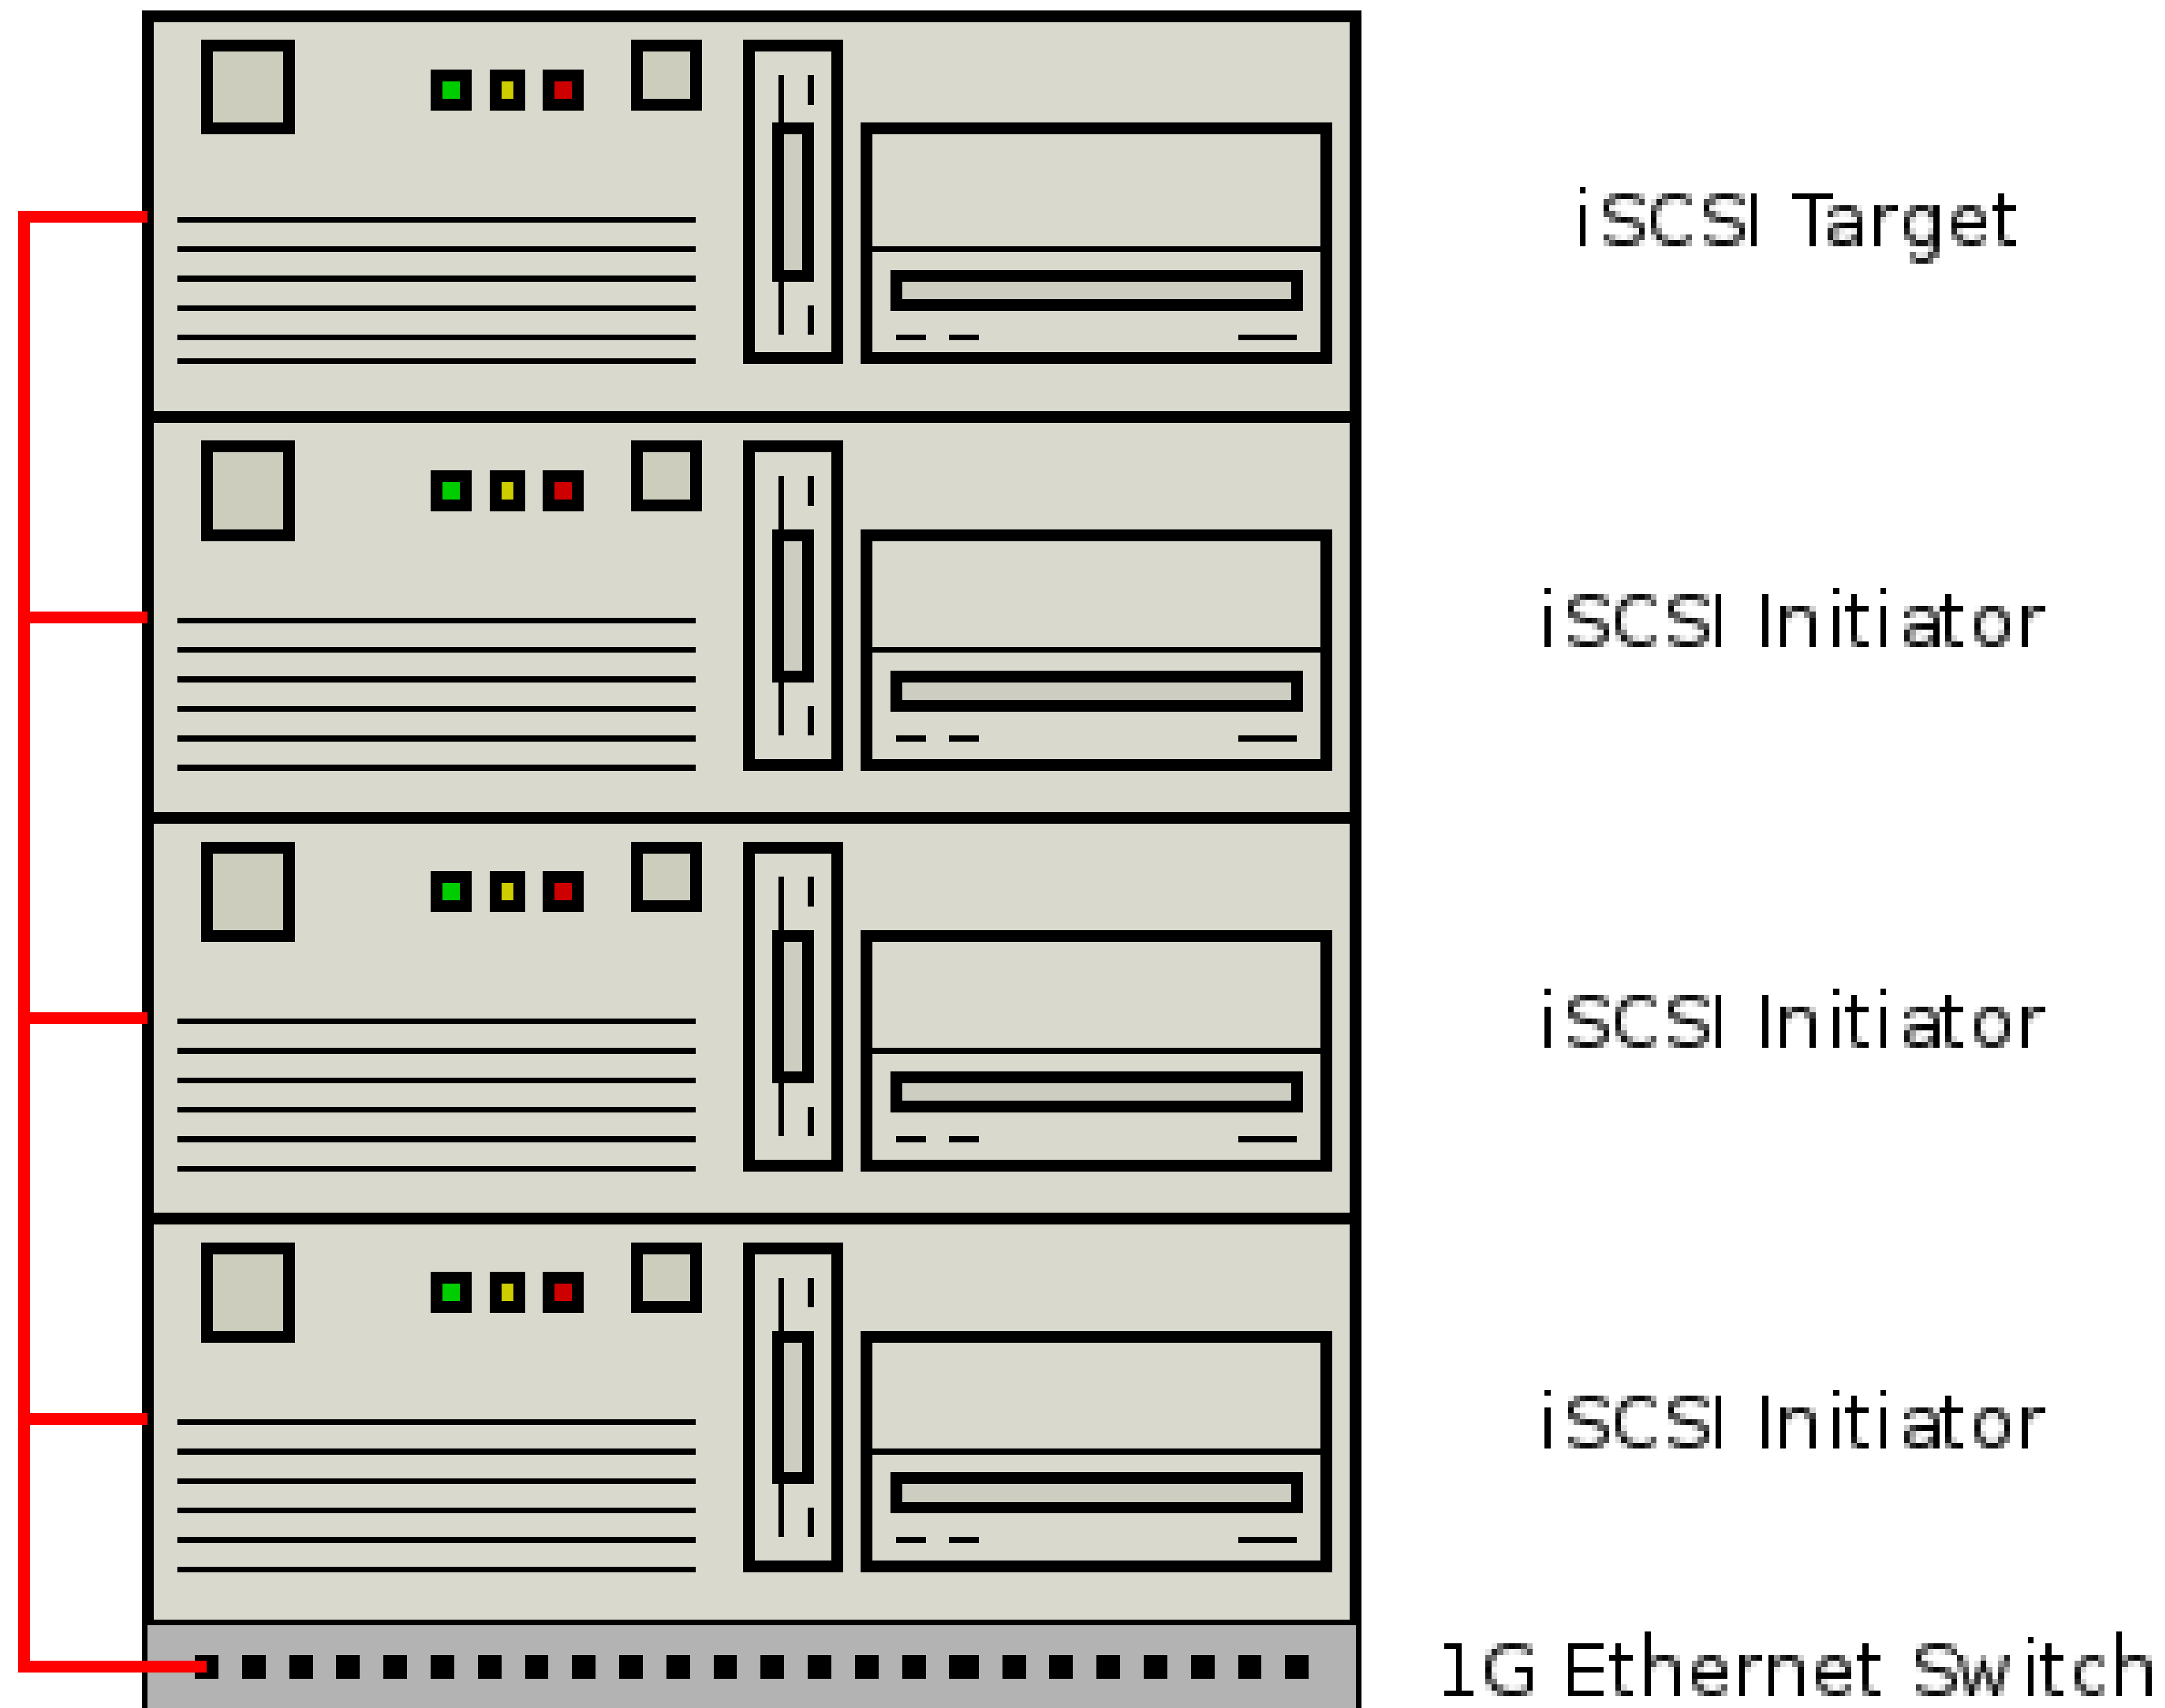
URL	URL	URL	URL
Mail	Mail	Mail	Mail
D'base	D'base	D'base	D'base
Web Site	Web Site	Web Site	Web Site
GFS2	GFS2	GFS2	GFS2

Cluster Software



Hardware





References

Source:

http://clusterlabs.org/doc/en-US/Pacemaker/1.1-pcs/html/Clusters_from_Scratch/_pacemaker_architecture.html

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

INSTALL THE CLUSTERING SOFTWARE

Make sure the correct repositories are set up

- The Cluster software is in the High Availability repo. You'll have access to it if you buy a subscription to this add-on.

Install corosync and pacemaker on the nodes

- `yum -y install lvm2-cluster corosync pacemaker pcs fence-agents-all`
 - lvm2-cluster provides cluster-aware logical volume capabilities
 - corosync and pacemaker (as described before)
 - pcs is the pacemaker and corosync administration tool. It can be used from the command line, and it also provides pcsd, which exposes a web-based UI. We'll use the web UI for today.
 - fence-agents-all provides fence agents for all supported fence devices

root@hideo:~				
File Edit View Search Terminal Help				
Dependencies Resolved				
Package	Arch	Version	Repository	Size
Installing:				
corosync	x86_64	2.3.3-2.el7	ha	186 k
fence-agents-all	x86_64	4.0.2-21.el7	rhel	7.3 k
lvm2-cluster	x86_64	7:2.02.105-13.el7	ha	492 k
pacemaker	x86_64	1.1.10-27.el7	ha	406 k
pcs	x86_64	0.9.115-31.el7	ha	5.0 M
Installing for dependencies:				
OpenIPMI-modalias	x86_64	2.0.19-11.el7	rhel	15 k
autogen-libopts	x86_64	5.18-5.el7	rhel	66 k
corosynclib	x86_64	2.3.3-2.el7	ha	108 k
dlm	x86_64	4.0.2-3.el7	ha	88 k
dlm-lib	x86_64	4.0.2-3.el7	ha	22 k
fence-agents-apc	x86_64	4.0.2-21.el7	rhel	12 k
fence-agents-apc-snmp	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-bladecenter	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-brocade	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-cisco-mds	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-cisco-ucs	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-common	x86_64	4.0.2-21.el7	rhel	44 k
fence-agents-drac5	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-eaton-snmp	x86_64	4.0.2-21.el7	rhel	12 k
fence-agents-eps	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-hpblade	x86_64	4.0.2-21.el7	rhel	10 k
fence-agents-ibmblade	x86_64	4.0.2-21.el7	rhel	10 k
fence-agents-ifmib	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-ilo-mp	x86_64	4.0.2-21.el7	rhel	9.8 k
fence-agents-ilo2	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-intelmodular	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-ipdu	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-ipmilan	x86_64	4.0.2-21.el7	rhel	22 k
fence-agents-kdump	x86_64	4.0.2-21.el7	rhel	21 k
fence-agents-rhev	x86_64	4.0.2-21.el7	rhel	11 k
fence-agents-rsb	x86_64	4.0.2-21.el7	rhel	10 k
fence-agents-scsi	x86_64	4.0.2-21.el7	rhel	14 k
fence-agents-vmware-soap	x86_64	4.0.2-21.el7	rhel	12 k
fence-agents-wti	x86_64	4.0.2-21.el7	rhel	12 k
fence-virt	x86_64	0.3.0-16.el7	rhel	39 k
gnutls-dane	x86_64	3.1.18-8.el7	rhel	51 k
gnutls-utils	x86_64	3.1.18-8.el7	rhel	206 k
ipmitool	x86_64	1.8.13-7.el7	rhel	381 k
ldns	x86_64	1.6.16-7.el7	rhel	473 k
libevent	x86_64	2.0.21-4.el7	rhel	214 k
libqb	x86_64	0.16.0-6.el7	ha	89 k
libtool-ltdl	x86_64	2.4.2-20.el7	rhel	49 k
libyaml	x86_64	0.1.4-8.el7	rhel	54 k
net-snmp-libs	x86_64	1:5.7.2-18.el7	rhel	745 k
net-snmp-utils	x86_64	1:5.7.2-18.el7	rhel	195 k
pacemaker-cli	x86_64	1.1.10-27.el7	ha	208 k
pacemaker-cluster-libs	x86_64	1.1.10-27.el7	ha	79 k
pacemaker-libs	x86_64	1.1.10-27.el7	ha	458 k
perl-TimeDate	noarch	1:2.30-2.el7	rhel	52 k


```
root@hideo:~
File Edit View Search Terminal Help

fence-agents-hpblade      x86_64      4.0.2-21.el7      rhel      10 k
fence-agents-ibmblade     x86_64      4.0.2-21.el7      rhel      10 k
fence-agents-ifmib        x86_64      4.0.2-21.el7      rhel      11 k
fence-agents-ilo-mp       x86_64      4.0.2-21.el7      rhel      9.8 k
fence-agents-ilo2         x86_64      4.0.2-21.el7      rhel      11 k
fence-agents-intelmodular x86_64      4.0.2-21.el7      rhel      11 k
fence-agents-ipdu         x86_64      4.0.2-21.el7      rhel      11 k
fence-agents-ipmilan      x86_64      4.0.2-21.el7      rhel      22 k
fence-agents-kdump        x86_64      4.0.2-21.el7      rhel      21 k
fence-agents-rhevm        x86_64      4.0.2-21.el7      rhel      11 k
fence-agents-rsb          x86_64      4.0.2-21.el7      rhel      10 k
fence-agents-scsi         x86_64      4.0.2-21.el7      rhel      14 k
fence-agents-vmware-soap  x86_64      4.0.2-21.el7      rhel      12 k
fence-agents-wti          x86_64      4.0.2-21.el7      rhel      12 k
fence-virt                x86_64      0.3.0-16.el7      rhel      39 k
gnutls-dane               x86_64      3.1.18-8.el7      rhel      51 k
gnutls-utils              x86_64      3.1.18-8.el7      rhel      206 k
ipmitool                  x86_64      1.8.13-7.el7      rhel      381 k
ldns                      x86_64      1.6.16-7.el7      rhel      473 k
libevent                  x86_64      2.0.21-4.el7      rhel      214 k
libqb                     x86_64      0.16.0-6.el7      ha        89 k
libtool-ltdl              x86_64      2.4.2-20.el7      rhel      49 k
libyaml                   x86_64      0.1.4-8.el7       rhel      54 k
net-snmp-libs              x86_64      1:5.7.2-18.el7    rhel      745 k
net-snmp-utils             x86_64      1:5.7.2-18.el7    rhel      195 k
pacemaker-cli              x86_64      1.1.10-27.el7     ha        208 k
pacemaker-cluster-libs    x86_64      1.1.10-27.el7     ha        79 k
pacemaker-libs             x86_64      1.1.10-27.el7     ha        458 k
perl-TimeDate              noarch      1:2.30-2.el7      rhel      52 k
pexpect                   noarch      2.3-11.el7        rhel      142 k
psmisc                     x86_64      22.20-8.el7       rhel      140 k
python-requests            noarch      1.1.0-8.el7       rhel      70 k
python-six                 noarch      1.3.0-4.el7       rhel      18 k
python-suds                noarch      0.4.1-5.el7       rhel      204 k
python-urllib3             noarch      1.5-8.el7         rhel      41 k
resource-agents            x86_64      3.9.5-26.el7      ha        257 k
ruby                       x86_64      2.0.0.353-19.el7  rhel      66 k
ruby-irb                   noarch      2.0.0.353-19.el7  rhel      87 k
ruby-libs                  x86_64      2.0.0.353-19.el7  rhel      2.8 M
rubygem-bigdecimal         x86_64      1.2.0-19.el7      rhel      78 k
rubygem-io-console         x86_64      0.4.2-19.el7      rhel      49 k
rubygem-json               x86_64      1.7.7-19.el7      rhel      74 k
rubygem-psych              x86_64      2.0.0-19.el7      rhel      76 k
rubygem-rdoc               noarch      4.0.0-19.el7      rhel      317 k
rubygems                   noarch      2.0.14-19.el7     rhel      211 k
sg3_utils                  x86_64      1.37-5.el7        rhel      640 k
telnet                     x86_64      1:0.17-59.el7     rhel      63 k
unbound-libs               x86_64      1.4.20-19.el7     rhel      294 k

Transaction Summary
=====
Install 5 Packages (+63 Dependent packages)

Total download size: 15 M
Installed size: 44 M
Is this ok [y/d/N]: 
```

```
root@hideo:~  
File Edit View Search Terminal Help  
Verifying : fence-agents-rhevm-4.0.2-21.el7.x86_64 37/68  
Verifying : resource-agents-3.9.5-26.el7.x86_64 38/68  
Verifying : fence-agents-apc-snmp-4.0.2-21.el7.x86_64 39/68  
Verifying : libevent-2.0.21-4.el7.x86_64 40/68  
Verifying : gnutls-dane-3.1.18-8.el7.x86_64 41/68  
Verifying : rubygem-io-console-0.4.2-19.el7.x86_64 42/68  
Verifying : fence-agents-rsb-4.0.2-21.el7.x86_64 43/68  
Verifying : l:telnet-0.17-59.el7.x86_64 44/68  
Verifying : OpenIPMI-modalias-2.0.19-11.el7.x86_64 45/68  
Verifying : libtool-ltdl-2.4.2-20.el7.x86_64 46/68  
Verifying : fence-agents-hpblade-4.0.2-21.el7.x86_64 47/68  
Verifying : pacemaker-cli-1.1.10-27.el7.x86_64 48/68  
Verifying : dlm-4.0.2-3.el7.x86_64 49/68  
Verifying : python-suds-0.4.1-5.el7.noarch 50/68  
Verifying : l:perl-TimeDate-2.30-2.el7.noarch 51/68  
Verifying : libqb-0.16.0-6.el7.x86_64 52/68  
Verifying : fence-agents-kdump-4.0.2-21.el7.x86_64 53/68  
Verifying : pacemaker-cluster-libs-1.1.10-27.el7.x86_64 54/68  
Verifying : corosynclib-2.3.3-2.el7.x86_64 55/68  
Verifying : fence-agents-common-4.0.2-21.el7.x86_64 56/68  
Verifying : fence-agents-bladecenter-4.0.2-21.el7.x86_64 57/68  
Verifying : dlm-lib-4.0.2-3.el7.x86_64 58/68  
Verifying : fence-virt-0.3.0-16.el7.x86_64 59/68  
Verifying : fence-agents-intelmodular-4.0.2-21.el7.x86_64 60/68  
Verifying : pexpect-2.3-11.el7.noarch 61/68  
Verifying : pacemaker-libs-1.1.10-27.el7.x86_64 62/68  
Verifying : fence-agents-eaton-snmp-4.0.2-21.el7.x86_64 63/68  
Verifying : python-urllib3-1.5-8.el7.noarch 64/68  
Verifying : fence-agents-drac5-4.0.2-21.el7.x86_64 65/68  
Verifying : fence-agents-cisco-mds-4.0.2-21.el7.x86_64 66/68  
Verifying : fence-agents-apc-4.0.2-21.el7.x86_64 67/68  
Verifying : fence-agents-ilo2-4.0.2-21.el7.x86_64 68/68  
  
Installed:  
corosync.x86_64 0:2.3.3-2.el7 fence-agents-all.x86_64 0:4.0.2-21.el7 lvm2-cluster.x86_64 7:2.02.105-13.el7 pacemaker.x86_64 0:1.1.10-27.el7 pcs.x86_64 0:0.9.115-31.el7  
  
Dependency Installed:  
OpenIPMI-modalias.x86_64 0:2.0.19-11.el7 autogen-libopts.x86_64 0:5.18-5.el7 corosynclib.x86_64 0:2.3.3-2.el7 dlm.x86_64 0:4.0.2-3.el7  
dlm-lib.x86_64 0:4.0.2-3.el7 fence-agents-apc.x86_64 0:4.0.2-21.el7 fence-agents-apc-snmp.x86_64 0:4.0.2-21.el7 fence-agents-bladecenter.x86_64 0:4.0.2-21.el7  
fence-agents-brocade.x86_64 0:4.0.2-21.el7 fence-agents-cisco-mds.x86_64 0:4.0.2-21.el7 fence-agents-cisco-ucs.x86_64 0:4.0.2-21.el7 fence-agents-common.x86_64 0:4.0.2-21.el7  
fence-agents-drac5.x86_64 0:4.0.2-21.el7 fence-agents-eaton-snmp.x86_64 0:4.0.2-21.el7 fence-agents-eps.x86_64 0:4.0.2-21.el7 fence-agents-hpblade.x86_64 0:4.0.2-21.el7  
fence-agents-ibmblade.x86_64 0:4.0.2-21.el7 fence-agents-ifmib.x86_64 0:4.0.2-21.el7 fence-agents-ilo-mp.x86_64 0:4.0.2-21.el7 fence-agents-ilo2.x86_64 0:4.0.2-21.el7  
fence-agents-intelmodular.x86_64 0:4.0.2-21.el7 fence-agents-ipdu.x86_64 0:4.0.2-21.el7 fence-agents-ipmilan.x86_64 0:4.0.2-21.el7 fence-agents-kdump.x86_64 0:4.0.2-21.el7  
fence-agents-rhevm.x86_64 0:4.0.2-21.el7 fence-agents-rsb.x86_64 0:4.0.2-21.el7 fence-agents-scsi.x86_64 0:4.0.2-21.el7 fence-agents-vmware-soap.x86_64 0:4.0.2-21.el7  
fence-agents-wti.x86_64 0:4.0.2-21.el7 fence-virt.x86_64 0:0.3.0-16.el7 gnutls-dane.x86_64 0:3.1.18-8.el7 gnutls-utils.x86_64 0:3.1.18-8.el7  
ipmitool.x86_64 0:1.8.13-7.el7 ldns.x86_64 0:1.6.16-7.el7 libevent.x86_64 0:2.0.21-4.el7 libqb.x86_64 0:0.16.0-6.el7  
libtool-ltdl.x86_64 0:2.4.2-20.el7 libyaml.x86_64 0:0.1.4-8.el7 net-snmp-libs.x86_64 1:5.7.2-18.el7 net-snmp-utils.x86_64 1:5.7.2-18.el7  
pacemaker-cli.x86_64 0:1.1.10-27.el7 pacemaker-cluster-libs.x86_64 0:1.1.10-27.el7 pacemaker-libs.x86_64 0:1.1.10-27.el7 perl-TimeDate.noarch 1:2.30-2.el7  
pexpect.noarch 0:2.3-11.el7 psmisc.x86_64 0:22.20-8.el7 python-requests.noarch 0:1.1.0-8.el7 python-six.noarch 0:1.3.0-4.el7  
python-suds.noarch 0:0.4.1-5.el7 python-urllib3.noarch 0:1.5-8.el7 resource-agents.x86_64 0:3.9.5-26.el7 ruby.x86_64 0:2.0.0.353-19.el7  
ruby-irb.noarch 0:2.0.0.353-19.el7 ruby-libs.x86_64 0:2.0.0.353-19.el7 rubygem-bigdecimal.x86_64 0:1.2.0-19.el7 rubygem-io-console.x86_64 0:0.4.2-19.el7  
rubygem-json.x86_64 0:1.7.7-19.el7 rubygem-psych.x86_64 0:2.0.0-19.el7 rubygem-rdoc.noarch 0:4.0.0-19.el7 rubygems.noarch 0:2.0.14-19.el7  
sg3_utils.x86_64 0:1.37-5.el7 telnet.x86_64 1:0.17-59.el7 unbound-libs.x86_64 0:1.4.20-19.el7  
  
Complete!  
[root@hideo ~]#
```



```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
Installed:  
corosync.x86_64 0:2.3.3-2.el7          fence-agents-all.x86_64 0:4.0.2-21.el7  
  
Dependency Installed:  
OpenIPMI-modalias.x86_64 0:2.0.19-11.el7          autogen-libopts.x86_64 0:  
dlm-lib.x86_64 0:4.0.2-3.el7                    fence-agents-apc.x86_64 0  
fence-agents-brocade.x86_64 0:4.0.2-21.el7        fence-agents-cisco-mds.x8  
fence-agents-drac5.x86_64 0:4.0.2-21.el7          fence-agents-eaton-snmp.x  
fence-agents-ibmblade.x86_64 0:4.0.2-21.el7       fence-agents-ifmib.x86_64  
fence-agents-intelmodular.x86_64 0:4.0.2-21.el7   fence-agents-ipdu.x86_64  
fence-agents-rhevm.x86_64 0:4.0.2-21.el7          fence-agents-rsb.x86_64 0  
fence-agents-wti.x86_64 0:4.0.2-21.el7            fence-virt.x86_64 0:0.3.0  
ipmitool.x86_64 0:1.8.13-7.el7                    ldns.x86_64 0:1.6.16-7.el  
libtool-ltdl.x86_64 0:2.4.2-20.el7               libyaml.x86_64 0:0.1.4-8.  
pacemaker-cli.x86_64 0:1.1.10-27.el7             pacemaker-cluster-libs.x8  
pexpect.noarch 0:2.3-11.el7                      psmisc.x86_64 0:22.20-8.e  
python-suds.noarch 0:0.4.1-5.el7                  python-urllib3.noarch 0:1  
ruby-irb.noarch 0:2.0.0.353-19.el7                ruby-libs.x86_64 0:2.0.0.  
rubygem-json.x86_64 0:1.7.7-19.el7               rubygem-psych.x86_64 0:2.  
sg3_utils.x86_64 0:1.37-5.el7                    telnet.x86_64 1:0.17-59.e  
  
Complete!  
[root@hideo ~]#
```

```
root@lady3jane:~  
File Edit View Search Terminal Help  
pacemaker-libs.x86_64 0:1.1.10-27.el7  
perl-TimeDate.noarch 1:2.30-2.el7  
pexpect.noarch 0:2.3-11.el7  
psmisc.x86_64 0:22.20-8.el7  
python-requests.noarch 0:1.1.0-8.el7  
python-six.noarch 0:1.3.0-4.el7  
python-suds.noarch 0:0.4.1-5.el7  
python-urllib3.noarch 0:1.5-8.el7  
resource-agents.x86_64 0:3.9.5-26.el7  
ruby.x86_64 0:2.0.0.353-19.el7  
ruby-irb.noarch 0:2.0.0.353-19.el7  
ruby-libs.x86_64 0:2.0.0.353-19.el7  
rubygem-bigdecimal.x86_64 0:1.2.0-19.el7  
rubygem-io-console.x86_64 0:0.4.2-19.el7  
rubygem-json.x86_64 0:1.7.7-19.el7  
rubygem-psych.x86_64 0:2.0.0-19.el7  
rubygem-rdoc.noarch 0:4.0.0-19.el7  
rubygems.noarch 0:2.0.14-19.el7  
sg3_utils.x86_64 0:1.37-5.el7  
telnet.x86_64 1:0.17-59.el7  
unbound-libs.x86_64 0:1.4.20-19.el7  
  
Complete!  
[root@lady3jane ~]#
```

```
root@maelcum:~  
File Edit View Search Terminal Help  
pacemaker-libs.x86_64 0:1.1.10-27.el7  
perl-TimeDate.noarch 1:2.30-2.el7  
pexpect.noarch 0:2.3-11.el7  
psmisc.x86_64 0:22.20-8.el7  
python-requests.noarch 0:1.1.0-8.el7  
python-six.noarch 0:1.3.0-4.el7  
python-suds.noarch 0:0.4.1-5.el7  
python-urllib3.noarch 0:1.5-8.el7  
resource-agents.x86_64 0:3.9.5-26.el7  
ruby.x86_64 0:2.0.0.353-19.el7  
ruby-irb.noarch 0:2.0.0.353-19.el7  
ruby-libs.x86_64 0:2.0.0.353-19.el7  
rubygem-bigdecimal.x86_64 0:1.2.0-19.el7  
rubygem-io-console.x86_64 0:0.4.2-19.el7  
rubygem-json.x86_64 0:1.7.7-19.el7  
rubygem-psych.x86_64 0:2.0.0-19.el7  
rubygem-rdoc.noarch 0:4.0.0-19.el7  
rubygems.noarch 0:2.0.14-19.el7  
sg3_utils.x86_64 0:1.37-5.el7  
telnet.x86_64 1:0.17-59.el7  
unbound-libs.x86_64 0:1.4.20-19.el7  
  
Complete!  
[root@maelcum ~]#
```

Enable and start the pcsd service

- `systemctl enable pcsd.service`
- `systemctl start pcsd.service`


```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
pexpect.noarch 0:2.3-11.el7  
python-suds.noarch 0:0.4.1-5.el7  
ruby-irb.noarch 0:2.0.0.353-19.el7  
rubygem-json.x86_64 0:1.7.7-19.el7  
sg3_utils.x86_64 0:1.37-5.el7  
psmisc.x86_64 0:22.20-8.e  
python-urllib3.noarch 0:1  
ruby-libs.x86_64 0:2.0.0.  
rubygem-psych.x86_64 0:2.  
telnet.x86_64 1:0.17-59.e  
  
Complete!  
[root@hideo ~]# systemctl enable pcsd.service  
ln -s '/usr/lib/systemd/system/pcsd.service' '/etc/systemd/system/multi-user.target.wants/pcsd.service'  
[root@hideo ~]# systemctl start pcsd.service  
[root@hideo ~]# systemctl status pcsd.service  
pcsd.service - PCS GUI and remote configuration interface  
Loaded: loaded (/usr/lib/systemd/system/pcsd.service; enabled)  
Active: active (running) since Sun 2014-03-23 17:52:13 CDT; 18s ago  
Main PID: 2432 (pcsd)  
CGroup: /system.slice/pcsd.service  
└─2432 /bin/sh /usr/lib/pcsd/pcsd start  
└─2436 /bin/bash -c ulimit -S -c 0 >/dev/null 2>&1 ; /usr/bin/ruby...  
└─2437 /usr/bin/ruby -I/usr/lib/pcsd /usr/lib/pcsd/ssl.rb  
  
Mar 23 17:52:13 hideo.tc.redhat.com systemd[1]: Started PCS GUI and remote c....  
Hint: Some lines were ellipsized, use -l to show in full.  
[root@hideo ~]#
```

```
root@lady3jane:~  
File Edit View Search Terminal Help  
rubygem-rdoc.noarch 0:4.0.0-19.el7  
rubygems.noarch 0:2.0.14-19.el7  
sg3_utils.x86_64 0:1.37-5.el7  
telnet.x86_64 1:0.17-59.el7  
unbound-libs.x86_64 0:1.4.20-19.el7  
  
Complete!  
[root@lady3jane ~]# systemctl enable pcsd.service  
ln -s '/usr/lib/systemd/system/pcsd.service' '/etc/systemd/system/multi-user.target.wants/pcsd.service'  
[root@lady3jane ~]# systemctl start pcsd.service  
[root@lady3jane ~]# systemctl status pcsd.service  
pcsd.service - PCS GUI and remote configuration interface  
Loaded: loaded (/usr/lib/systemd/system/pcsd.service; enabled)  
Active: active (running) since Sun 2014-03-23 17:52:16 CDT; 6s ago  
Main PID: 27240 (pcsd)  
CGroup: /system.slice/pcsd.service  
└─27240 /bin/sh /usr/lib/pcsd/pcsd start  
└─27244 /bin/bash -c ulimit -S -c 0 >/dev/null 2>&1 ; /usr/bin/rub...  
└─27245 /usr/bin/ruby -I/usr/lib/pcsd /usr/lib/pcsd/ssl.rb  
  
Mar 23 17:52:16 lady3jane.tc.redhat.com systemd[1]: Started PCS GUI and remot...  
Hint: Some lines were ellipsized, use -l to show in full.  
[root@lady3jane ~]#
```

```
root@maelcum:~  
File Edit View Search Terminal Help  
rubygem-rdoc.noarch 0:4.0.0-19.el7  
rubygems.noarch 0:2.0.14-19.el7  
sg3_utils.x86_64 0:1.37-5.el7  
telnet.x86_64 1:0.17-59.el7  
unbound-libs.x86_64 0:1.4.20-19.el7  
  
Complete!  
[root@maelcum ~]# systemctl enable pcsd.service  
ln -s '/usr/lib/systemd/system/pcsd.service' '/etc/systemd/system/multi-user.target.wants/pcsd.service'  
[root@maelcum ~]# systemctl start pcsd.service  
[root@maelcum ~]# systemctl status pcsd.service  
pcsd.service - PCS GUI and remote configuration interface  
Loaded: loaded (/usr/lib/systemd/system/pcsd.service; enabled)  
Active: active (running) since Sun 2014-03-23 17:52:15 CDT; 15s ago  
Main PID: 27219 (pcsd)  
CGroup: /system.slice/pcsd.service  
└─27219 /bin/sh /usr/lib/pcsd/pcsd start  
└─27223 /bin/bash -c ulimit -S -c 0 >/dev/null 2>&1 ; /usr/bin/rub...  
└─27224 /usr/bin/ruby -I/usr/lib/pcsd /usr/lib/pcsd/ssl.rb  
  
Mar 23 17:52:15 maelcum.tc.redhat.com systemd[1]: Started PCS GUI and remote ...  
Hint: Some lines were ellipsized, use -l to show in full.  
[root@maelcum ~]#
```

Administrative user

- For now, the clustering software used the “hacluster” account for administration. Set this user's password across all the nodes:
- `echo [password] | passwd --stdin hacluster`

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# echo red22hat | passwd --stdin hacluster  
Changing password for user hacluster.  
passwd: all authentication tokens updated successfully.  
[root@hideo ~]#
```

```
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# echo red22hat | passwd --stdin hacluster  
Changing password for user hacluster.  
passwd: all authentication tokens updated successfully.  
[root@lady3jane ~]#
```

```
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# echo red22hat | passwd --stdin hacluster  
Changing password for user hacluster.  
passwd: all authentication tokens updated successfully.  
[root@maelcum ~]#
```




10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE THE CLUSTER

Authorize the Nodes

- From any node: `pcs cluster auth [node, node, node]`
- Note that the user configuration file and security token files are created in `/var/lib/pcsd`

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # target - do not install pacemaker, corosync, etc  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# ls /var/lib/pcsd/  
pcsd.cookiesecret pcsd.crt pcsd.key  
[root@hideo ~]# pcs cluster auth hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com  
Username: hacluster  
Password:  
hideo.tc.redhat.com: Authorized  
lady3jane.tc.redhat.com: Authorized  
maelcum.tc.redhat.com: Authorized  
[root@hideo ~]# ls /var/lib/pcsd/  
pcsd.cookiesecret pcsd.crt pcsd.key pcs_users.conf tokens  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# date  
Fri Mar 21 16:58:32 CDT 2014  
[root@maelcum ~]# ls -l /var/lib/pcsd/  
total 20  
-rwx-----. 1 root root 60 Mar 21 16:57 pcsd.cookiesecret  
-rwx-----. 1 root root 1224 Mar 21 16:57 pcsd.crt  
-rwx-----. 1 root root 1675 Mar 21 16:57 pcsd.key  
-rw-r--r--. 1 root root 536 Mar 21 16:58 pcs_users.conf  
-rw-----. 1 root root 195 Mar 21 16:58 tokens  
[root@maelcum ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# date  
Fri Mar 21 16:58:44 CDT 2014  
[root@lady3jane ~]# ls -l /var/lib/pcsd/  
total 20  
-rwx-----. 1 root root 60 Mar 21 16:57 pcsd.cookiesecret  
-rwx-----. 1 root root 1229 Mar 21 16:57 pcsd.crt  
-rwx-----. 1 root root 1679 Mar 21 16:57 pcsd.key  
-rw-r--r--. 1 root root 536 Mar 21 16:58 pcs_users.conf  
-rw-----. 1 root root 195 Mar 21 16:58 tokens  
[root@lady3jane ~]#
```


Set up the cluster

- From any node: `pcs cluster setup --name [node, node, node]`
- Note that there is now a `corosync.conf` file in `/etc/corosync`

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# ls /etc/corosync/
corosync.conf.example      corosync.xml.example
corosync.conf.example.udpu uidgid.d
[root@hideo ~]# pcs cluster setup --name summit hideo.tc.redhat.com maelcum.tc.r
edhat.com lady3jane.tc.redhat.com
Shutting down pacemaker/corosync services...
Redirecting to /bin/systemctl stop pacemaker.service
Redirecting to /bin/systemctl stop corosync.service
Killing any remaining services...
Removing all cluster configuration files...
hideo.tc.redhat.com: Succeeded
maelcum.tc.redhat.com: Succeeded
lady3jane.tc.redhat.com: Succeeded
[root@hideo ~]# ls /etc/corosync/
corosync.conf      corosync.conf.example.udpu uidgid.d
corosync.conf.example corosync.xml.example
[root@hideo ~]#
```

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # target - do not install pacemaker, corosync, etc  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# pcs cluster setup --name summit hideo.tc.redhat.com maelcum.tc.r  
edhat.com lady3jane.tc.redhat.com  
Shutting down pacemaker/corosync services...  
Redirecting to /bin/systemctl stop pacemaker.service  
Redirecting to /bin/systemctl stop corosync.service  
Killing any remaining services...  
Removing all cluster configuration files...  
hideo.tc.redhat.com: Succeeded  
maelcum.tc.redhat.com: Succeeded  
lady3jane.tc.redhat.com: Succeeded  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# pcs status  
Error: cluster is not currently running on this node  
[root@maelcum ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# pcs status  
Error: cluster is not currently running on this node  
[root@lady3jane ~]#
```

This generates the corosync config file

- `/etc/corosync/corosync.conf`

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# cat /etc/corosync/corosync.conf  
totem {  
  version: 2  
  secauth: off  
  cluster_name: summit  
  transport: udpu  
}  
  
nodelist {  
  node {  
    ring0_addr: hideo.tc.redhat.com  
    nodeid: 1  
  }  
  node {  
    ring0_addr: maelcum.tc.redhat.com  
    nodeid: 2  
  }  
  node {  
    ring0_addr: lady3jane.tc.redhat.com  
    nodeid: 3  
  }  
}  
  
quorum {  
  provider: corosync_votequorum  
}  
  
logging {  
  to_syslog: yes  
}  
[root@hideo ~]#
```

Enable the cluster services on all nodes

- From any node: `pcs cluster enable --all`
- Note that this does not start the services, it only enables them


```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # target - do not install pacemaker, corosync, etc  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# pcs cluster enable --all  
hideo.tc.redhat.com: Cluster Enabled  
maelcum.tc.redhat.com: Cluster Enabled  
lady3jane.tc.redhat.com: Cluster Enabled  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# pcs status  
Error: cluster is not currently running on this node  
[root@maelcum ~]# pcs status  
Error: cluster is not currently running on this node  
[root@maelcum ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# pcs status  
Error: cluster is not currently running on this node  
[root@lady3jane ~]# pcs status  
Error: cluster is not currently running on this node  
[root@lady3jane ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# systemctl status corosync  
corosync.service - Corosync Cluster Engine  
    Loaded: loaded (/usr/lib/systemd/system/corosync.service; enabled)  
    Active: inactive (dead)  
  
Mar 22 12:14:48 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Mar 22 12:15:23 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Mar 22 12:15:48 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Mar 22 12:15:58 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Mar 22 12:17:50 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Mar 22 12:17:57 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Mar 22 12:18:04 hideo.tc.redhat.com systemd[1]: Stopped Corosync Cluster Engine.  
Hint: Some lines were ellipsized, use -l to show in full.  
[root@hideo ~]#
```

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# systemctl status pacemaker
```

```
pacemaker.service - Pacemaker High Availability Cluster Manager
```

```
Loaded: loaded (/usr/lib/systemd/system/pacemaker.service; enabled)
```

```
Active: inactive (dead)
```

```
Mar 22 12:14:48 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Mar 22 12:15:23 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Mar 22 12:15:48 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Mar 22 12:15:58 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Mar 22 12:17:50 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Mar 22 12:17:57 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Mar 22 12:18:03 hideo.tc.redhat.com systemd[1]: Stopped Pacemaker High Avail....
```

```
Hint: Some lines were ellipsized, use -l to show in full.
```

```
[root@hideo ~]#
```

Start the cluster

- From any node: `pcs cluster start --all`

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# pcs cluster start --all
maelcum.tc.redhat.com: Starting Cluster...
lady3jane.tc.redhat.com: Starting Cluster...
hideo.tc.redhat.com: Starting Cluster...
[root@hideo ~]#
```

It will take a few seconds for the nodes to sync

- This is normal and expected


```
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# pcs status  
Cluster name: summit  
WARNING: no stonith devices and stonith-enabled is not false  
Last updated: Fri Mar 21 17:08:16 2014  
Last change: Fri Mar 21 17:07:55 2014 via crmd on maelcum.tc.redhat.com  
Current DC: NONE  
3 Nodes configured  
0 Resources configured  
  
Node hideo.tc.redhat.com (1): UNCLEAN (offline)  
Node lady3jane.tc.redhat.com (3): UNCLEAN (offline)  
Node maelcum.tc.redhat.com (2): UNCLEAN (offline)  
  
Full list of resources:  
  
PCSD Status:  
  hideo.tc.redhat.com: Online  
  maelcum.tc.redhat.com: Online  
  lady3jane.tc.redhat.com: Online  
  
Daemon Status:  
  corosync: active/enabled  
  pacemaker: active/enabled  
  pcsd: active/enabled  
[root@maelcum ~]#
```

```
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# pcs status  
Cluster name: summit  
WARNING: no stonith devices and stonith-enabled is not false  
Last updated: Fri Mar 21 17:08:27 2014  
Last change: Fri Mar 21 17:08:16 2014 via crmd on maelcum.tc.redhat.com  
Stack: corosync  
Current DC: maelcum.tc.redhat.com (2) - partition with quorum  
Version: 1.1.10-27.el7-368c726  
3 Nodes configured  
0 Resources configured  
  
Online: [ hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com ]  
  
Full list of resources:  
  
PCSD Status:  
  hideo.tc.redhat.com: Online  
  maelcum.tc.redhat.com: Online  
  lady3jane.tc.redhat.com: Online  
  
Daemon Status:  
  corosync: active/enabled  
  pacemaker: active/enabled  
  pcsd: active/enabled  
[root@lady3jane ~]#
```

Verify Corosync Installation

- `corosync-cfgtool -s`

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# corosync-cfgtool -s
```

Printing ring status.

Local node ID 1

RING ID 0

id = 172.31.100.22

status = ring 0 active with no faults

```
[root@hideo ~]#
```

Verify Corosync Installation

- `corosync-cmapctl | grep members`

```
[root@hideo ~]# corosync-cmapctl | grep members
runtime.totem.pg.mrp.srp.members.1.config_version (u64) = 0
runtime.totem.pg.mrp.srp.members.1.ip (str) = r(0) ip(172.31.100.22)
runtime.totem.pg.mrp.srp.members.1.join_count (u32) = 1
runtime.totem.pg.mrp.srp.members.1.status (str) = joined
runtime.totem.pg.mrp.srp.members.2.config_version (u64) = 0
runtime.totem.pg.mrp.srp.members.2.ip (str) = r(0) ip(172.31.100.10)
runtime.totem.pg.mrp.srp.members.2.join_count (u32) = 1
runtime.totem.pg.mrp.srp.members.2.status (str) = joined
runtime.totem.pg.mrp.srp.members.3.config_version (u64) = 0
runtime.totem.pg.mrp.srp.members.3.ip (str) = r(0) ip(172.31.100.21)
runtime.totem.pg.mrp.srp.members.3.join_count (u32) = 1
runtime.totem.pg.mrp.srp.members.3.status (str) = joined
[root@hideo ~]#
```


Verify Corosync Installation

- `crm_verify -L -V`

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# crm_verify -L -V  
    error: unpack_resources:      Resource start-up disabled since no STONITH reso  
urces have been defined  
    error: unpack_resources:      Either configure some or disable STONITH with th  
e stonith-enabled option  
    error: unpack_resources:      NOTE: Clusters with shared data need STONITH to  
ensure data integrity  
Errors found during check: config not valid  
[root@hideo ~]#
```



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

LOG INTO THE WEB UI

Connect via https to port 2224

- Any system in the cluster can be used (no more single management node!).

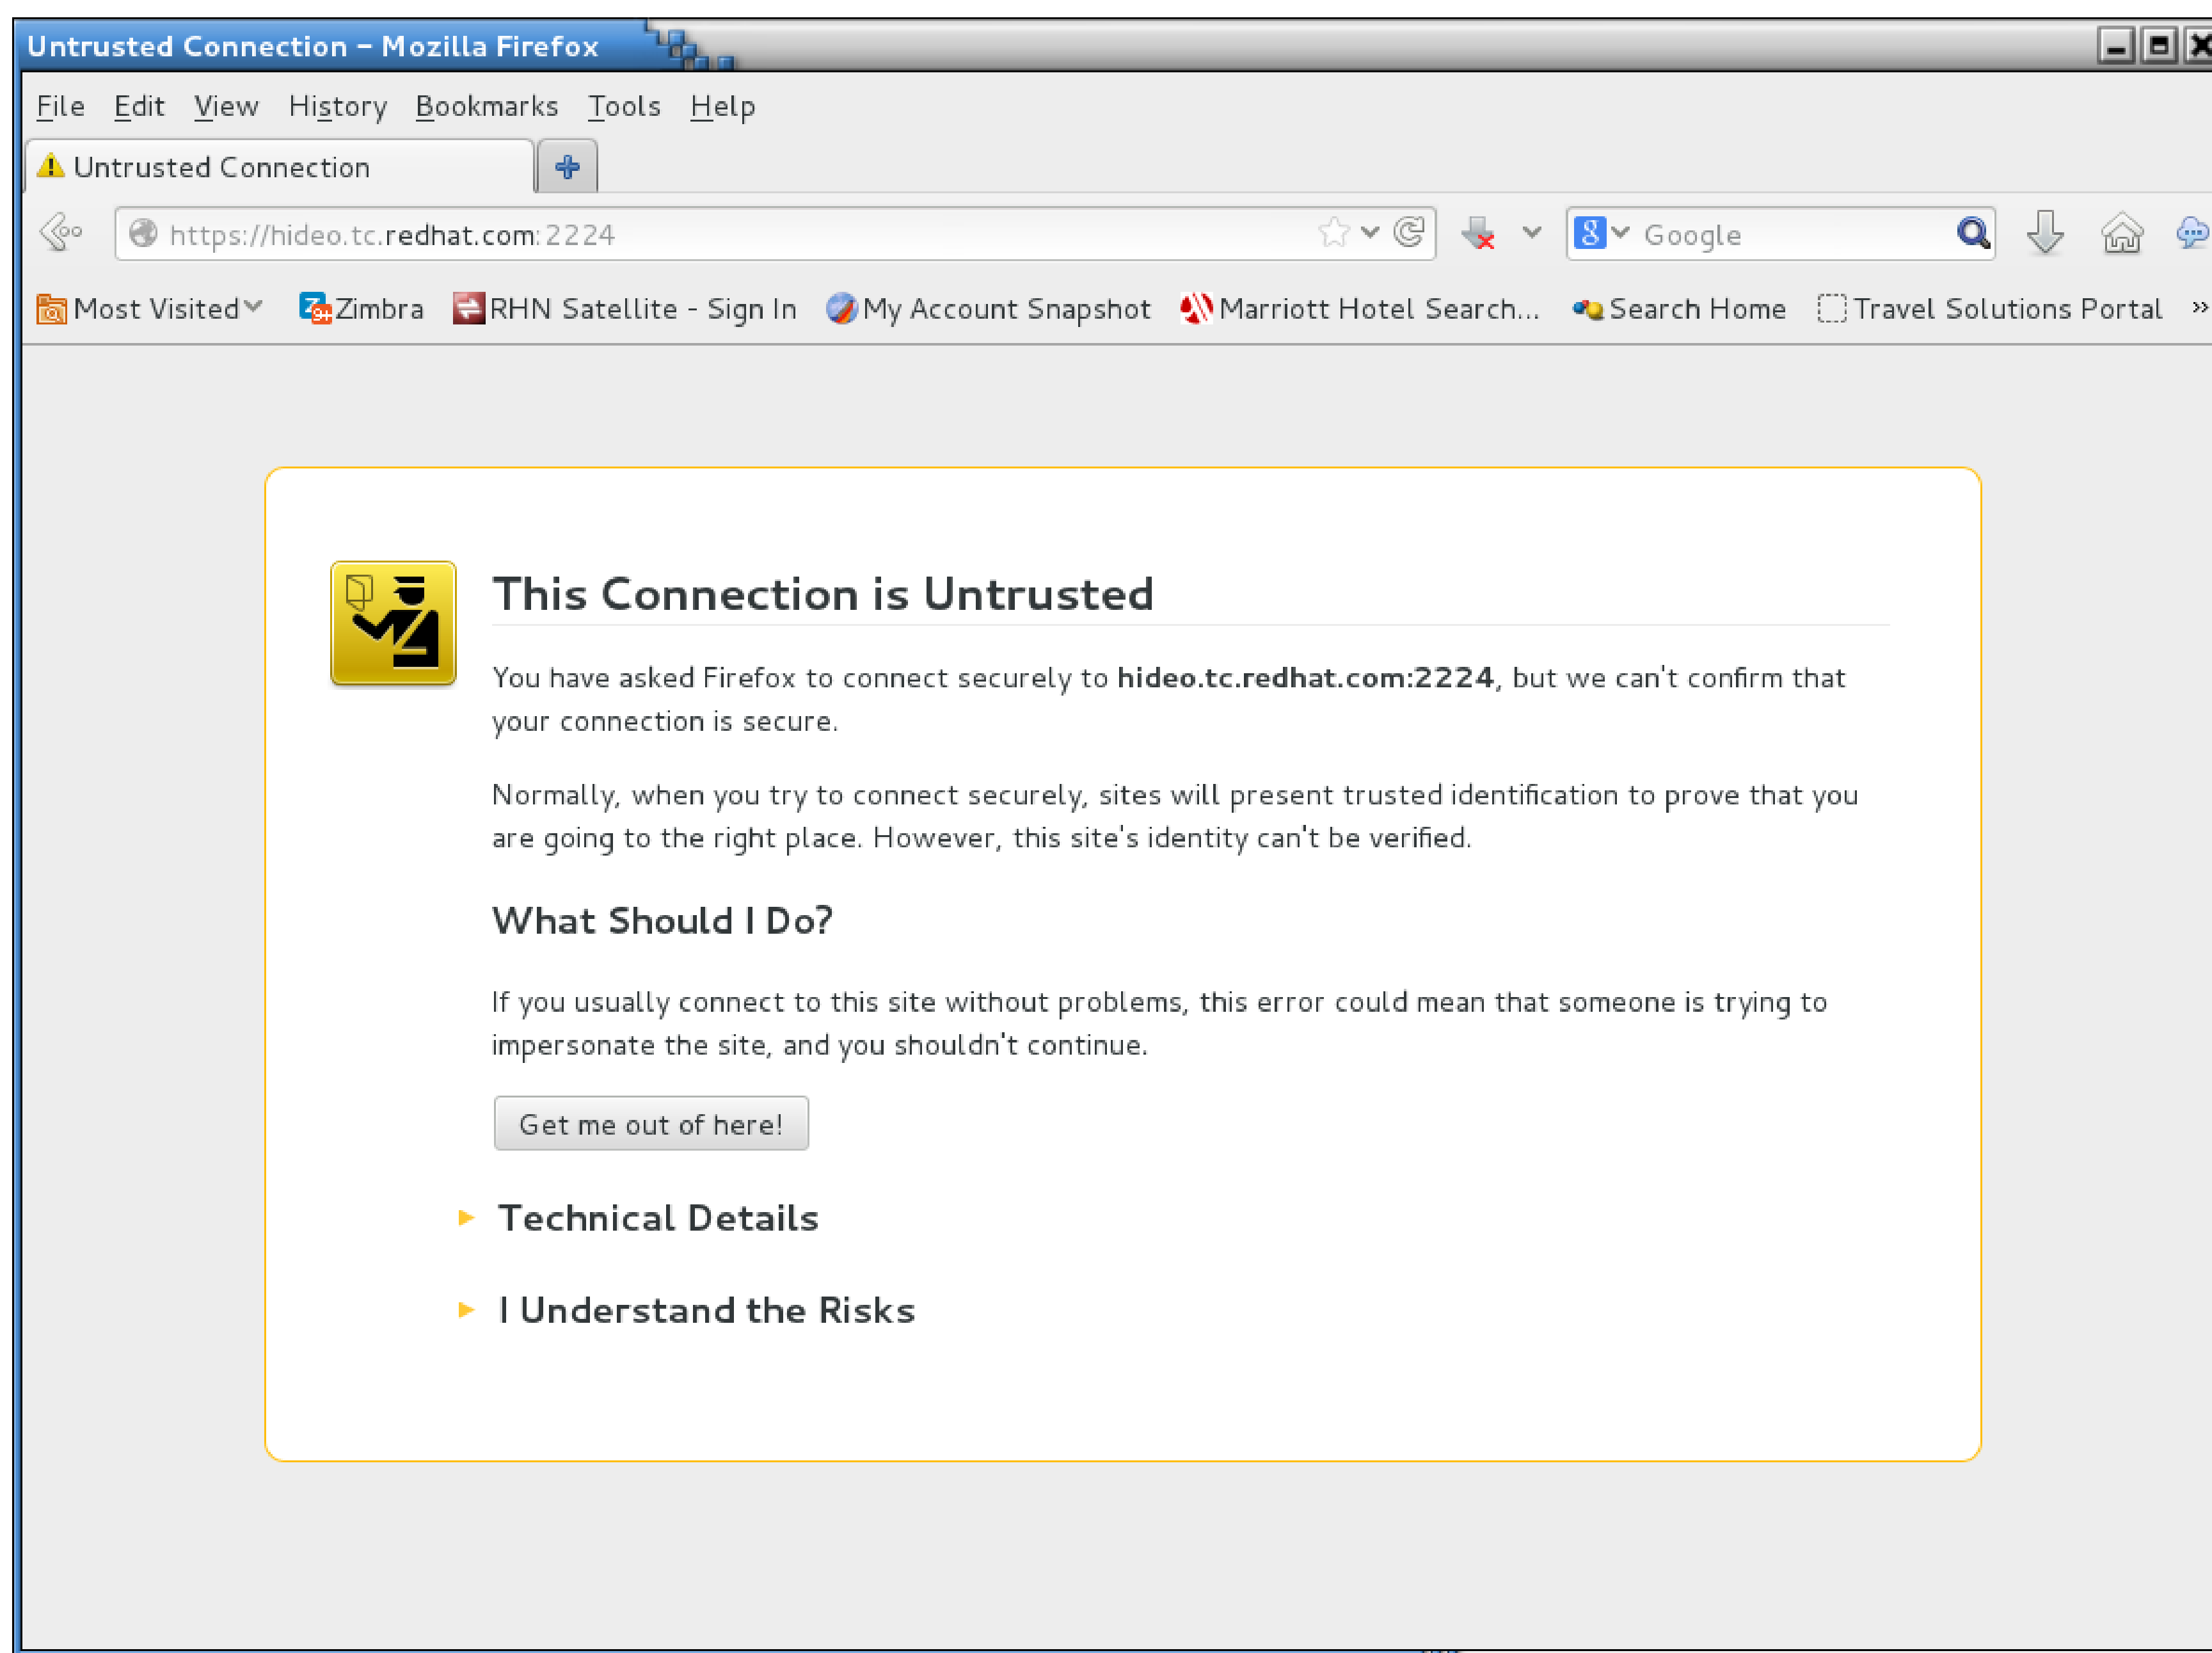
Connect via https to port 2224

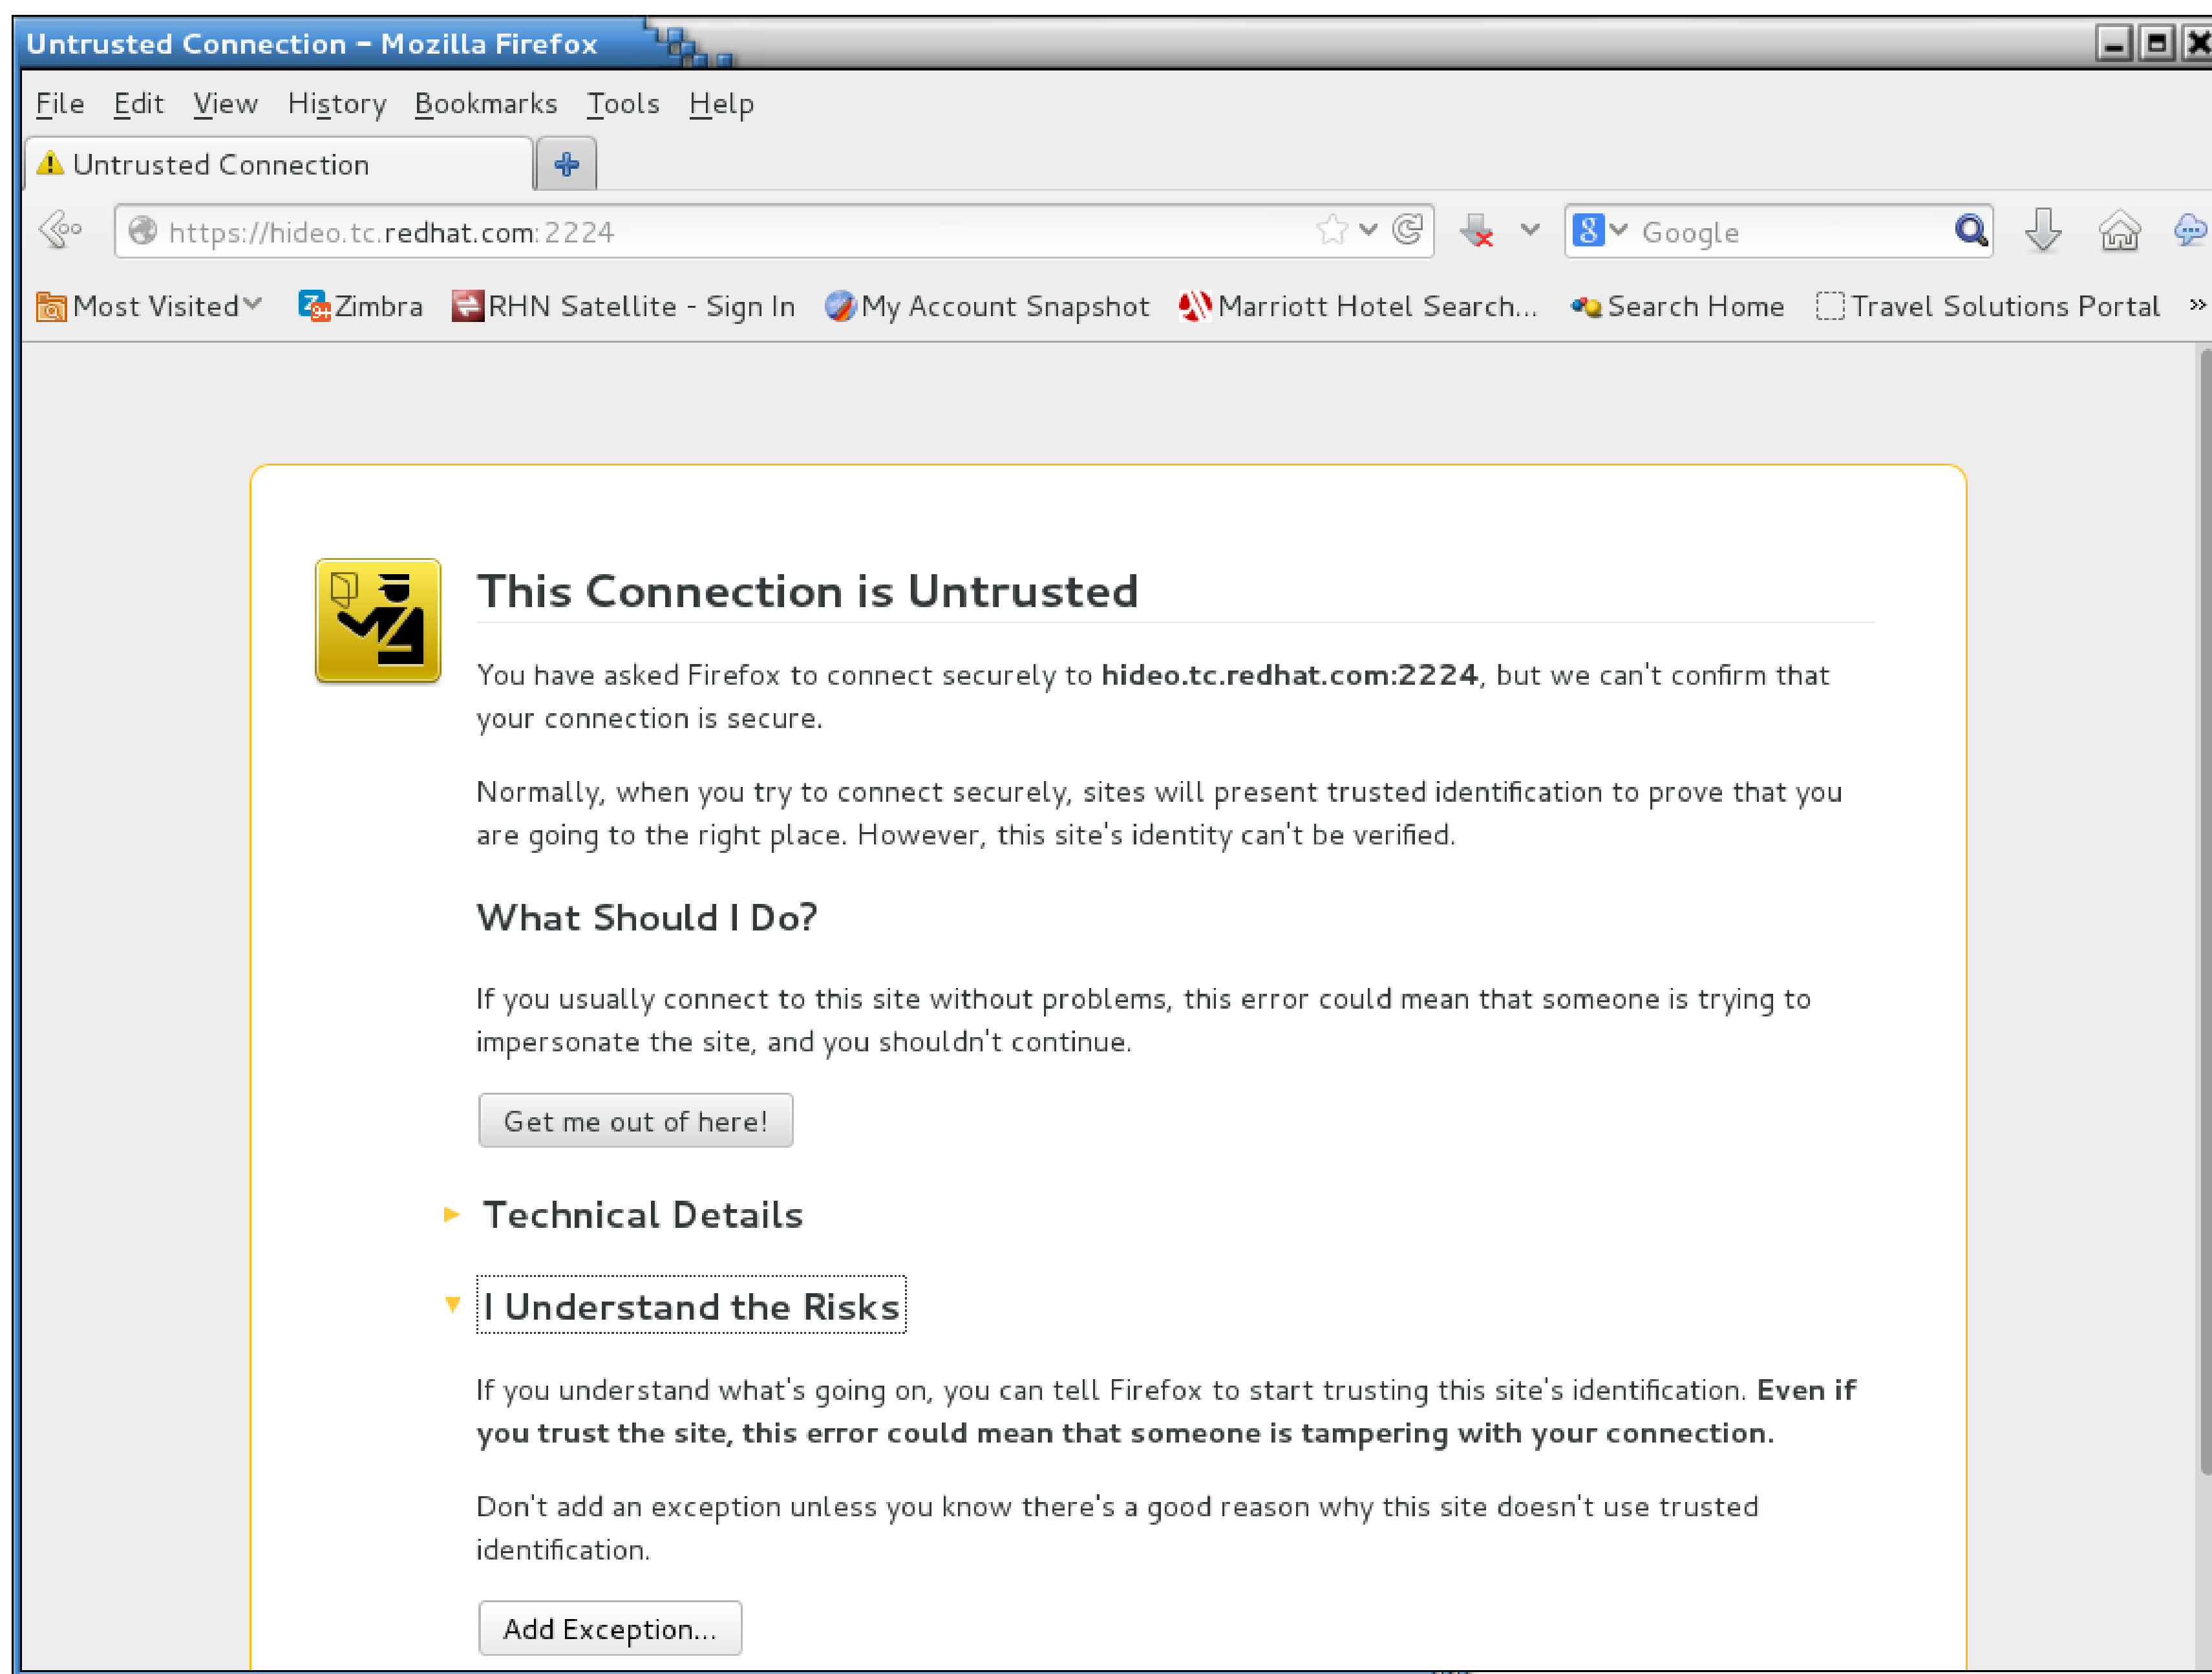
- Any system in the cluster can be used (no more single management node!).
- And there was much rejoicing!!!

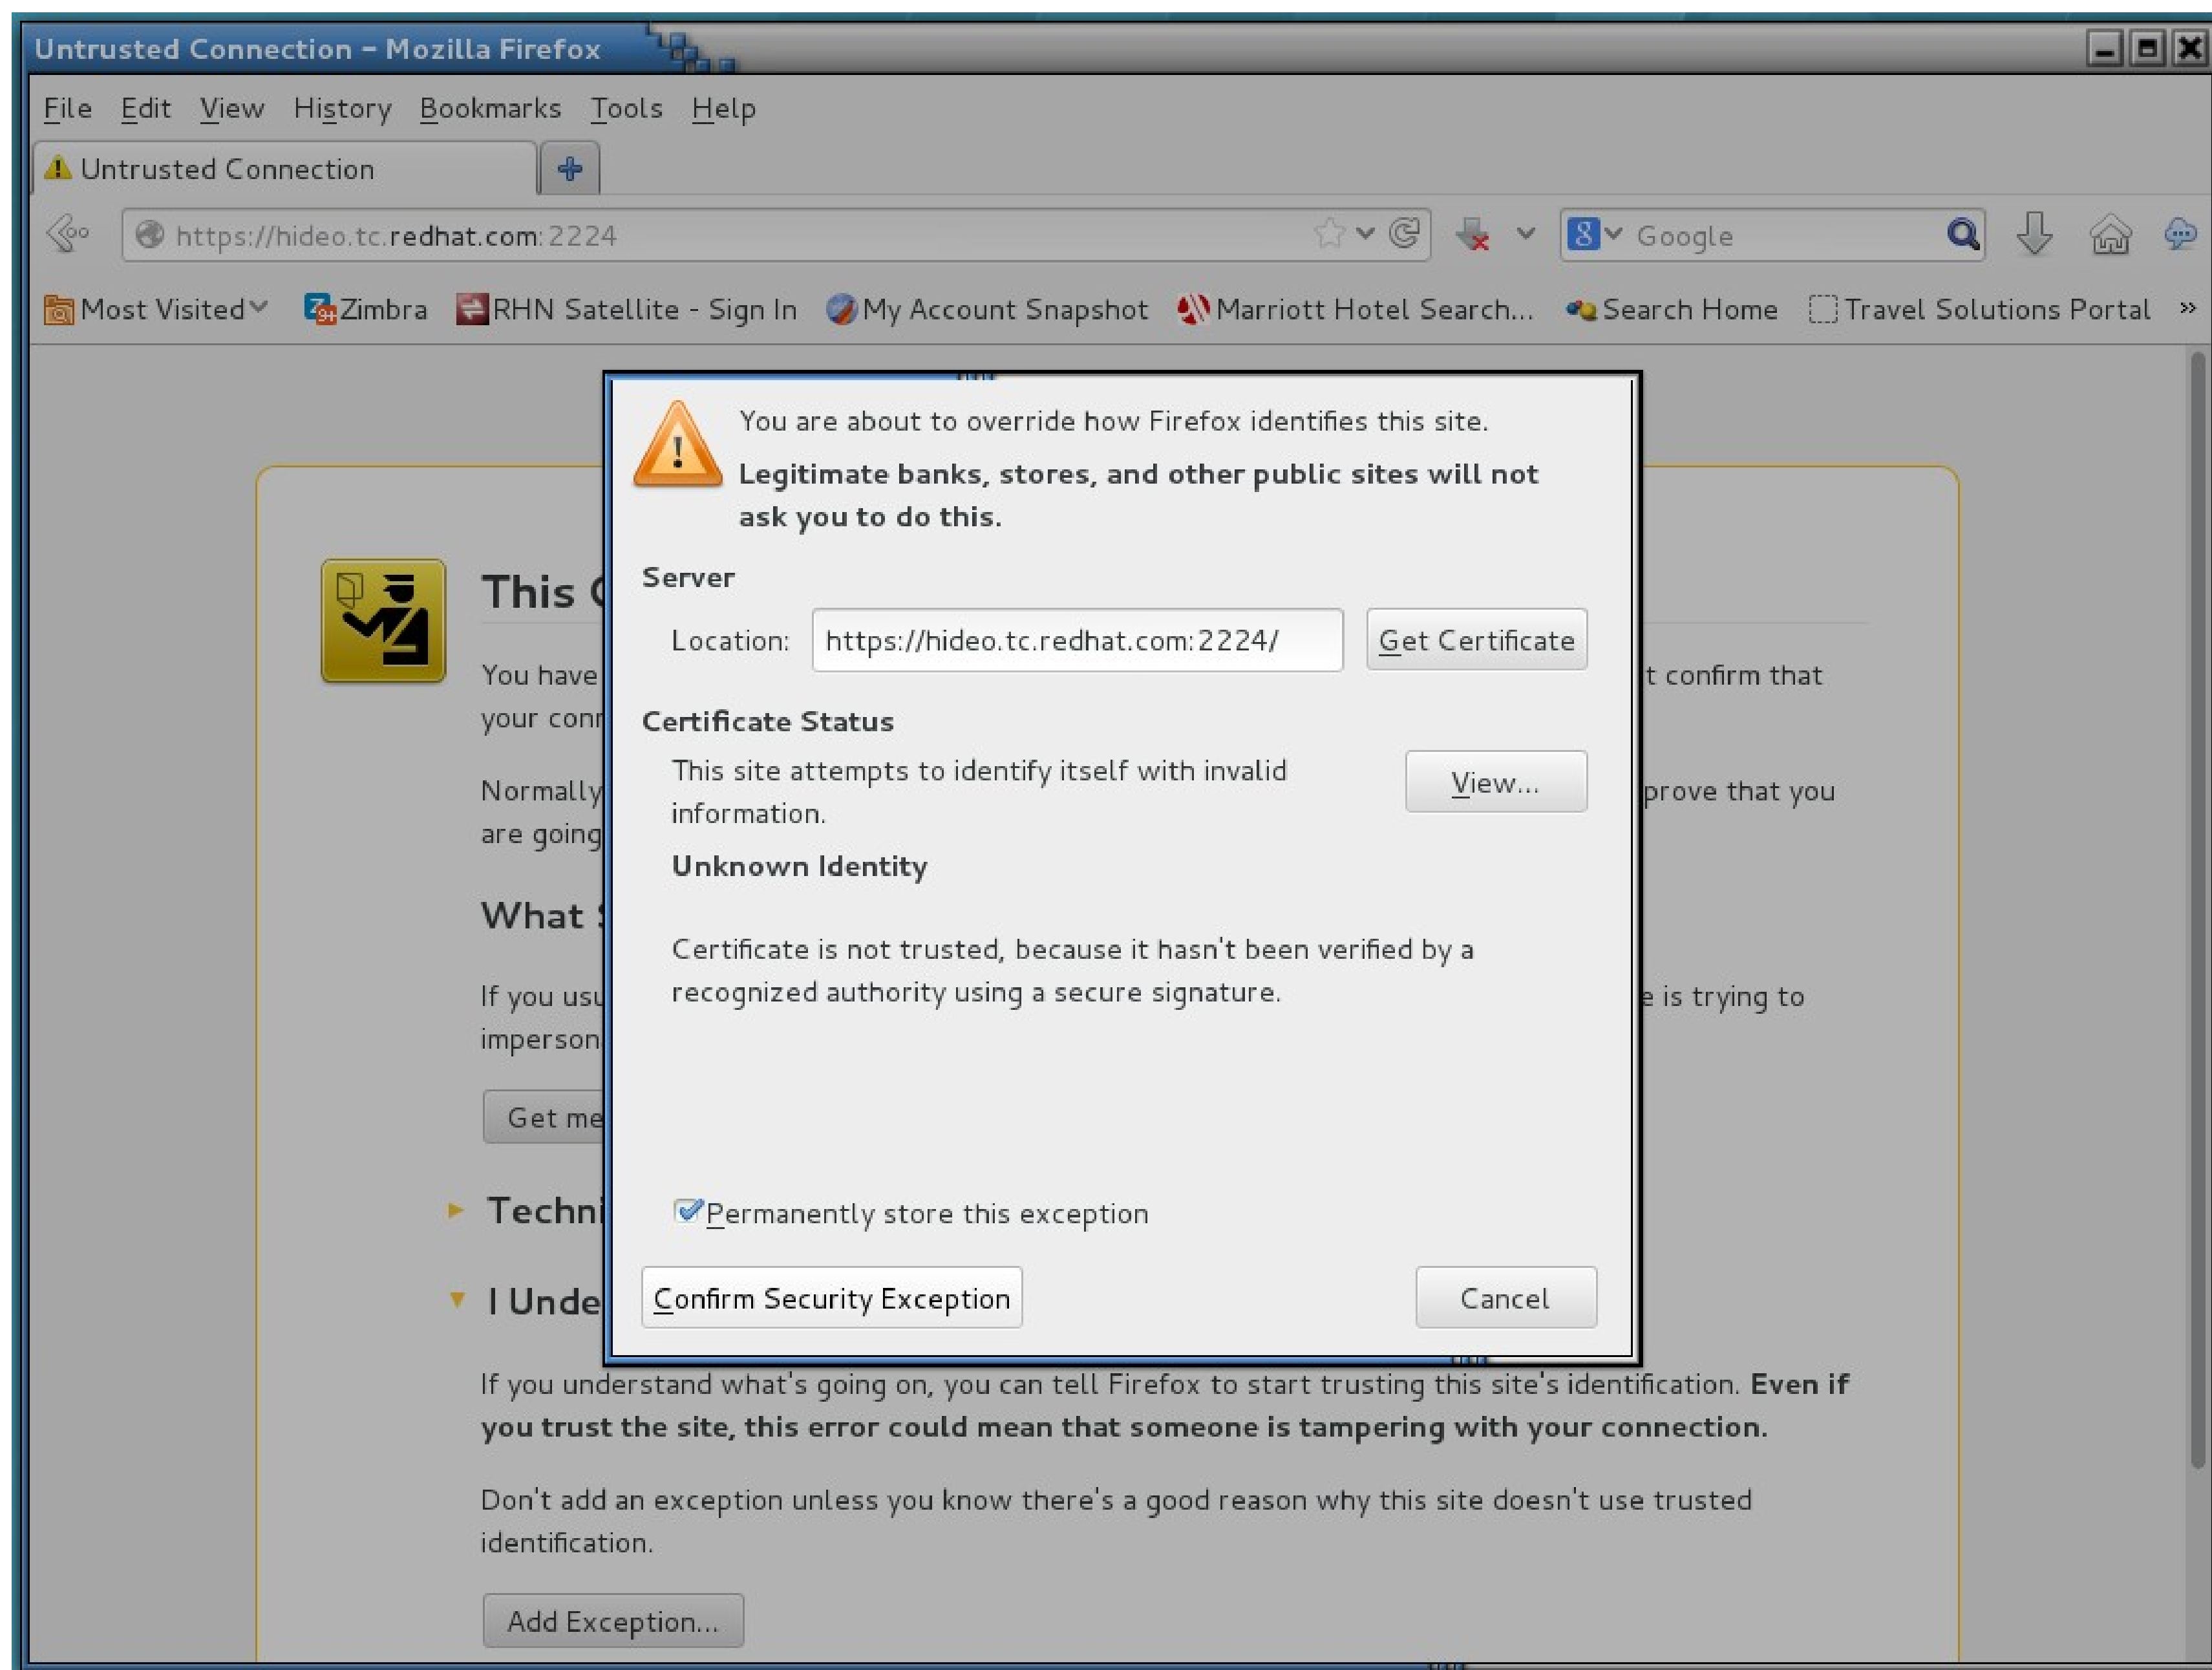


Accept the self-signed certificate

- Pretty standard process

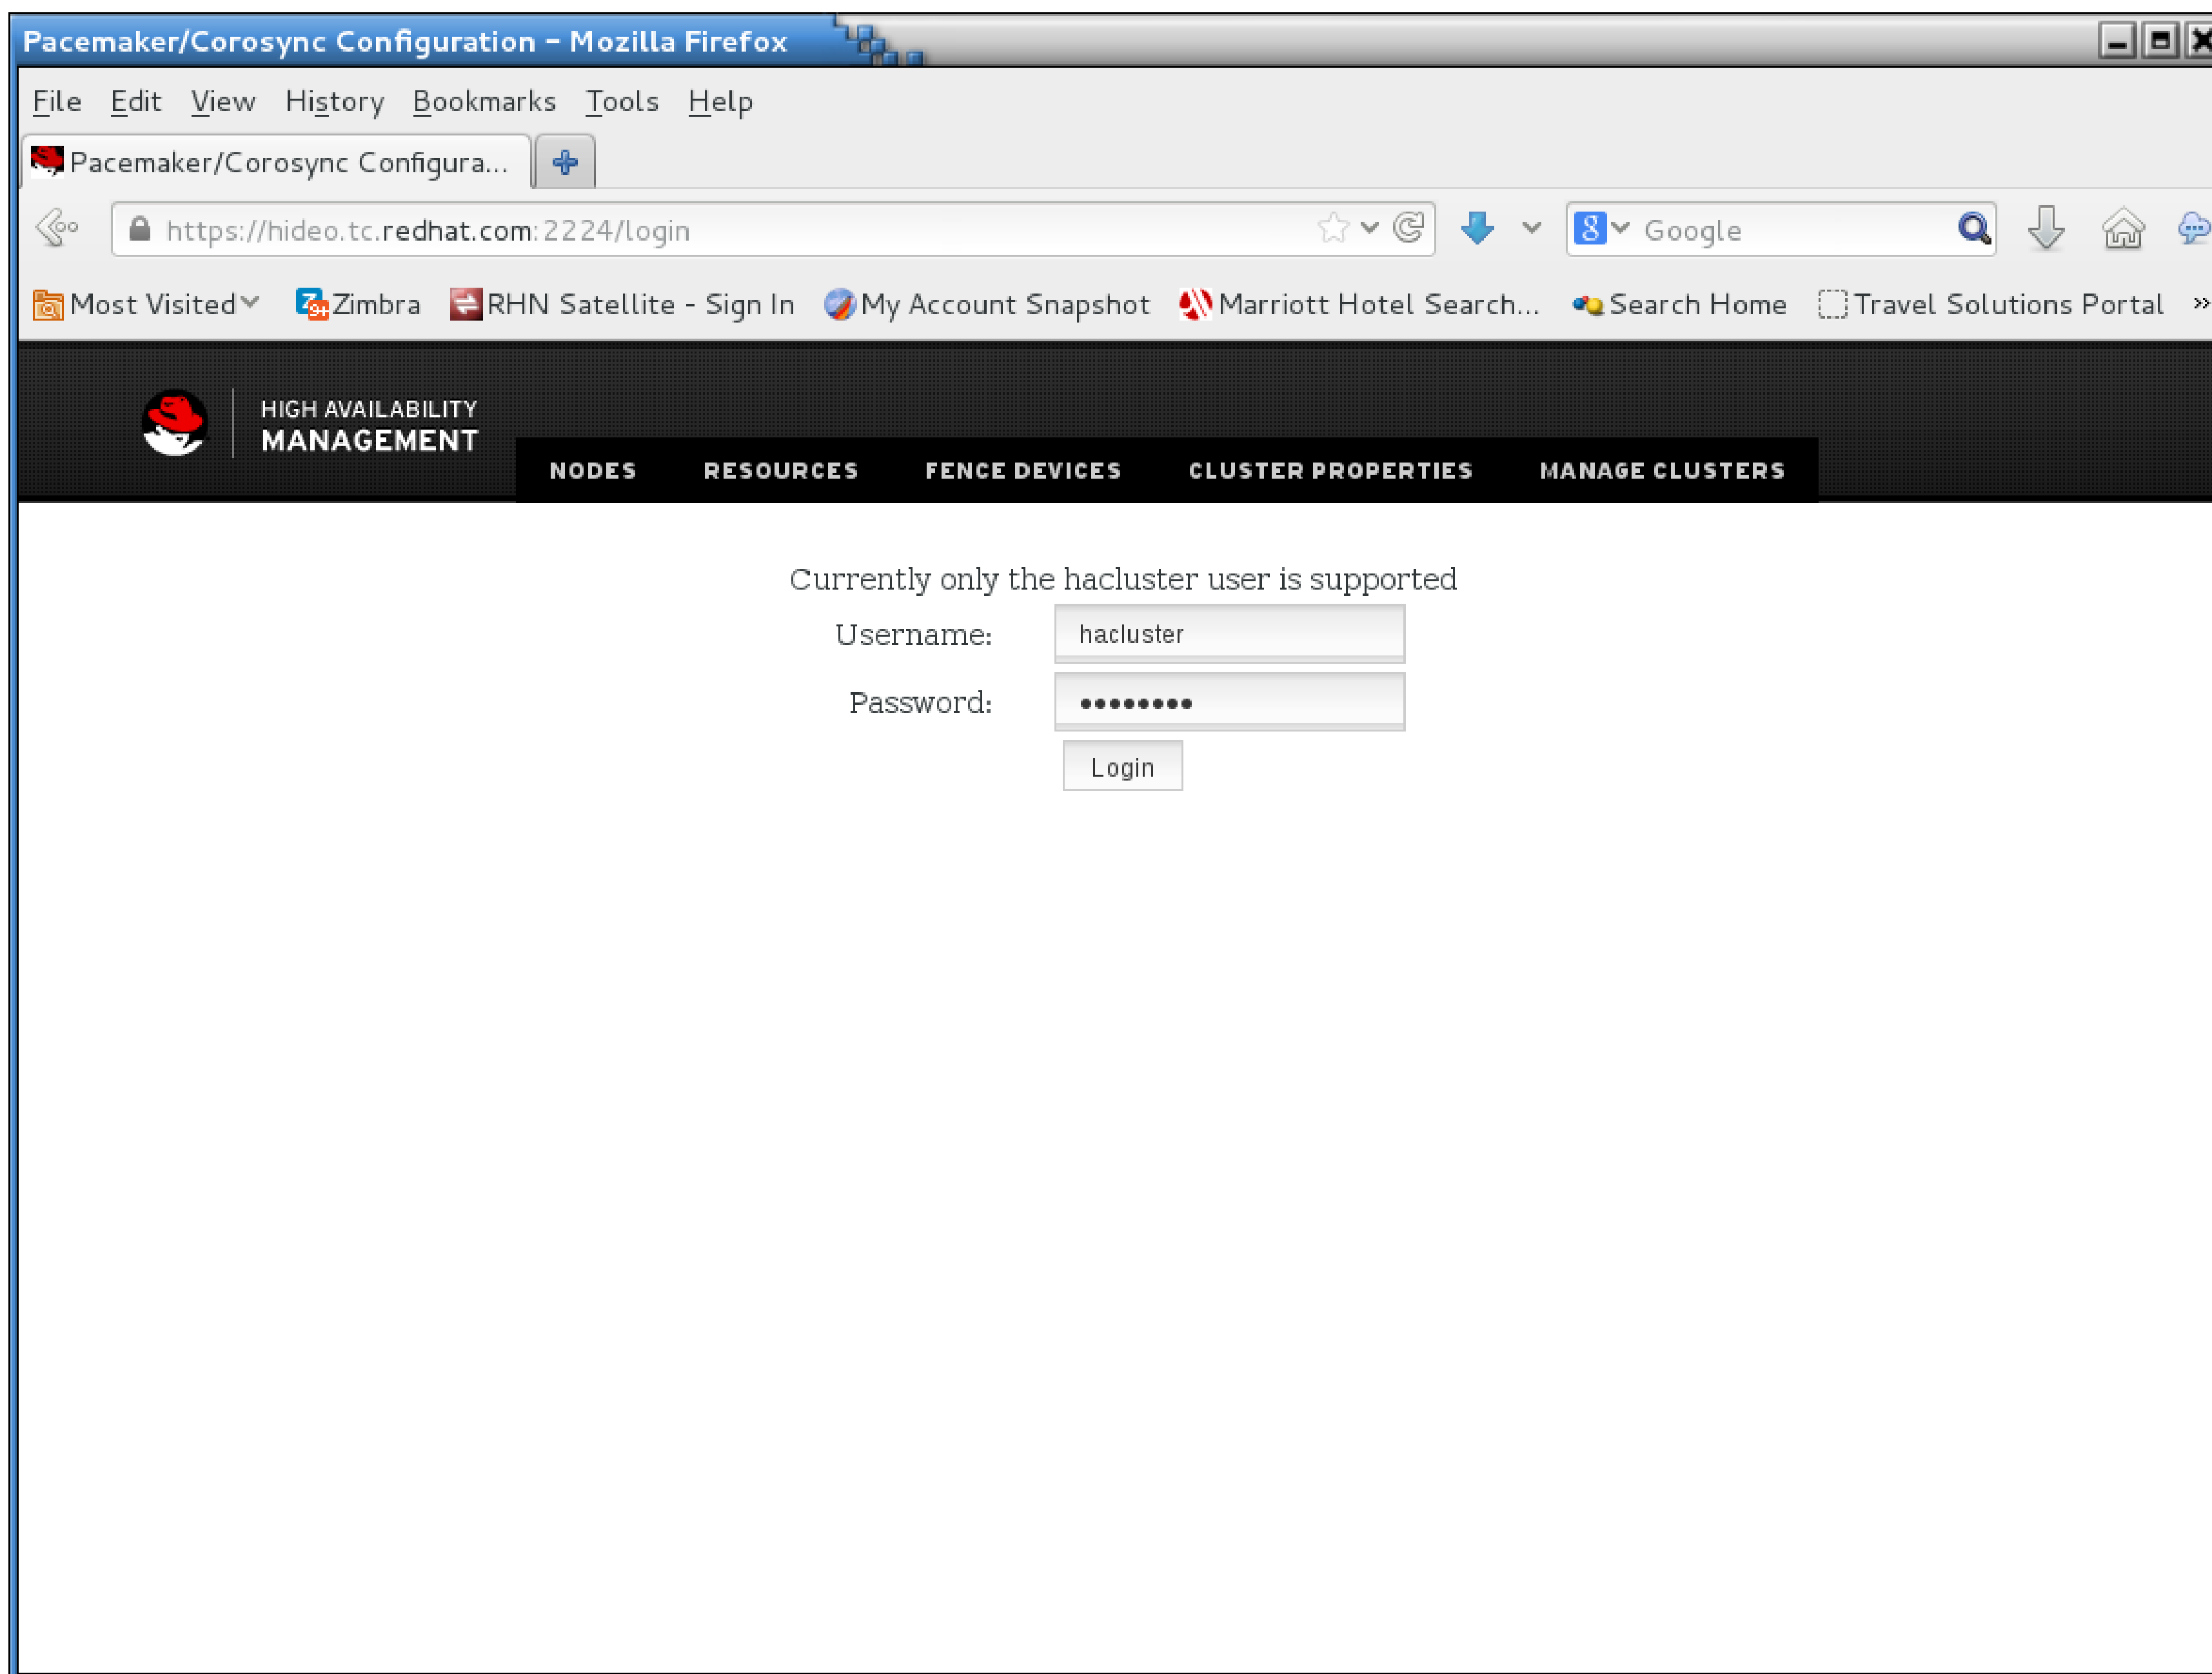






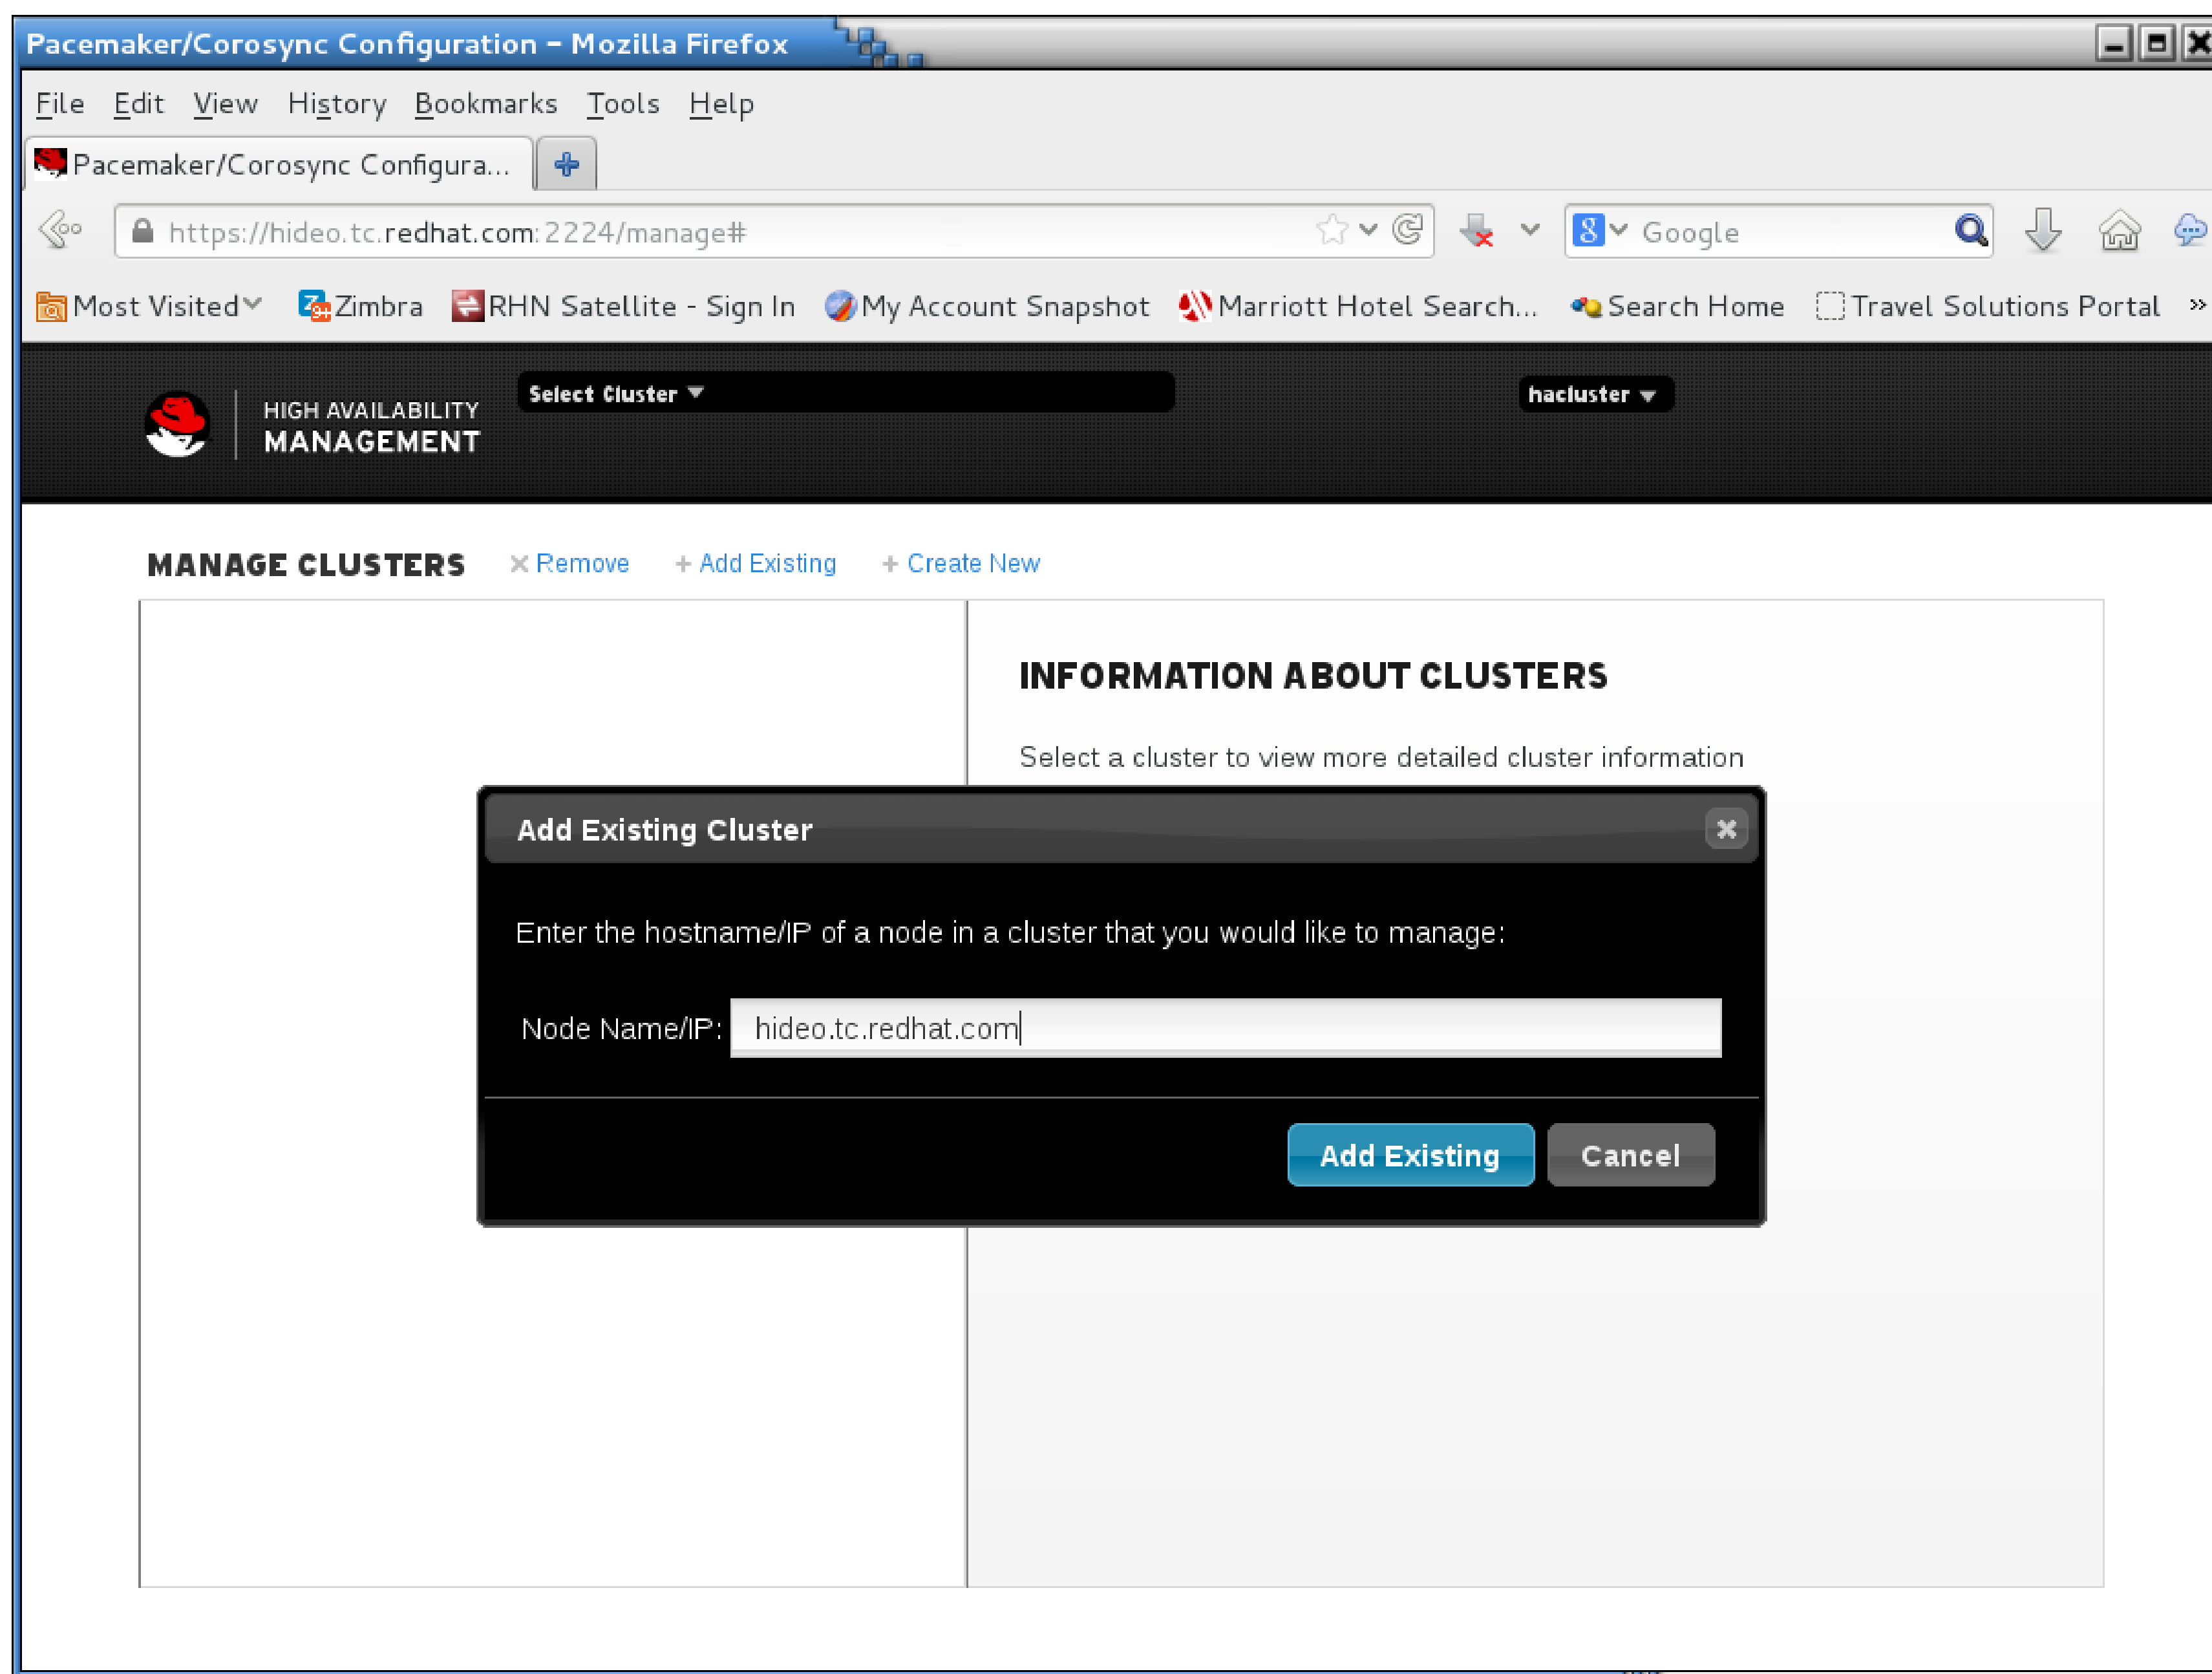
Log in as hacluster

- Use the password set earlier for hacluster



Connect to any of the nodes

- In this case, `hideo.tc.redhat.com`
- Choose “add existing”




Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/manage

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

 HIGH AVAILABILITY MANAGEMENT

Select Cluster ▼ hacluster ▼

MANAGE CLUSTERS

✕ Remove + Add Existing + Create New

	NAME	NODES
<input type="checkbox"/> ✓	summit	3

INFORMATION ABOUT CLUSTERS

Cluster: summit

Nodes: hideo.tc.redhat.com
lady3jane.tc.redhat.com
maelcum.tc.redhat.com



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

ACCESSING THE NODES

Navigate through and check each node

- They should all have pacemaker, corosync, and pcsd running

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

NODES

☐ x Remove + Add Edit Node

- ☐ ✓ hideo.tc.redhat.com
- ☐ ✓ lady3jane.tc.redhat.com
- ☐ ✓ maelcum.tc.redhat.com

Edit Node hideo.tc.redhat.com

hideo.tc.redhat.com ✓ Pacemaker Connected ✓ Corosync Connected

Start Stop Restart Standby Configure Fencing

Node ID: 1 Uptime: 0 days, 00:39:24

Cluster Daemons

NAME	STATUS
pacemaker	✓ Running (Enabled)
corosync	✓ Running (Enabled)
pcsd	✓ Running (Enabled)

Running Resources

NAME
NONE

Resource Location Preferences


Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/nodes/lady3jane.tc.redhat. Google

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

 HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

NODES

☐ x Remove + Add Edit Node

- ☐ ✓ hideo.tc.redhat.com
- ☐ ✓ lady3jane.tc.redhat.com
- ☐ ✓ maelcum.tc.redhat.com

Edit Node lady3jane.tc.redhat.com

lady3jane.tc.redhat.com ✓ Pacemaker Connected
✓ Corosync Connected

Start Stop Restart Standby Configure Fencing

Node ID: 3 Uptime: 0 days, 00:39:24

Cluster Daemons

NAME	STATUS
pacemaker	✓ Running (Enabled)
corosync	✓ Running (Enabled)
pcsd	✓ Running (Enabled)

Running Resources

NAME
NONE

Resource Location Preferences

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura...

https://hideo.tc.redhat.com:2224/managec/summit/main#/nodes/maelcum.tc.redhat.c

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

NODES

☐ x Remove + Add Edit Node

<input type="checkbox"/>	✓	hideo.tc.redhat.com
<input type="checkbox"/>	✓	lady3jane.tc.redhat.com
<input type="checkbox"/>	✓	maelcum.tc.redhat.com

Edit Node maelcum.tc.redhat.com

maelcum.tc.redhat.com

✓ Pacemaker Connected
✓ Corosync Connected

Start Stop Restart Standby Configure Fencing

Node ID: 2 Uptime: 0 days, 00:39:23

Cluster Daemons

NAME	STATUS
pacemaker	✓ Running (Enabled)
corosync	✓ Running (Enabled)
pcsd	✓ Running (Enabled)

Running Resources

NAME
NONE

Resource Location Preferences



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

ACCESSING THE CLUSTER

Inspect the Cluster Properties Tab

- This is where general options which affect cluster-wide settings


Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/configure

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

 HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

CLUSTER PROPERTIES

Batch Limit:	30 jobs
No Quorum Policy:	Stop
Symmetric:	<input checked="" type="checkbox"/>
Stonith Enabled:	<input checked="" type="checkbox"/>
Stonith Action:	Reboot
Cluster Delay:	60s
Stop Orphan Resources:	<input checked="" type="checkbox"/>
Stop Orphan Actions:	<input checked="" type="checkbox"/>
Start Failure is Fatal:	<input checked="" type="checkbox"/>
PE Error Storage:	-1
PE Warning Storage:	5000
PE Input Storage:	4000

Context-dependent Help Available

- Hover the mouse over a setting for more information

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura...

https://hideo.tc.redhat.com:2224/managec/summit/main#/configure

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

CLUSTER PROPERTIES

Batch Limit:	30	jobs
No Quorum Policy:	Stop	
Symmetric:	<input checked="" type="checkbox"/>	
Stonith Enabled:	<input checked="" type="checkbox"/>	
Stonith Action:	Reboot	
Cluster Delay:	60s	
Stop Orphan Resources:	<input checked="" type="checkbox"/>	
Stop Orphan Actions:	<input checked="" type="checkbox"/>	
Start Failure is Fatal:	<input checked="" type="checkbox"/>	
PE Error Storage:	-1	
PE Warning Storage:	5000	
PE Input Storage:	4000	

Apply Changes

What to do when the cluster does not have quorum. Allowed values:

- * ignore - continue all resource management
- * freeze - continue resource management, but don't recover resources from nodes not in the affected partition
- * stop - stop all resources in the affected cluster partition
- * suicide - fence all nodes in the affected cluster partition

RED HAT
SUMMIT

10 YEARS *and counting*

SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE FENCING

What is Fencing?

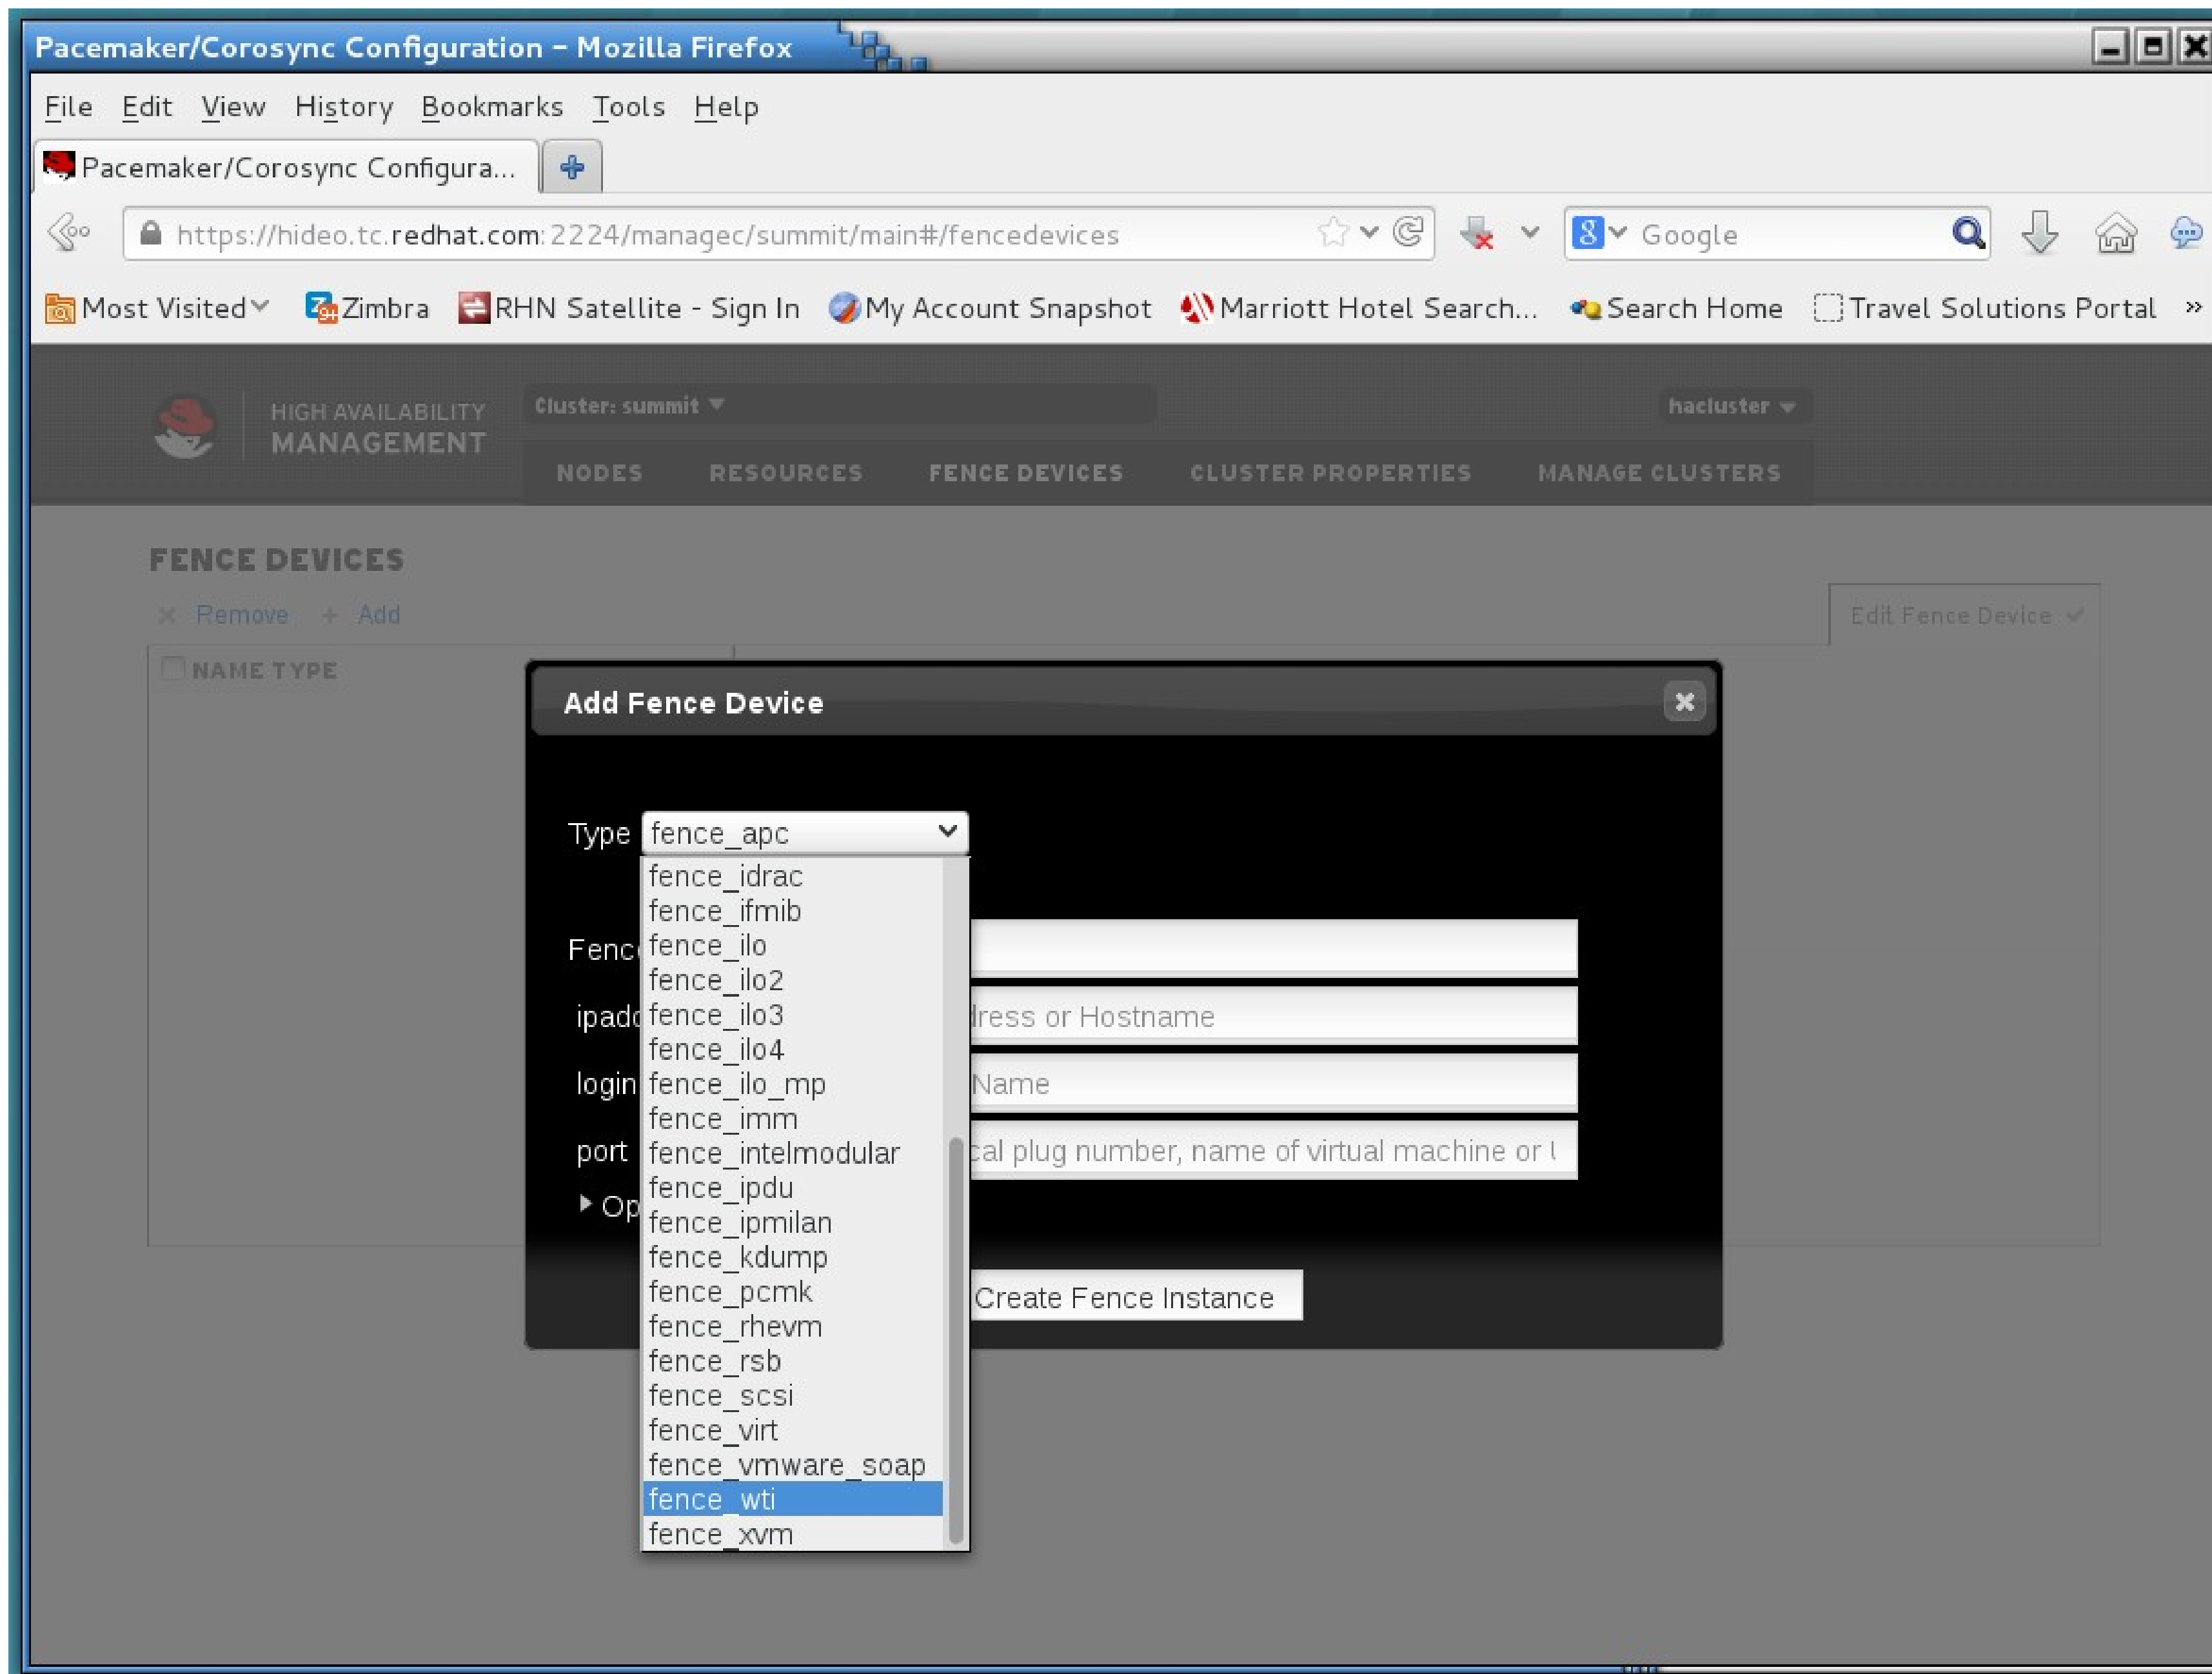
- If a node stops responding, the cluster will attempt to remove that node from the cluster.
- This is referred to as STONITH (**S**hoot **T**he **O**ther **N**ode **I**n **T**he **H**ead)
- You don't want multiple machines to e.g. write to the same datastore without doing distributed lock management, so it's best to just take the unresponsive node out of the equation.

What is Fencing?

- There are numerous ways to do this.
 - Managed power devices from WTI, APC and others
 - Fiber channel fencing
 - IPMI
 - Out of band management devices (DRAC, iLO, RSA, etc.)
 - SCSI reservations
 - Virtual Machine fencing

Configure Fence Devices

- Choose +Add
- Choose the fence device (in this case, my lab has a manageable WTI ips-800-d20 power switch - they are typically a couple of hundred bucks on eBay, they work amazingly well, and they are very rugged).



Configure Fence Devices

- Assign a friendly name (in this case, “summit-wti”)
- Enter the IP address or hostname (in this case, “wti.tc.redhat.com”)


Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/fencedevices/summit-wti ☆ Google

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

 HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

FENCE DEVICES

× Remove + Add Edit Fence Device ✓

<input type="checkbox"/>	NAME	TYPE
<input checked="" type="checkbox"/>	summit-wti	stonith:fence_wti ▶

Edit Fence Device summit-wti

summit-wti ✓ Inactive

× Remove ↻ Cleanup

Current Location:

Type: stonith:fence_wti

ipaddr wti.tc.redhat.com

port Physical plug number, name of virtual machine

▶ Optional Arguments:

Apply Changes

Configure Fence Devices

- Go into “Optional Arguments” to set the following extra settings:
 - “pcmk_host_map” in the format [hostname]:[power_port],[hostname]:[power_port],[hostname]:[power_port]. So for my lab:
 - lady3jane.tc.redhat.com:5;hideo.tc.redhat.com:4;maelcum.tc.redhat.com:6

Configure Fence Devices

- Go into “Optional Arguments” to set the following extra settings:
 - For the WTI, you don't need to define a user, only “passwd”

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/fenceddevices/summit-wti ☆ ↺ ↻ ↵ Google

Most Visited ▾ Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

Type: stonith:fence_wti

ipaddr	wti.tc.redhat.com
port	Physical plug number, name of virtual machine
▼ Optional Arguments:	
pcmk_host_list	
pcmk_host_map	ideo.tc.redhat.com:4;maelcum.tc.redhat.com:6
pcmk_host_check	
login	Login Name
passwd	red22hat
cmd_prompt	Force Python regex for command prompt
secure	SSH connection
ipport	TCP/UDP port to use for connection with device
inet4_only	Forces agent to use IPv4 addresses only
inet6_only	Forces agent to use IPv6 addresses only
passwd_script	Script to retrieve password
identity_file	Identity file for ssh
ssh_options	SSH options to use
verbose	Verbose mode
debug	Write debug information to given file
version	Display version information and exit

Configure Fence Devices

- Go into “Optional Arguments” to set the following extra settings:
 - I also prefer to set `power_wait` to 5 seconds so the system being power cycled doesn't get damaged.
 - I set “delay” to about 5 seconds, as well. We don't want to fence systems instantly when it's only temporarily unavailable.

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/fencedevices/summit-wti ☆ Google

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

secure	SSH connection
ipport	TCP/UDP port to use for connection with device
inet4_only	Forces agent to use IPv4 addresses only
inet6_only	Forces agent to use IPv6 addresses only
passwd_script	Script to retrieve password
identity_file	Identity file for ssh
ssh_options	SSH options to use
verbose	Verbose mode
debug	Write debug information to given file
version	Display version information and exit
help	Display help and exit
separator	Separator for CSV created by operation list
power_timeout	Test X seconds for status change after ON/OFF
shell_timeout	Wait X seconds for cmd prompt after issuing c
login_timeout	Wait X seconds for cmd prompt after login
power_wait	5
delay	5
retry_on	Count of attempts to retry power on

Apply Changes

Configure Fence Devices

- After a few seconds, the fence device should turn green, indicating all of the systems are aware of it and it's active.


Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/fencedevices/summit-wti ☆ Google

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

 HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

FENCE DEVICES

× Remove + Add Edit Fence Device ✓

<input type="checkbox"/>	NAME	TYPE
<input checked="" type="checkbox"/>	summit-wti	stonith:fence_wti

Edit Fence Device summit-wti

summit-wti ✓ Running

× Remove ↻ Cleanup

Current Location: hideo.tc.redhat.com

Type: stonith:fence_wti

ipaddr wti.tc.redhat.com

port Physical plug number, name of virtual machine

► **Optional Arguments:**

Apply Changes

Test Fence Devices

- You can test the fence device and host mapping using stonith_admin

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # target - do not install pacemaker, corosync, etc  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# stonith_admin --reboot hideo.tc.redhat.com  
[root@maelcum ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
Mar 22 12:38:29 lady3jane.tc.redhat.com corosync[4819]: [TOTEM ] A processor failed, forming new configuration.  
Mar 22 12:38:30 lady3jane.tc.redhat.com corosync[4819]: [TOTEM ] A new membership (172.31.100.10:16) was formed. Members left: 1  
Mar 22 12:38:30 lady3jane.tc.redhat.com corosync[4819]: [QUORUM] Members[2]: 3 2  
Mar 22 12:38:30 lady3jane.tc.redhat.com corosync[4819]: [MAIN ] Completed service synchronization, ready to provide service.  
Mar 22 12:38:30 lady3jane.tc.redhat.com crmd[4840]: notice: crm_update_peer_state: pcmk_quorum_notification: Node hideo.tc.redhat.com[1] - state is now lost (was member)  
Mar 22 12:38:30 lady3jane.tc.redhat.com pacemakerd[4834]: notice: crm_update_peer_state: pcmk_quorum_notification: Node hideo.tc.redhat.com[1] - state is now lost (was member)  
Mar 22 12:38:39 lady3jane.tc.redhat.com stonith-ng[4836]: notice: log_operation: Operation 'reboot' [6188] (call 2 from stonith_admin.6171) for host 'hideo.tc.redhat.com' with device 'summit-wti' returned: 0 (OK)  
Mar 22 12:38:39 lady3jane.tc.redhat.com stonith-ng[4836]: notice: remote_op_done: Operation reboot of hideo.tc.redhat.com by maelcum.tc.redhat.com for stonith_admin.6171@maelcum.tc.redhat.com.0accfd40: OK  
Mar 22 12:38:39 lady3jane.tc.redhat.com crmd[4840]: notice: tengine_stonith_notify: Peer hideo.tc.redhat.com was terminated (reboot) by maelcum.tc.redhat.com for maelcum.tc.redhat.com: OK (ref=0accfd40-b41c-43d7-a195-0c519b7e2045) by client stonith_admin.6171  
Mar 22 12:38:39 lady3jane.tc.redhat.com stonith-ng[4836]: notice: stonith_device
```

root@lady3jane:~

File Edit View Search Terminal Help

fy: Peer hideo.tc.redhat.com was terminated (reboot) by maelcum.tc.redhat.com for maelcum.tc.redhat.com: OK (ref=0accfd40-b41c-43d7-a195-0c519b7e2045) by client stonith_admin.6171

Mar 22 12:38:39 lady3jane.tc.redhat.com stonith-ng[4836]: notice: stonith_device_register: Device 'summit-wti' already existed in device list (1 active devices)

Mar 22 12:38:41 lady3jane.tc.redhat.com crmd[4840]: notice: process_lrm_event: LRM operation summit-wti_start_0 (call=23, rc=0, cib-update=26, confirmed=true) ok

Mar 22 12:38:42 lady3jane.tc.redhat.com crmd[4840]: notice: process_lrm_event: LRM operation summit-wti_monitor_60000 (call=24, rc=0, cib-update=27, confirmed=false) ok

Mar 22 12:39:21 lady3jane.tc.redhat.com corosync[4819]: [TOTEM] A new membership (172.31.100.10:20) was formed. Members joined: 1

Mar 22 12:39:21 lady3jane.tc.redhat.com corosync[4819]: [QUORUM] Members[3]: 3 2 1

Mar 22 12:39:21 lady3jane.tc.redhat.com corosync[4819]: [MAIN] Completed service synchronization, ready to provide service.

Mar 22 12:39:21 lady3jane.tc.redhat.com crmd[4840]: notice: crm_update_peer_state: pcmk_quorum_notification: Node hideo.tc.redhat.com[1] - state is now member (was lost)

Mar 22 12:39:21 lady3jane.tc.redhat.com pacemakerd[4834]: notice: crm_update_peer_state: pcmk_quorum_notification: Node hideo.tc.redhat.com[1] - state is now member (was lost)

Mar 22 12:40:01 lady3jane.tc.redhat.com systemd[1]: Starting Session 136 of user

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE RESOURCES

Create Resources

- We'll start out very simple: with an Apache web server instance which does not use shared storage. We'll add in shared storage later.
- There will be two components we need to define:
 - The floating IP address
 - The apache

Configure the floating IP address

- Choose Add
- Choose Open Cluster Framework (OCF) heartbeat Class/Provider
- Choose IPAddr2 (this is new, and Linux specific, don't use the old IPAddr) type
- Give it a Resource ID (friendly name)
- Assign the IP address
- Note that you get context-sensitive hover help!

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

RESOURCES

Remove Add Create Group

NAME TYPE

Add Resource

Class/Provider ocf:heartbeat

Type IPAddr2

Description: Manages virtual IPv4 and IPv6 addresses (Linux specific version) i

Resource Group: None

Clone: ☐

Master/Slave: ☐

Disabled: ☐

Resource ID summit-ip

ip IPv4 or IPv6 address

Optional Arguments: The IPv4 (dotted quad notation) or IPv6 address (colon hexadecimal notation)
example IPv4 "192.168.1.1".
example IPv6 "2001:db8:DC28:0:0:FC57:D4C8:1FFF".


Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

 HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

RESOURCES

[Remove](#) [Add](#) [Create Group](#) [Edit Resource](#)

<input type="checkbox"/>	NAME	TYPE
<input checked="" type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-ip

summit-ip ✓ Running

[✓ Enable](#) [✗ Disable](#) [↻ Cleanup](#) [✗ Remove](#)

Current Location: hideo.tc.redhat.com

Resource Location Preferences (0)

Node	Score	Remove
NONE		
<input type="text"/>	<input type="text"/>	<input type="button" value="Add"/>

Resource Ordering Preferences (0)

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# ip addr show  
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN  
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00  
    inet 127.0.0.1/8 scope host lo  
        valid_lft forever preferred_lft forever  
    inet6 ::1/128 scope host  
        valid_lft forever preferred_lft forever  
2: p14p1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP q  
len 1000  
    link/ether 90:2b:34:ad:2c:fe brd ff:ff:ff:ff:ff:ff  
    inet 172.31.100.22/24 brd 172.31.100.255 scope global dynamic p14p1  
        valid_lft 17989sec preferred_lft 17989sec  
    inet 172.31.100.250/24 brd 172.31.100.255 scope global secondary p14p1  
        valid_lft forever preferred_lft forever  
    inet6 fe80::922b:34ff:fead:2cfe/64 scope link  
        valid_lft forever preferred_lft forever  
[root@hideo ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# pcs resource show  
summit-ip (ocf::heartbeat:IPaddr2): Started  
[root@hideo ~]#
```

Install httpd and wget on all the nodes

- `yum -y install httpd wget`


```
root@neuromancer:~
File Edit View Search Terminal Help
[root@neuromancer ~]# # iSCSI target - don't install anything else here
[root@neuromancer ~]#

root@hideo:~
File Edit View Search Terminal Help
Transaction Summary
=====
Install 1 Package (+4 Dependent packages)

Total download size: 1.4 M
Installed size: 4.3 M
Downloading packages:
(1/5): apr-1.4.8-3.el7.x86_64.rpm | 103 kB 00:00
(2/5): apr-util-1.5.2-6.el7.x86_64.rpm | 92 kB 00:00
(3/5): httpd-tools-2.4.6-17.el7.x86_64.rpm | 76 kB 00:00
(4/5): mailcap-2.1.41-2.el7.noarch.rpm | 31 kB 00:00
(5/5): httpd-2.4.6-17.el7.x86_64.rpm | 1.2 MB 00:00
-----
Total 4.7 MB/s | 1.4 MB 00:00
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
  Installing : apr-1.4.8-3.el7.x86_64 1/5
  Installing : apr-util-1.5.2-6.el7.x86_64 2/5
  Installing : httpd-tools-2.4.6-17.el7.x86_64 3/5
  Installing : mailcap-2.1.41-2.el7.noarch 4/5
  Installing : httpd-2.4.6-17.el7.x86_64 5/5

root@lady3jane:~
File Edit View Search Terminal Help
[root@lady3jane ~]# yum -y install httpd wget

root@maelcum:~
File Edit View Search Terminal Help
Transaction Summary
=====
Install 1 Package (+4 Dependent packages)

Total download size: 1.4 M
Installed size: 4.3 M
Downloading packages:
(1/5): apr-1.4.8-3.el7.x86_64.rpm | 103 kB 00:00
(2/5): apr-util-1.5.2-6.el7.x86_64.rpm | 92 kB 00:00
(3/5): httpd-tools-2.4.6-17.el7.x86_64.rpm | 76 kB 00:00
(4/5): httpd-2.4.6-17.el7.x86_64.rpm | 1.2 MB 00:00
(5/5): mailcap-2.1.41-2.el7.noarch.rpm | 31 kB 00:00
-----
Total 5.7 MB/s | 1.4 MB 00:00
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
  Installing : apr-1.4.8-3.el7.x86_64 1/5
  Installing : apr-util-1.5.2-6.el7.x86_64 2/5
  Installing : httpd-tools-2.4.6-17.el7.x86_64 3/5
  Installing : mailcap-2.1.41-2.el7.noarch 4/5
  Installing : httpd-2.4.6-17.el7.x86_64 [#####] 5/5
```

Install httpd and wget on all the nodes

- Confirm that httpd is disabled - we want it started by the cluster software, not at boot time!

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# systemctl status httpd  
httpd.service - The Apache HTTP Server  
    Loaded: loaded (/usr/lib/systemd/system/httpd.service; disabled)  
    Active: inactive (dead)  
  
[root@hideo ~]#
```

Configure the Apache service

- For testing, echo the hostname into index.html on each node:
 - `hostname > /var/www/html/index.html`

<pre>root@neuromancer:~ File Edit View Search Terminal Help [root@neuromancer ~]# # iSCSI target - don't install anything else here [root@neuromancer ~]#</pre>	<pre>root@hideo:~ File Edit View Search Terminal Help [root@hideo ~]# hostname > /var/www/html/index.html [root@hideo ~]# cat /var/www/html/index.html hideo.tc.redhat.com [root@hideo ~]#</pre>
<pre>root@lady3jane:~ File Edit View Search Terminal Help [root@lady3jane ~]# hostname > /var/www/html/index.html [root@lady3jane ~]#</pre>	<pre>root@maelcum:~ File Edit View Search Terminal Help [root@maelcum ~]# hostname > /var/www/html/index.html [root@maelcum ~]#</pre>

Enable Apache monitoring

```
cat > /etc/httpd/conf.d/status.conf << EOF
```

```
<Location /server-status>
```

```
    SetHandler server-status
```

```
    Order deny,allow
```

```
    Deny from all
```

```
    Allow from 127.0.0.1
```

```
</Location>
```

```
EOF
```

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# cat > /etc/httpd/conf.d/status.conf << EOF
```

```
> <Location /server-status>
```

```
>     SetHandler server-status
```

```
>     Order deny,allow
```

```
>     Deny from all
```

```
>     Allow from 127.0.0.1
```

```
> </Location>
```

```
> EOF
```

```
[root@hideo ~]# cat /etc/httpd/conf.d/status.conf
```

```
<Location /server-status>
```

```
    SetHandler server-status
```

```
    Order deny,allow
```

```
    Deny from all
```

```
    Allow from 127.0.0.1
```

```
</Location>
```

```
[root@hideo ~]#
```


Configure the Apache service

- Choose Add
- Choose Open Cluster Framework (OCF) heartbeat Class/Provider
- Choose the apache type
- Give it a Resource ID (friendly name)

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-apache

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

RESOURCES

Remove Add Create Group

	NAME	TYPE
<input type="checkbox"/>	summit-ip	ocf::heartbeat

Add Resource

Class/Provider: ocf:heartbeat

Type: apache

Description: Manages an Apache Web server instance ⓘ

Resource Group: None

Clone: ☐

Master/Slave: ☐

Disabled: ☐

Resource ID: summit-apache

Optional Arguments:

Create Resource

Resource Ordering Preferences (0)

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-apache

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

RESOURCES

Remove Add Create Group Edit Resource

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-apache

summit-apache Running

Enable Disable Cleanup Remove

Current Location: maelcum.tc.redhat.com

Resource Location Preferences (0)

Node	Score	Remove
NONE		
		Add

Resource Ordering Preferences (0)

root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started hideo.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started maelcum.tc.redhat.com

PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online


Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@hideo ~]#

Wait, what?

- You see anything wrong with this picture?



root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started <u>hideo.tc.redhat.com</u>
summit-apache	(ocf::heartbeat:apache):	Started <u>maelcum.tc.redhat.com</u>

PCSD Status:

hideo.tc.redhat.com: Online

lady3jane.tc.redhat.com: Online

maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled

pacemaker: active/enabled

pcsd: active/enabled

[root@hideo ~]#

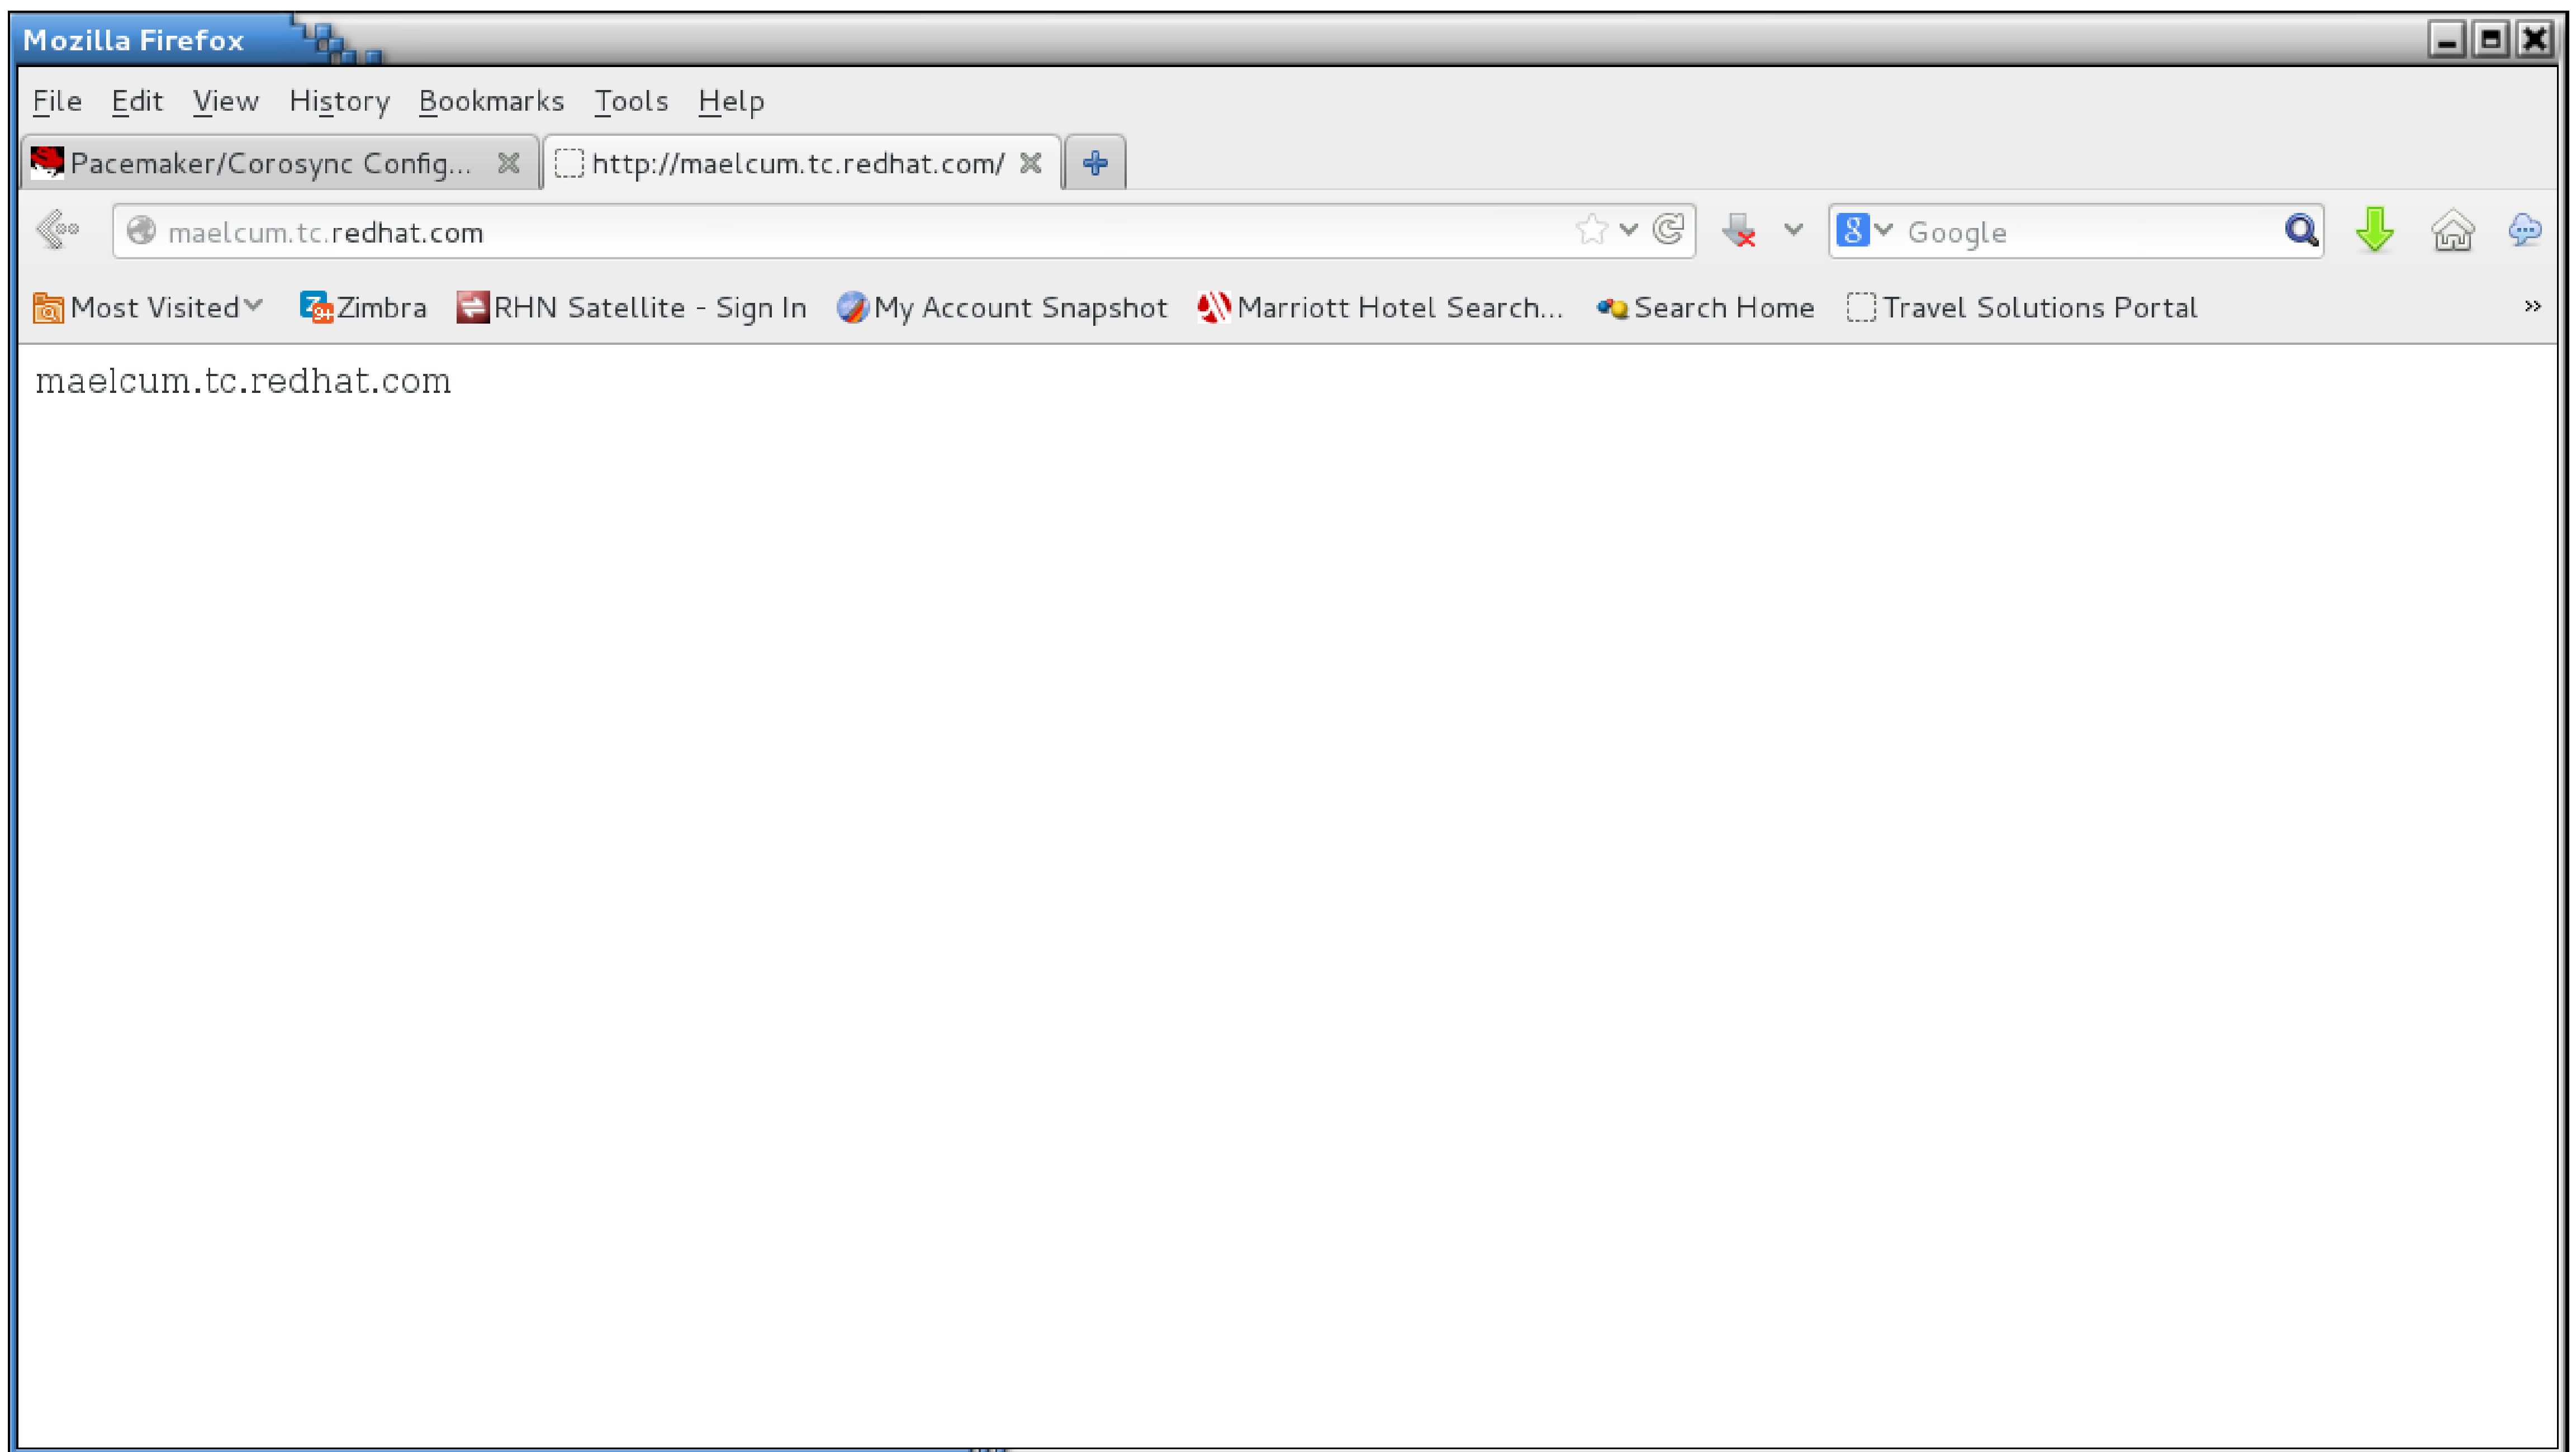


Unable to connect

Firefox can't establish a connection to the server at 172.31.100.250.

- The site could be temporarily unavailable or too busy. Try again in a few moments.
- If you are unable to load any pages, check your computer's network connection.
- If your computer or network is protected by a firewall or proxy, make sure that Firefox is permitted to access the Web.

Try Again



We need to set two resource features

- Resource Ordering Preferences
- Resource Colocation Preferences

Resource Ordering Preferences

- Choose the summit-ip resource
- Go to Resource Ordering Preferences
- Add in the resource summit-apache
- Set summit-apache to start after summit-ip
- Click add

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource ✓

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input checked="" type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-ip

summit-ip ✓ Running

✓ Enable ✕ Disable ↻ Cleanup ✕ Remove

Current Location: hideo.tc.redhat.com

Resource Location Preferences (0)

Resource Ordering Preferences (0)

Resource	Before/After	Score	Remove
NONE			
summit-apache	After		Add

Resource Colocation Preferences (0)

Resource	Together/Apart	Score	Remove
NONE			

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura...

+

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

☆ ↺

📄 ↵

🔍 Google

🔍

⬇️

🏠

💬

Most Visited ▾

📧 Zimbra

🔑 RHN Satellite - Sign In

👤 My Account Snapshot

🏨 Marriott Hotel Search...

🌐 Search Home

🗂️ Travel Solutions Portal

»

✕ Remove

+

Add

+

Create Group

Edit Resource ▾

<input type="checkbox"/>	NAME	TYPE
<input type="checkbox"/>	✓ summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	✓ summit-ip	ocf::heartbeat:IPaddr2 ▶

summit-ip

✓ Running

✓ Enable

✕ Disable

🔄 Cleanup

✕ Remove

Current Location:

hideo.tc.redhat.com

Resource Location Preferences (0)

▶

Resource Ordering Preferences (1)

▼

Resource	Before/After	Score	Remove
summit-apache	Starts after summit-ip		✕
<input type="text"/>	After ▾	<input type="text"/>	Add

Resource Colocation Preferences (0)

▼

Resource	Together/Apart	Score	Remove
NONE			

Resource Colocation Preferences

- Choose the summit-apache resource
- Choose Resource Colocation Preferences
- Enter summit-ip and set it to start together with summit-apache
- Click add

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-apache

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource ✓

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-apache

summit-apache ✓ Running

✓ Enable ✕ Disable ↻ Cleanup ✕ Remove

Current Location: maelcum.tc.redhat.com

Resource Location Preferences (0)

Resource Ordering Preferences (1)

Resource Colocation Preferences (0)

Resource	Together/Apart	Score	Remove
NONE			
summit-ip	Together		Add

Resource Meta Attributes (0)

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-apache

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-apache

summit-apache Running

Enable Disable Cleanup Remove

Current Location: maelcum.tc.redhat.com

Resource Location Preferences (0)

Resource Ordering Preferences (1)

Resource Colocation Preferences (1)

Resource	Together/Apart	Score	Remove
summit-ip	Together	INFINITY	X
	Together		Add

Resource Meta Attributes (0)

root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started hideo.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started hideo.tc.redhat.com


PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@hideo ~]#



Apache Monitoring

- Go into the summit-apache resource and choose Optional Arguments
- Because we set up monitoring, use the URL we defined - `http://localhost/server-status`

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-apache

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

configfile	configuration file path
httpd	httpd binary path
port	httpd port
statusurl	http://localhost/server-status
testregex	monitor regular expression
client	http client
testurl	test url
testregex10	extended monitor regular expression
testconffile	test configuration file
testname	test name
options	command line options
envfiles	environment settings files
use_ipv6	use ipv6 with http clients

Apply Changes

Test Apache

- You can kill the httpd process to test that it gets restarted

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# ps ax | grep httpd  
9039 ?          Ss      0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c  
PidFile /var/run/httpd.pid  
9043 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c  
PidFile /var/run/httpd.pid  
9045 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c  
PidFile /var/run/httpd.pid  
9047 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c  
PidFile /var/run/httpd.pid  
9049 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c  
PidFile /var/run/httpd.pid  
9052 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c  
PidFile /var/run/httpd.pid  
9315 pts/0      S+      0:00 grep --color=auto httpd  
[root@hideo ~]# pkill -9 httpd  
[root@hideo ~]# ps ax | grep httpd  
9514 pts/0      S+      0:00 grep --color=auto httpd  
[root@hideo ~]#
```


Mar 24 11:29:23 hideo.tc.redhat.com apache(summit-apache)[11054]: INFO: apache not running
Mar 24 11:29:23 hideo.tc.redhat.com crmd[2106]: notice: process_lrm_event: LRM operation summit-apache_monitor_10000 (call=68, rc=7, cib-update=54, confirmed=false) not running
Mar 24 11:29:23 hideo.tc.redhat.com attrd[2103]: notice: attrd_cs_dispatch: Update relayed from maelcum.tc.redhat.com
Mar 24 11:29:23 hideo.tc.redhat.com attrd[2103]: notice: attrd_trigger_update: Sending flush op to all hosts for: fail-count-summit-apache (3)
Mar 24 11:29:23 hideo.tc.redhat.com attrd[2103]: notice: attrd_perform_update: Sent update 65: fail-count-summit-apache=3
Mar 24 11:29:23 hideo.tc.redhat.com attrd[2103]: notice: attrd_cs_dispatch: Update relayed from maelcum.tc.redhat.com
Mar 24 11:29:23 hideo.tc.redhat.com attrd[2103]: notice: attrd_trigger_update: Sending flush op to all hosts for: last-failure-summit-apache (1395678563)
Mar 24 11:29:23 hideo.tc.redhat.com attrd[2103]: notice: attrd_perform_update: Sent update 67: last-failure-summit-apache=1395678563
Mar 24 11:29:23 hideo.tc.redhat.com apache(summit-apache)[11099]: INFO: apache is not running.
Mar 24 11:29:24 hideo.tc.redhat.com crmd[2106]: notice: process_lrm_event: LRM operation summit-apache_stop_0 (call=70, rc=0, cib-update=55, confirmed=true) ok
Mar 24 11:29:24 hideo.tc.redhat.com systemd[1]: pacemaker.service: Got notification message from PID 11146, but reception only permitted for PID 1749
Mar 24 11:29:24 hideo.tc.redhat.com apache(summit-apache)[11165]: INFO: Successfully retrieved http header at http://localhost:80

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# ps ax | grep httpd
```

```
21426 ?          Ss      0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
```

```
  PidFile /var/run/httpd.pid
```

```
21430 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
```

```
  PidFile /var/run/httpd.pid
```

```
21432 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
```

```
  PidFile /var/run/httpd.pid
```

```
21433 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
```

```
  PidFile /var/run/httpd.pid
```

```
21435 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
```

```
  PidFile /var/run/httpd.pid
```

```
21439 ?          S       0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
```

```
  PidFile /var/run/httpd.pid
```

```
21559 pts/0      S+      0:00 grep --color=auto httpd
```

```
[root@hideo ~]#
```

Resource Location Preferences

- You can also set up host affinity via Resource Location Preferences
- Add each of the hosts you want to run the service on and add a score. The higher the score, the more likely the service is to run on that node.

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input checked="" type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-ip

summit-ip Running

Enable Disable Cleanup Remove

Current Location: hideo.tc.redhat.com

Resource Location Preferences (0)

Node	Score	Remove
NONE		
hideo.tc.redhat.com	10	Add

Resource Ordering Preferences (1)

Resource Colocation Preferences (1)

Resource Meta Attributes (0)

root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started hideo.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started hideo.tc.redhat.com


PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@hideo ~]#



Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input checked="" type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-ip

summit-ip Running

Enable Disable Cleanup Remove

Current Location: hideo.tc.redhat.com

Resource Location Preferences (2)

Node	Score	Remove
hideo.tc.redhat.com	10	X
lady3jane.tc.redhat.com	20	X
		Add

Resource Ordering Preferences (1)

Resource Colocation Preferences (1)

Resource Meta Attributes (0)

root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started lady3jane.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started lady3jane.tc.redhat.com

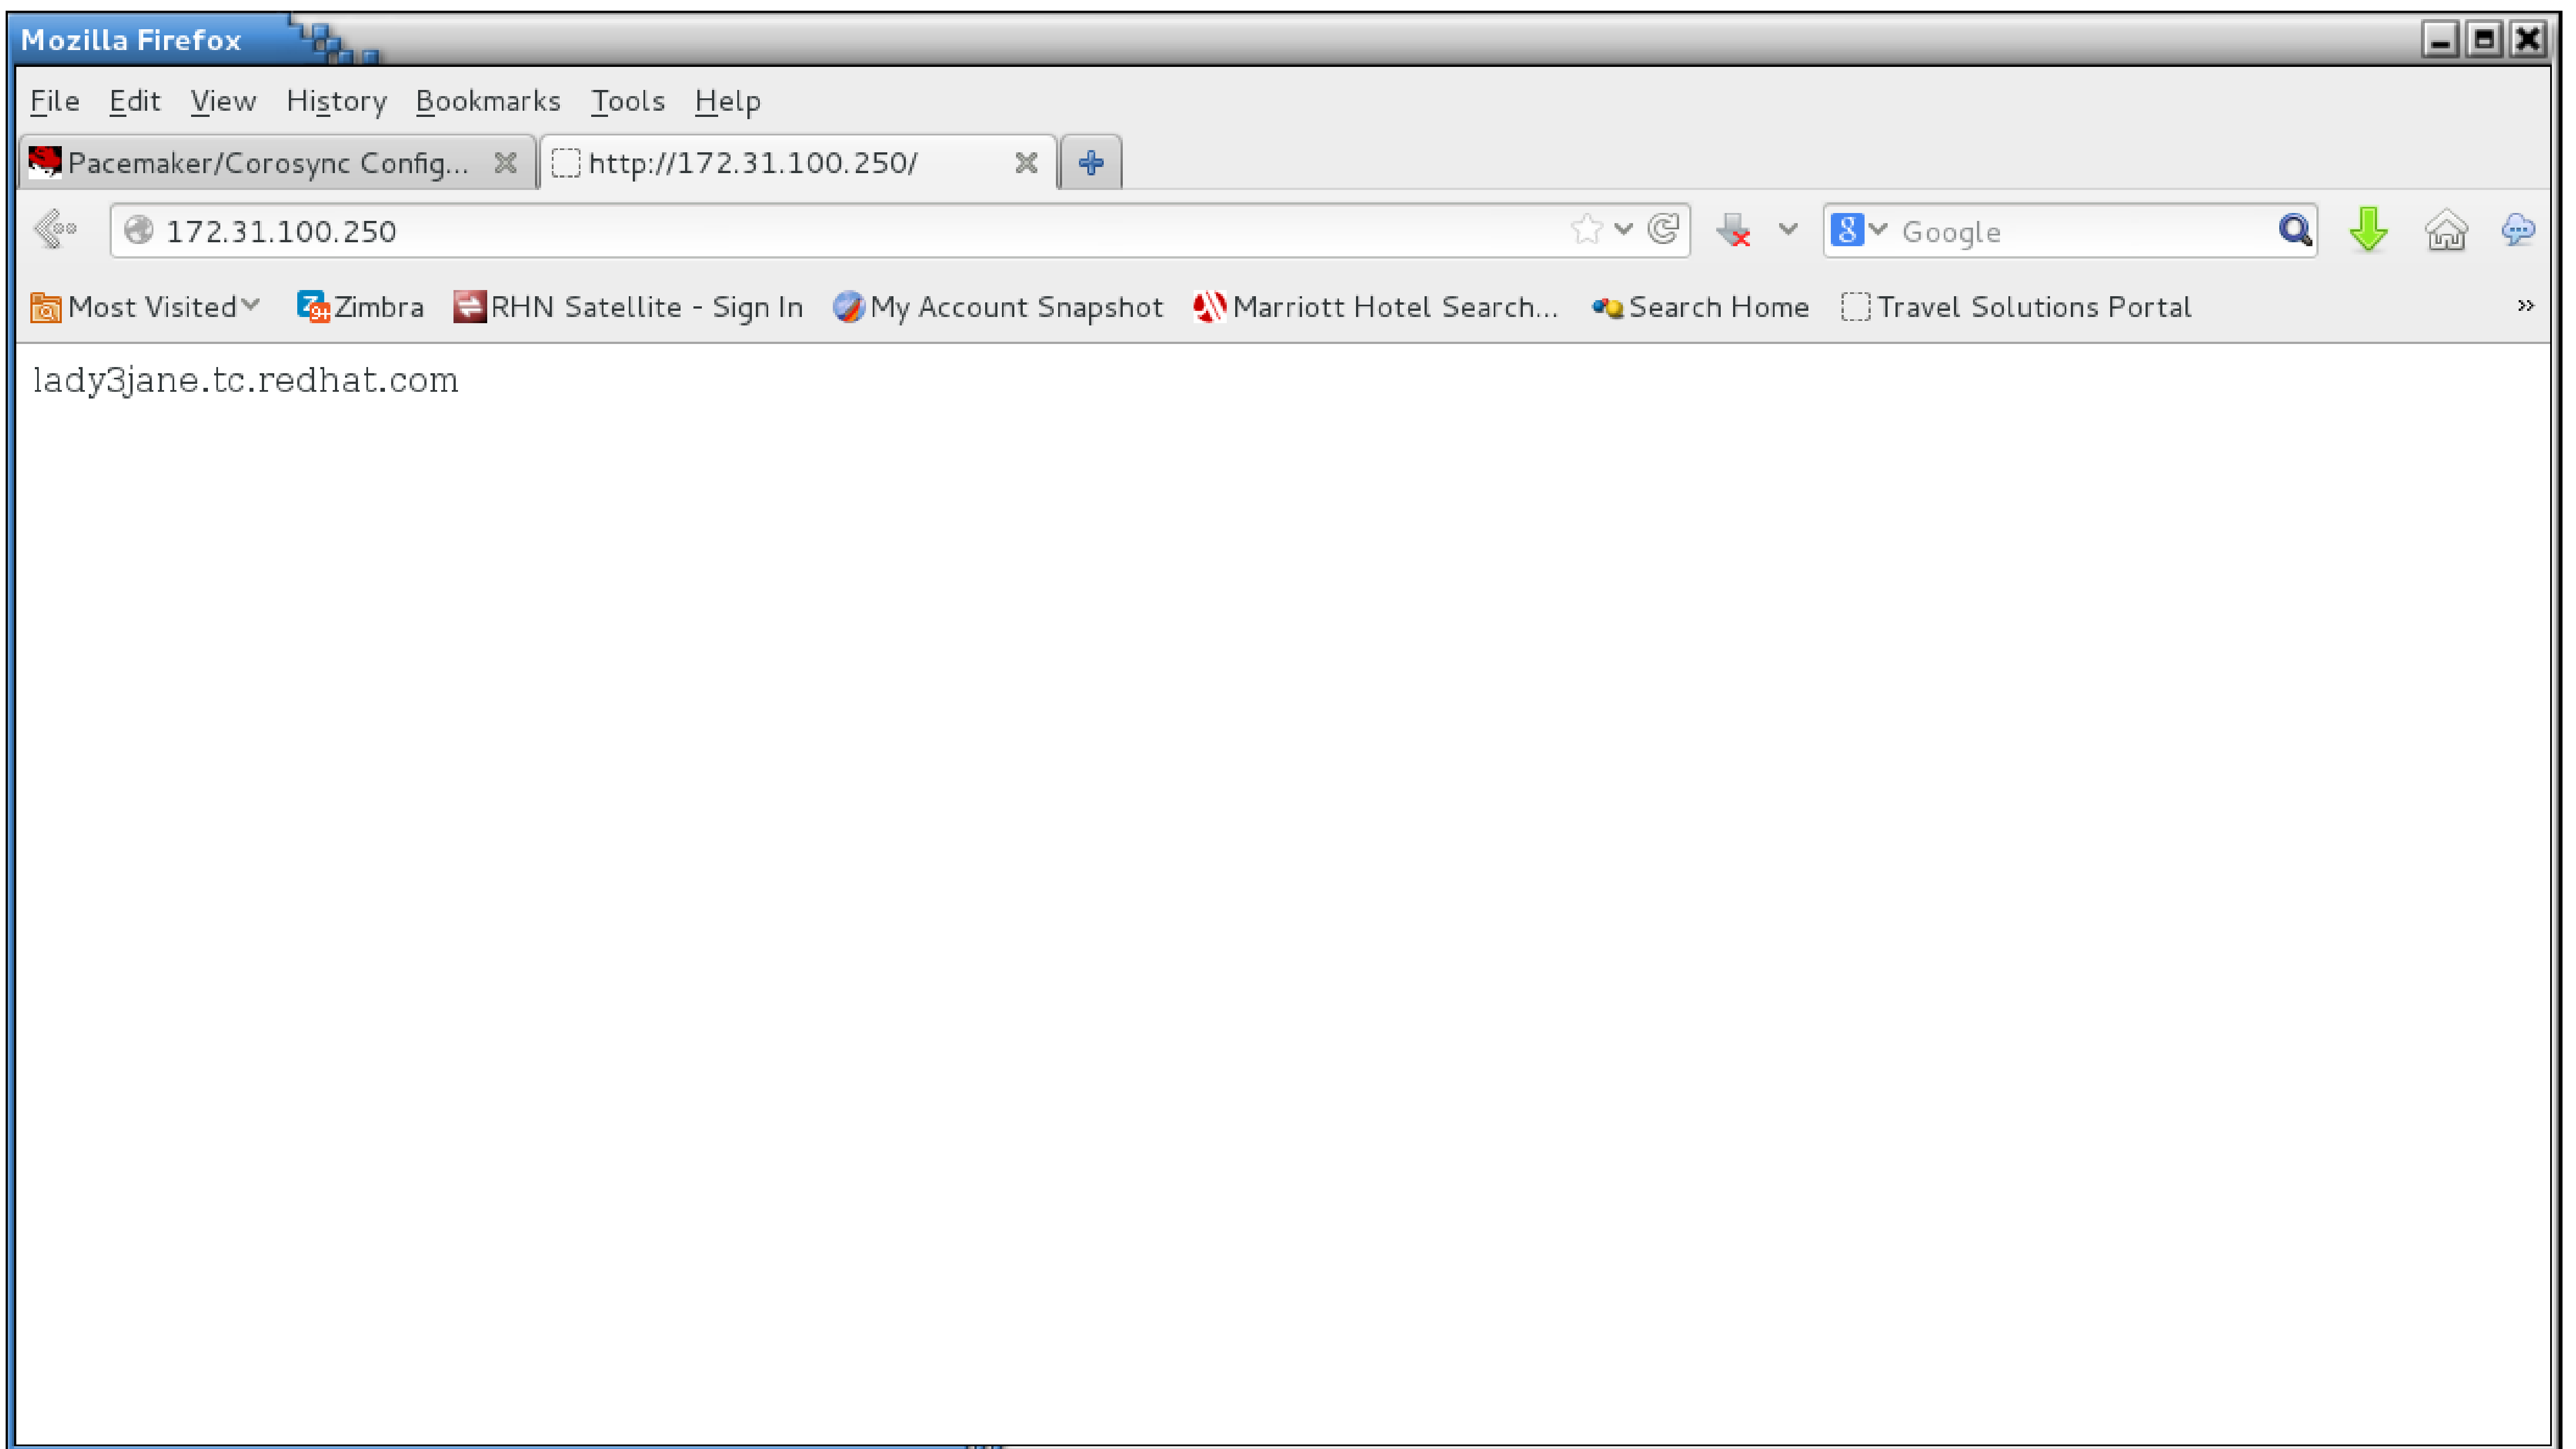
PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@hideo ~]#



Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource ✓

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input checked="" type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-ip

summit-ip ✓ Running

✓ Enable ✕ Disable ↻ Cleanup ✕ Remove

Current Location: lady3jane.tc.redhat.com

Resource Location Preferences (3)

Node	Score	Remove
hideo.tc.redhat.com	10	✕
lady3jane.tc.redhat.com	20	✕
maelcum.tc.redhat.com	30	✕
<input type="text"/>	<input type="text"/>	Add

Resource Ordering Preferences (1)

Resource Colocation Preferences (1)

root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started maelcum.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started maelcum.tc.redhat.com

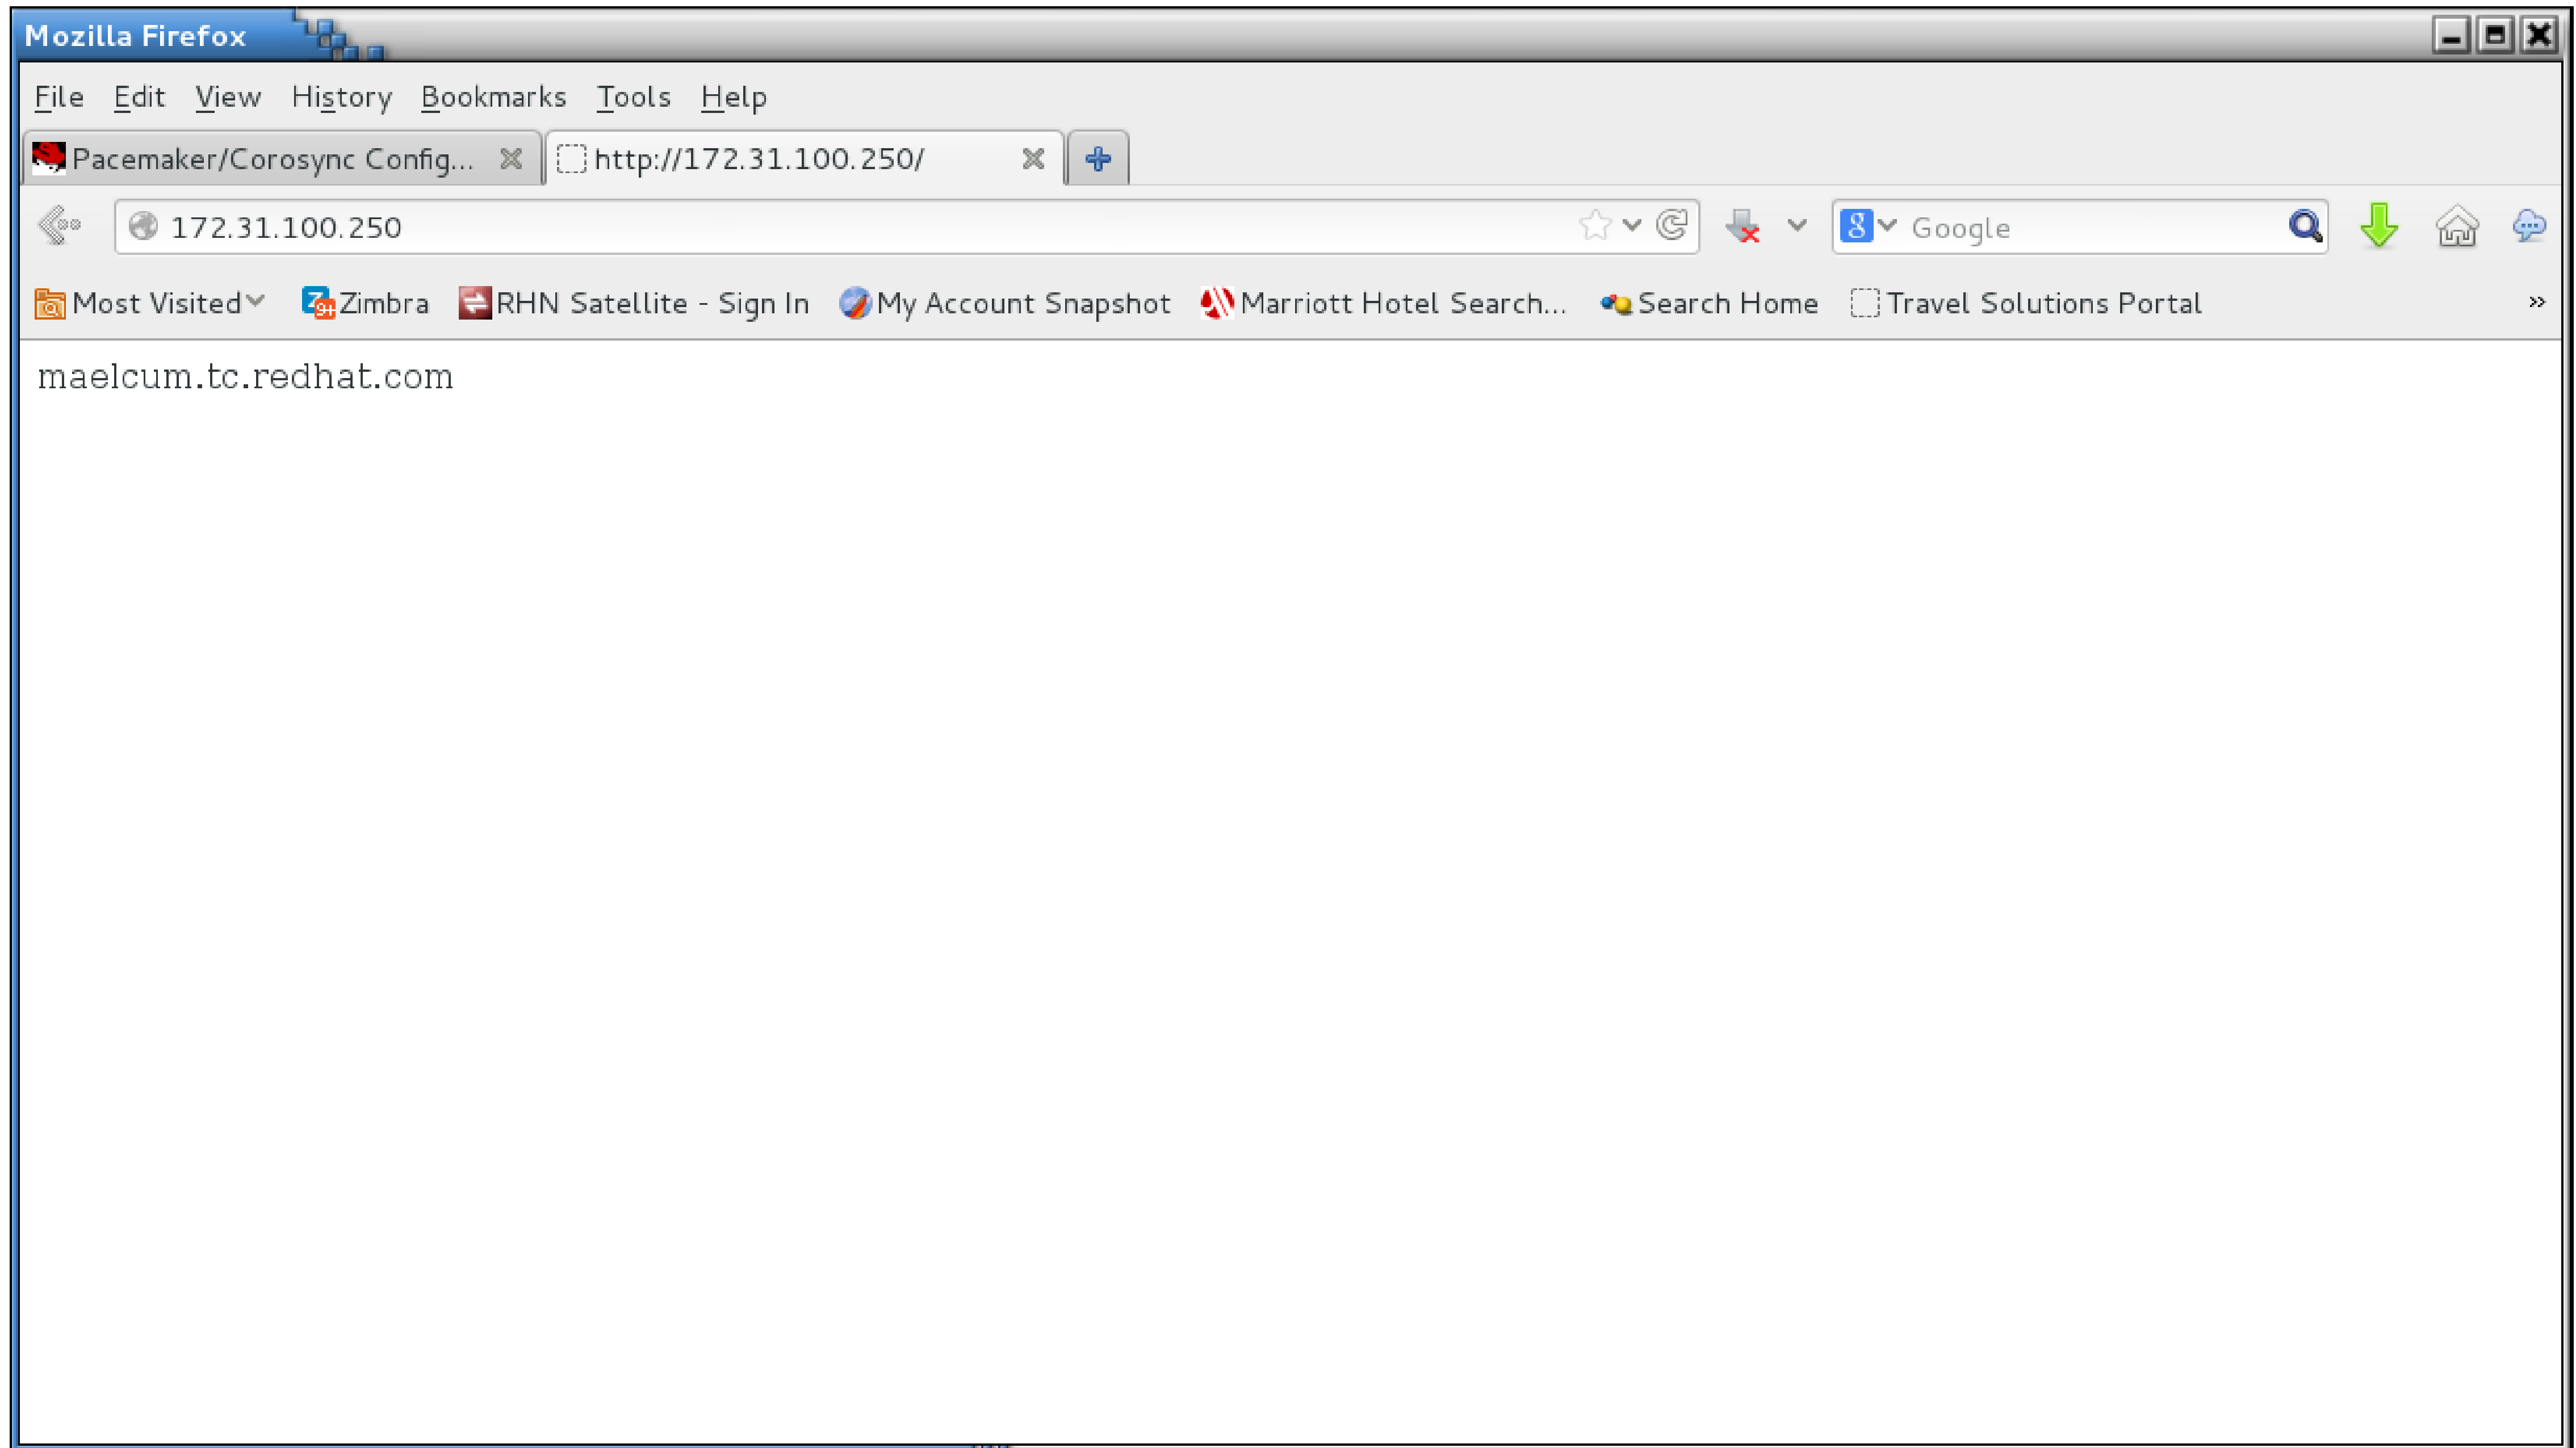
PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@hideo ~]#



To force a service to a host

- Use a score of INFINITY

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

Remove Add Create Group Edit Resource

	NAME	TYPE
<input type="checkbox"/>	summit-apache	ocf::heartbeat:apache
<input checked="" type="checkbox"/>	summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-ip

summit-ip Running

Enable Disable Cleanup Remove

Current Location: hideo.tc.redhat.com

Resource Location Preferences (3)

Node	Score	Remove
lady3jane.tc.redhat.com	20	X
maelcum.tc.redhat.com	30	X
hideo.tc.redhat.com	INFINITY	X
		Add

Resource Ordering Preferences (1)

Resource Colocation Preferences (1)

root@hideo:~

File Edit View Search Terminal Help

Current DC: maelcum.tc.redhat.com (3) - partition with quorum

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started hideo.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started hideo.tc.redhat.com


PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@hideo ~]#



To offline a host

- Go to Nodes
- Choose the host
- Choose “Stop”

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/nodes/hideo.tc.redhat.com

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

NODES

☐ X Remove + Add Edit Node ✓

- ☐ ✓ hideo.tc.redhat.com
- ☐ ✓ lady3jane.tc.redhat.com
- ☐ ✓ maelcum.tc.redhat.com

Edit Node hideo.tc.redhat.com

hideo.tc.redhat.com

✓ Pacemaker Not Connected
✓ Corosync Not Connected

Start Stop Restart Standby Configure Fencing

Node ID: Uptime: 0 days, 04:17:13

Cluster Daemons

NAME	STATUS
pacemaker	✓ Stopped (Enabled)
corosync	✓ Stopped (Enabled)
pcsd	✓ Running (Enabled)

root@maelcum:~

File Edit View Search Terminal Help

Version: 1.1.10-27.el7-368c726

3 Nodes configured

3 Resources configured

Online: [lady3jane.tc.redhat.com maelcum.tc.redhat.com]

OFFLINE: [hideo.tc.redhat.com]

Full list of resources:

summit-wti	(stonith:fence_wti):	Started lady3jane.tc.redhat.com
summit-ip	(ocf::heartbeat:IPaddr2):	Started maelcum.tc.redhat.com
summit-apache	(ocf::heartbeat:apache):	Started maelcum.tc.redhat.com

PCSD Status:

hideo.tc.redhat.com: Online
lady3jane.tc.redhat.com: Online
maelcum.tc.redhat.com: Online

Daemon Status:

corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled

[root@maelcum ~]#

To online a host

- Go to Nodes
- Choose the host
- Choose “Start”

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/nodes/hideo.tc.redhat.com

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

NODES

☐ x Remove + Add Edit Node ✓

- ☐ ✓ hideo.tc.redhat.com
- ☐ ✓ lady3jane.tc.redhat.com
- ☐ ✓ maelcum.tc.redhat.com

Edit Node hideo.tc.redhat.com

hideo.tc.redhat.com

- ✓ Pacemaker Connected
- ✓ Corosync Connected

Start Stop Restart Standby Configure Fencing

Node ID: 1 Uptime: 0 days, 04:20:59

Cluster Daemons

NAME	STATUS
pacemaker	✓ Running (Enabled)
corosync	✓ Running (Enabled)
pcsd	✓ Running (Enabled)


```
root@hideo:~  
File Edit View Search Terminal Help  
Current DC: maelcum.tc.redhat.com (3) - partition with quorum  
Version: 1.1.10-27.el7-368c726  
3 Nodes configured  
3 Resources configured  
  
Online: [ hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com ]  
  
Full list of resources:  
  
summit-wti      (stonith:fence_wti):      Started lady3jane.tc.redhat.com  
summit-ip       (ocf::heartbeat:IPaddr2):  Started maelcum.tc.redhat.com  
summit-apache   (ocf::heartbeat:apache):   Started maelcum.tc.redhat.com  
  
PCSD Status:  
hideo.tc.redhat.com: Online  
lady3jane.tc.redhat.com: Online  
maelcum.tc.redhat.com: Online  
  
Daemon Status:  
corosync: active/enabled  
pacemaker: active/enabled  
pcsd: active/enabled  
[root@hideo ~]#
```

To reboot a host

- Go to Nodes
- Choose the host
- Choose “Restart”
 - Note: this will execute a controlled reboot within the OS (shutdown -r now), **not** a fencing event

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE DISTRIBUTED LOCK MANAGEMENT

Enable distributed lock management

- Create a new `ocf:pacemaker` class resource of control
- Check the box for “clone” - we want this service cloned across all the nodes
- Give it a Resource ID - in this case, `summit-dlm`
- After a few seconds it should start and turn green


Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-dlm

Most Visited ▾ Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal >>

 HIGH AVAILABILITY MANAGEMENT Cluster: summit ▾ hacluster ▾

NODES RESOURCES FENCE DEVICES CLUSTER PROPERTIES MANAGE CLUSTERS

RESOURCES

✕ Remove + Add + Create Group Edit Resource ✓

<input type="checkbox"/>	NAME	TYPE
<input type="checkbox"/>	✓ summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	✓ summit-dlm (Clone)	ocf::pacemaker:controld ▶
<input type="checkbox"/>	✓ summit-ip	ocf::heartbeat:IPaddr2

Edit Resource **summit-dlm**

summit-dlm ✓ Running

✓ Enable ✕ Disable ↻ Cleanup ✕ Remove

Current Location: hideo.tc.redhat.com, lady3jane.tc.redhat.com, maelcum.tc.redhat.com

Resource Location Preferences (0)

Node	Score	Remove
NONE		
<input type="text"/>	<input type="text"/>	Add

Resource Ordering Preferences (0)

Enable distributed lock management

- You can also make sure the `dlm_controld` process is running on all nodes

```
root@hideo:~
File Edit View Search Terminal Help
1552 ? Ss 0:00 /usr/sbin/pacemakerd -f
1576 ? S 0:00 [scsi_eh_8]
1577 ? S< 0:00 [scsi_tmf_8]
1578 ? S< 0:00 [iscsi_q_8]
1580 ? S< 0:00 [scsi_wq_8]
1757 ? Ss 0:01 /usr/libexec/pacemaker/cib
1758 ? Ss 0:00 /usr/libexec/pacemaker/stonithd
1760 ? Ss 0:00 /usr/libexec/pacemaker/lrmd
1763 ? Ss 0:00 /usr/libexec/pacemaker/attrd
1764 ? Ss 0:00 /usr/libexec/pacemaker/pengine
1765 ? Ss 0:00 /usr/libexec/pacemaker/crmd
1907 ? Ss 0:00 /usr/libexec/postfix/master -w
1962 ? S 0:00 pickup -l -t unix -u
1963 ? S 0:00 qmgr -l -t unix -u
2349 ? S 0:00 [kworker/5:0]
2382 ? Ss 0:00 sshd: root@pts/0
2385 pts/0 Ss 0:00 -bash
2426 ? S 0:00 [kworker/4:1]
3428 ? S 0:00 [kworker/0:0]
3475 ? S 0:00 [kworker/u16:0]
3477 ? R 0:00 [kworker/4:2]
5753 ? Ssl 0:00 dlm_control -q 0 -s 0
6321 pts/0 R+ 0:00 ps ax
[root@hideo ~]#
```


RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE DISTRIBUTED LOGICAL VOLUME MANAGEMENT

Enable Clustered Logical Volume Management

- Add an `ofc:heartbeat` Class/Provider for `clvm`
- Set it as cloned - we want this running on all nodes
- Give it a Resource ID - in this case, `summit-clvmd`

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

HIGH AVAILABILITY MANAGEMENT Cluster: summit hacluster

RESOURCES

Remove Add Create Group

	NAME
<input type="checkbox"/>	summit-apache
<input checked="" type="checkbox"/>	summit-dlm (Clone)
<input checked="" type="checkbox"/>	summit-ip

Add Resource

Class/Provider: ocf:heartbeat

Type: clvm

Description: clvmd ⓘ

Resource Group: None

Clone: ☒

Master/Slave: ☐

Disabled: ☐

Resource ID: summit-clvmd

Optional Arguments:

Create Resource

Running

Resource Ordering Preferences (1)

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura...

+

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-clvmd

☆

🔄

📄

🔍 Google

📁

🏠

💬

Most Visited

Zimbra

RHN Satellite - Sign In


My Account Snapshot

Marriott Hotel Search...

Search Home

Travel Solutions Portal

>>

 HIGH AVAILABILITY MANAGEMENT

Cluster: summit

hacluster

NODES

RESOURCES

FENCE DEVICES

CLUSTER PROPERTIES

MANAGE CLUSTERS

RESOURCES

✕ Remove

+ Add

+ Create Group

Edit Resource ✓

<input type="checkbox"/>	NAME	TYPE
<input type="checkbox"/>	✓ summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	✓ summit-clvmd (Clone)	ocf::heartbeat:clvm
<input type="checkbox"/>	✓ summit-dlm (Clone)	ocf::pacemaker:controld
<input type="checkbox"/>	✓ summit-ip	ocf::heartbeat:IPaddr2

Edit Resource summit-clvmd

summit-clvmd ✓ Running

✓ Enable ✕ Disable 🔄 Cleanup ✕ Remove


Current Location: hideo.tc.redhat.com, lady3jane.tc.redhat.com, maelcum.tc.redhat.com

Resource Location Preferences (0)

Node	Score	Remove
NONE		
<input type="text"/>	<input type="text"/>	Add

Resource Ordering Preferences (0)

#redhat #rhsummit

 redhat.

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# ps ax|grep lvm  
  333 ?          Ss      0:00 /usr/sbin/lvmetad -f  
 8801 ?          SLsl    0:00 /usr/sbin/clvmd -T90  
 9489 pts/0      S+      0:00 grep --color=auto lvm  
[root@hideo ~]#
```

RED HAT
SUMMIT

10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE CLUSTERED LOGICAL VOLUME MANAGEMENT

Change `locking_type` from 1 to 3 in `/etc/lvm/lvm.conf`

- `lvmconf --enable-cluster`

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# grep locking_type /etc/lvm/lvm.conf
# supported in clustered environment. If use_lvmetad=1 and locking_type=3
locking_type = 1
# NB. This option only affects locking_type = 1 viz. local file-based
# The external locking library to load if locking_type is set to 2.
# supported in clustered environment. If use_lvmetad=1 and locking_type=3
[root@hideo ~]# lvmconf --enable-cluster
[root@hideo ~]# grep locking_type /etc/lvm/lvm.conf
# supported in clustered environment. If use_lvmetad=1 and locking_type=3
locking_type = 3
# NB. This option only affects locking_type = 1 viz. local file-based
# The external locking library to load if locking_type is set to 2.
# supported in clustered environment. If use_lvmetad=1 and locking_type=3
[root@hideo ~]#
```


Note About the LVM Metadata Cache Daemon (lvmetad)

- From `lvmetad(8)`: `lvmetad` is a metadata caching daemon for LVM. The daemon receives notifications from `udev` rules (which must be installed for LVM to work correctly when `lvmetad` is in use). Through these notifications, `lvmetad` has an up-to-date and consistent image of the volume groups available in the system.
- From `/etc/lvm/lvm.conf`: Don't use `lvmetad` with locking type 3 as `lvmetad` is not yet supported in clustered environment. If `use_lvmetad=1` and `locking_type=3` is set at the same time, LVM always issues a warning message about this and then it automatically disables `lvmetad` use.

Note About the LVM Metadata Cache Daemon (lvmetad)

- Change `use_lvmetad = 1` to `use_lvmetad = 0`
- `perl -pi.orig -e 's/use_lvmetad = 1/use_lvmetad = 0/' /etc/lvm/lvm.conf`

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# perl -pi.orig -e 's/use_lvmetad = 1/use_lvmetad = 0/' /etc/lvm/l  
vm.conf  
[root@hideo ~]# diff /etc/lvm/lvm.conf /etc/lvm/lvm.conf.orig  
665c665  
<     use_lvmetad = 0  
---  
>     use_lvmetad = 1  
[root@hideo ~]#
```



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

CONFIGURE SHARED STORAGE

Create a partition on the iSCSI block device

- Use your favorite partitioning tool. Here, I use fdisk to create a 10G partition

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# fdisk /dev/sdb  
Welcome to fdisk (util-linux 2.23.2).  
  
Changes will remain in memory only, until you decide to write them.  
Be careful before using the write command.  
  
Device does not contain a recognized partition table  
Building a new DOS disklabel with disk identifier 0x624f48b2.  
  
Command (m for help): n  
Partition type:  
   p   primary (0 primary, 0 extended, 4 free)  
   e   extended  
Select (default p): p  
Partition number (1-4, default 1): 1  
First sector (8192-950966319, default 8192):  
Using default value 8192  
Last sector, +sectors or +size{K,M,G} (8192-950966319, default 950966319): +10G  
Partition 1 of type Linux and of size 10 GiB is set  
  
Command (m for help): w  
The partition table has been altered!  
  
Calling ioctl() to re-read partition table.  
Syncing disks.  
[root@hideo ~]#
```

Log out and back in to discover the partitions on the other nodes

- `iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --logout`
- `iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --login`


```
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# cat /proc/partitions  
major minor #blocks name  
  
8          0  976762584 sda  
8          1    512000 sda1  
8          2  10240000 sda2  
8          3   2150400 sda3  
8         16 475483160 sdb  
[root@lady3jane ~]# iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi  
i.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --logout  
Logging out of session [sid: 1, target: iqn.2003-01.org.linux-iscsi.neuromancer.  
x8664:sn.f63ec35cd646, portal: 172.31.100.14,3260]  
Logout of [sid: 1, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:sn.f63e  
c35cd646, portal: 172.31.100.14,3260] successful.  
[root@lady3jane ~]# iscsiadm --mode node --targetname iqn.2003-01.org.linux-iscsi  
i.neuromancer.x8664:sn.f63ec35cd646 --portal neuromancer.tc.redhat.com --login  
Logging in to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x  
8664:sn.f63ec35cd646, portal: 172.31.100.14,3260] (multiple)  
Login to [iface: default, target: iqn.2003-01.org.linux-iscsi.neuromancer.x8664:  
sn.f63ec35cd646, portal: 172.31.100.14,3260] successful.  
[root@lady3jane ~]# cat /proc/partitions  
major minor #blocks name  
  
8          0  976762584 sda  
8          1    512000 sda1  
8          2  10240000 sda2  
8          3   2150400 sda3  
8         16 475483160 sdb  
8         17  10485760 sdb1  
[root@lady3jane ~]#
```


Create a Physical Volume

- `pvccreate /dev/sdb1`

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# pvcreate /dev/sdb1  
Physical volume "/dev/sdb1" successfully created  
[root@hideo ~]#
```

Create a Volume Group

- `vgcreate --clustered y vg_summit /dev/sdb1`
 - Don't forget the “--clustered y” argument!

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# vgcreate -c y vg_summit /dev/sdb1  
  Clustered volume group "vg_summit" successfully created  
[root@hideo ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# vgcreate -c y vg_summit /dev/sdb1  
  Clustered volume group "vg_summit" successfully created  
[root@hideo ~]# vgsdisplay  
  --- Volume group ---  
  VG Name                vg_summit  
  System ID  
  Format                 lvm2  
  Metadata Areas         1  
  Metadata Sequence No   1  
  VG Access              read/write  
  VG Status              resizable  
  Clustered              yes  
  Shared                 no  
  MAX LV                 0  
  Cur LV                 0  
  Open LV                0  
  Max PV                 0  
  Cur PV                 1  
  Act PV                 1  
  VG Size                 10.00 GiB  
  PE Size                 4.00 MiB  
  Total PE               2559  
  Alloc PE / Size        0 / 0  
  Free PE / Size         2559 / 10.00 GiB  
  VG UUID                oFSCbH-kszB-fx9x-AmMF-WiSP-Pgu7-kQfZ2J  
  
[root@hideo ~]#
```

Create a Logical Volume

- `lvcreate -l 2559 -n lv_web vg_summit`

You can use pvs, lvs and vgs to scan LVM components from the other nodes

- From each node, just run those commands:
 - pvs (reports info about physical volumes)
 - vgs (reports information about volume groups)
 - lvs (reports information about logical volumes)

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#  
  
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# lvcreate -l 2559 -n lv_web vg_summit  
Logical volume "lv_web" created  
[root@hideo ~]#  
  
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# pvs  
PV          VG          Fmt Attr PSize  PFree  
/dev/sdb1   vg_summit lvm2 a-- 10.00g  0  
[root@lady3jane ~]# vgs  
VG          #PV #LV #SN Attr   VSize  VFree  
vg_summit   1   1   0 wz--nc 10.00g  0  
[root@lady3jane ~]# lvs  
LV          VG          Attr              LSize  Pool Origin Data%  Move Log Cpy%Sync Conver  
t  
lv_web      vg_summit -wi-a----- 10.00g  
[root@lady3jane ~]#  
  
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# pvs  
PV          VG          Fmt Attr PSize  PFree  
/dev/sdb1   vg_summit lvm2 a-- 10.00g 10.00g  
[root@maelcum ~]# vgs  
VG          #PV #LV #SN Attr   VSize  VFree  
vg_summit   1   0   0 wz--nc 10.00g 10.00g  
[root@maelcum ~]# lvs  
LV          VG          Attr              LSize  Pool Origin Data%  Move Log Cpy%Sync Conver  
t  
lv_web      vg_summit -wi-a----- 10.00g  
[root@maelcum ~]#
```


Install gfs2-utils

- `yum -y install gfs2-utils`

```
root@maelcum:~  
File Edit View Search Terminal Help  
=====
```

Installing:				
gfs2-utils	x86_64	3.1.6-13.el7	rhel	247 k

```
Transaction Summary  
=====
```

Install 1 Package

Total download size: 247 k
Installed size: 695 k
Downloading packages:

gfs2-utils-3.1.6-13.el7.x86_64.rpm	247 kB	00:00
------------------------------------	--------	-------

```
Running transaction check  
Running transaction test  
Transaction test succeeded  
Running transaction  
  Installing : gfs2-utils-3.1.6-13.el7.x86_64      1/1  
  Verifying   : gfs2-utils-3.1.6-13.el7.x86_64      1/1  
  
Installed:  
  gfs2-utils.x86_64 0:3.1.6-13.el7  
  
Complete!  
[root@maelcum ~]#
```

Create a GFS2 Filesystem on the Clustered Logical Volume

- `mkfs.gfs2 -j 3 -t summit:gfs0 /dev/vg_summit/lv_web`
 - `-j 3` is the number of journals - one per node. Extras are fine, too.
 - `-t summit:gfs0` is `[clustername]:[fsname]`. “summit” is the name of the cluster we defined, and “gfs0” is the name I gave the filesystem being created.
 - `/dev/vg_summit/lv_web` is the block device being formatted. In this case, a clustered logical volume

root@hideo:~

File Edit View Search Terminal Help

```
[root@hideo ~]# mkfs.gfs2 -j 3 -t summit:gfs0 /dev/vg_summit/lv_web
```

/dev/vg_summit/lv_web is a symbolic link to /dev/dm-0

This will destroy any data on /dev/dm-0

Are you sure you want to proceed? [y/n]y

Device: /dev/vg_summit/lv_web

Block size: 4096

Device size: 10.00 GB (2620416 blocks)

Filesystem size: 10.00 GB (2620413 blocks)

Journals: 3

Resource groups: 40

Locking protocol: "lock_dlm"

Lock table: "summit:gfs0"

UUID: 959505b8-8eef-760e-4b11-b742bc9f6e33

```
[root@hideo ~]#
```



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

ADD SHARED STORAGE AS A RESOURCE

Create a New Filesystem Resource

- Create a new ofc:heartbeat resource of type Filesystem
- Check the box to clone the resource - we wanted it mounted on all the servers
- Give it a Resource ID - in this case, “summit-gfs0”
- Define the block device - in this case, the clustered logical volume “/dev/vg_summit/lv_web”
- Define the filesystem type - in this case, gfs2
- Click “Create Resource”



RESOURCES

x Remove + Add + Create Group

<input type="checkbox"/>	NAME
<input type="checkbox"/>	summit-apache
<input type="checkbox"/>	summit-clvmd (Clone)
<input type="checkbox"/>	summit-dlm (Clone)
<input type="checkbox"/>	summit-ip

Add Resource



Class/Provider ocf:heartbeat

Type Filesystem

Description: Manages filesystem mounts ⓘ

Resource Group: None

Clone: ☒Master/Slave: ☐Disabled: ☐

Resource ID summit-gfs0

device /dev/vg_summit/lv_web

directory /var/www/html

fstype gfs2

▶ Optional Arguments:

Create Resource

CLUSTERS

Edit Resource ✓

✓ Running

remove

Add

Score

Remove

RESOURCES

[✕ Remove](#) [+ Add](#) [+ Create Group](#)[Edit Resource](#) ✓

<input type="checkbox"/>	NAME	TYPE
<input type="checkbox"/>	✓ summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	✓ summit-clvmd (Clone)	ocf::heartbeat:clvm
<input type="checkbox"/>	✓ summit-dlm (Clone)	ocf::pacemaker:controld
<input type="checkbox"/>	✓ summit-gfs0 (Clone)	ocf::heartbeat:Filesystem ▶
<input type="checkbox"/>	✓ summit-ip	ocf::heartbeat:IPaddr2

Edit Resource **summit-gfs0**

✓ Running

[✓ Enable](#) [✕ Disable](#) [🔄 Cleanup](#) [✕ Remove](#)**Current Location:** hideo.tc.redhat.com, lady3jane.tc.redhat.com, maelcum.tc.redhat.com

Resource Location Preferences (0)



Node	Score	Remove
NONE		
<input type="text"/>	<input type="text"/>	<input type="button" value="Add"/>

Resource Ordering Preferences (0)



Resource	Before/After	Score	Remove
NONE			

The GFS2 Mountpoint Should Be Visible on All Nodes

- Run “mount | grep gfs2” on each node

```
root@neuromancer:~  
File Edit View Search Terminal Help  
[root@neuromancer ~]# # iSCSI target - don't install anything else here  
[root@neuromancer ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# mount | grep gfs2  
/dev/mapper/vg_summit-lv_web on /var/www/html type gfs2 (rw,relatime,seclabel)  
[root@hideo ~]#
```

```
root@lady3jane:~  
File Edit View Search Terminal Help  
[root@lady3jane ~]# mount | grep gfs2  
/dev/mapper/vg_summit-lv_web on /var/www/html type gfs2 (rw,relatime,seclabel)  
[root@lady3jane ~]#
```

```
root@maelcum:~  
File Edit View Search Terminal Help  
[root@maelcum ~]# mount | grep gfs2  
/dev/mapper/vg_summit-lv_web on /var/www/html type gfs2 (rw,relatime,seclabel)  
[root@maelcum ~]#
```

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# pcs status  
Cluster name: summit  
Last updated: Mon Mar 24 21:33:37 2014  
Last change: Mon Mar 24 21:30:09 2014 via cibadmin on hideo.tc.redhat.com  
Stack: corosync  
Current DC: lady3jane.tc.redhat.com (2) - partition with quorum  
Version: 1.1.10-27.el7-368c726  
3 Nodes configured  
12 Resources configured  
  
Online: [ hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com ]  
  
Full list of resources:  
  
summit-wti (stonith:fence_wti): Started hideo.tc.redhat.com  
summit-ip (ocf::heartbeat:IPaddr2): Started maelcum.tc.redhat.com  
summit-apache (ocf::heartbeat:apache): Started maelcum.tc.redhat.com  
Clone Set: summit-dlm-clone [summit-dlm]  
Started: [ hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com ]  
Clone Set: summit-clvmd-clone [summit-clvmd]  
Started: [ hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com ]  
Clone Set: summit-gfs0-clone [summit-gfs0]  
Started: [ hideo.tc.redhat.com lady3jane.tc.redhat.com maelcum.tc.redhat.com ]  
  
PCSD Status:  
hideo.tc.redhat.com: Online  
lady3jane.tc.redhat.com: Online  
maelcum.tc.redhat.com: Online  
  
Daemon Status:  
corosync: active/enabled  
pacemaker: active/enabled  
pcsd: active/enabled  
[root@hideo ~]#
```

Test Your Web Site

- Create an index.html
 - `echo gfs > /var/www/html/index.html`

```
root@hideo:~  
File Edit View Search Terminal Help  
[root@hideo ~]# echo gfs > /var/www/html/index.html  
[root@hideo ~]#
```

Migrate the Web Site

- Set the priority of the ip address resource to INFINITY (one at a time) on each node

Pacemaker/Corosync Configuration - Mozilla Firefox

File Edit View History Bookmarks Tools Help

Pacemaker/Corosync Configura... +

https://hideo.tc.redhat.com:2224/managec/summit/main#/resources/summit-ip

Most Visited Zimbra RHN Satellite - Sign In My Account Snapshot Marriott Hotel Search... Search Home Travel Solutions Portal

RESOURCES

[x Remove](#) [+ Add](#) [+ Create Group](#) [Edit Resource ✓](#)

<input type="checkbox"/>	NAME	TYPE
<input type="checkbox"/>	✓ summit-apache	ocf::heartbeat:apache
<input type="checkbox"/>	✓ summit-clvmd (Clone)	ocf::heartbeat:clvm
<input type="checkbox"/>	✓ summit-dlm (Clone)	ocf::pacemaker:controld
<input type="checkbox"/>	✓ summit-gfs0 (Clone)	ocf::heartbeat:Filesystem
<input type="checkbox"/>	✓ summit-ip	ocf::heartbeat:IPaddr2 ▶

Edit Resource **summit-ip**

summit-ip ✓ Running

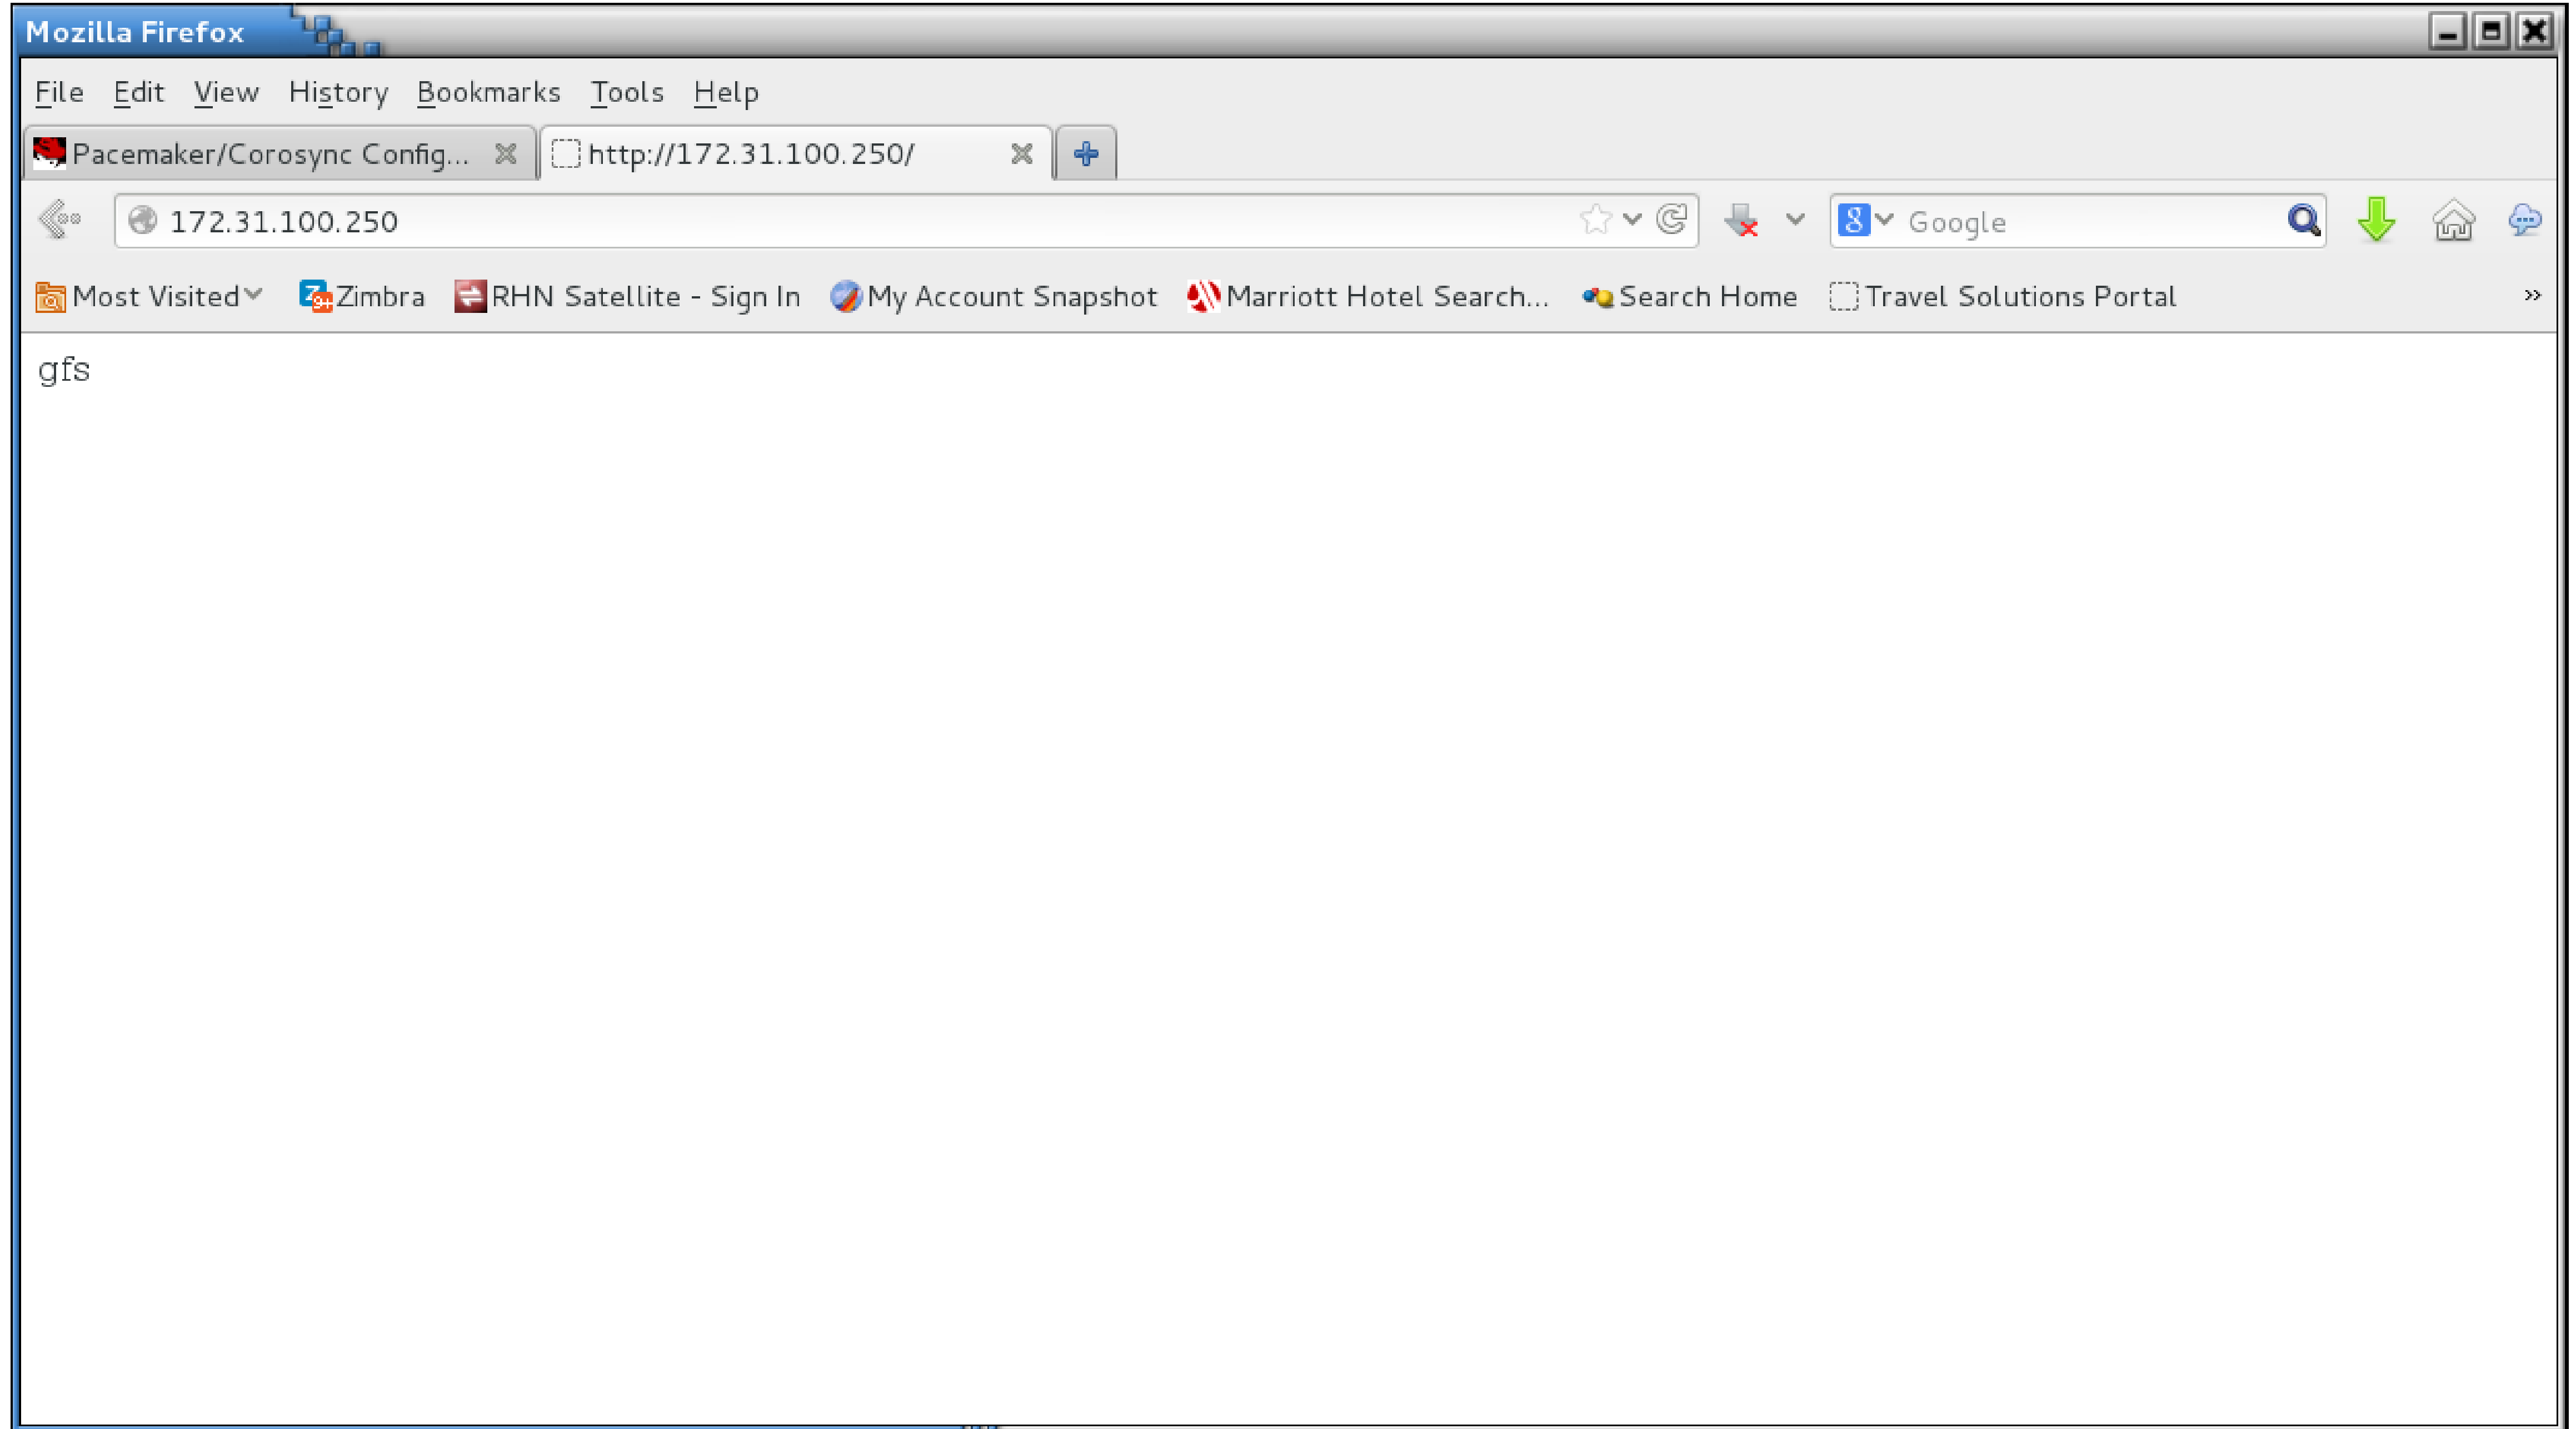
✓ Enable ✕ Disable ↻ Cleanup ✕ Remove

Current Location: maelcum.tc.redhat.com

Resource Location Preferences (3)

Node	Score	Remove
maelcum.tc.redhat.com	30	X
hideo.tc.redhat.com	10	X
lady3jane.tc.redhat.com	INFINITY	X

Resource Ordering Preferences (1)





10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

TEST CRASHES

We've Already Crashed httpd

- Now let's crash the cluster nodes!
- sync the filesystem first!
- `sync; echo s > /proc/sysrq-trigger`
- `echo c > /proc/sysrq-trigger`

root@lady3jane:~

```
File Edit View Search Terminal Help
16134 ? S< 0:00 [bioset]
16142 ? S 0:00 [kworker/1:2]
20201 ? S 0:00 pickup -l -t unix -u
21188 ? S< 0:00 [glock_workqueue]
21189 ? S< 0:00 [delete_workqueue]
21190 ? S< 0:00 [gfs_recovery]
21241 ? S 0:00 [kworker/3:0]
21242 ? S< 0:00 [kworker/0:2H]
21250 ? S 0:00 [kworker/u16:1]
21934 ? S< 0:00 [dlm_callback]
21935 ? S 0:00 [dlm_recoverd]
21972 ? S< 0:00 [kworker/2:2H]
21973 ? S< 0:00 [kworker/3:2H]
21974 ? S 0:00 [gfs2_logd]
21975 ? S 0:00 [gfs2_quotad]
27282 ? Ss 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
27285 ? S 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
27287 ? S 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
27289 ? S 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
27291 ? S 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
27292 ? S 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
27390 ? S 0:00 /sbin/httpd -DSTATUS -f /etc/httpd/conf/httpd.conf -c
29177 pts/0 R+ 0:00 ps ax
[root@lady3jane ~]# echo c > /proc/sysrq-trigger
```


[illegible]

```
root@maelcum:~  
File Edit View Search Terminal Help  
ation message from PID 30553, but reception only permitted for PID 1140  
Mar 24 21:52:07 maelcum.tc.redhat.com apache(summit-apache)[30566]: INFO: waitin  
g for apache /etc/httpd/conf/httpd.conf to come up  
Mar 24 21:52:08 maelcum.tc.redhat.com dlm_control[2542]: 13896 fence status 2 r  
eceive 0 from 1 walltime 1395715928 local 13896  
Mar 24 21:52:08 maelcum.tc.redhat.com kernel: GFS2: fsid=summit:gfs0.2: recover  
generation 3 jid 0  
Mar 24 21:52:08 maelcum.tc.redhat.com kernel: GFS2: fsid=summit:gfs0.2: jid=0: T  
rying to acquire journal lock...  
Mar 24 21:52:08 maelcum.tc.redhat.com kernel: GFS2: fsid=summit:gfs0.2: jid=0: B  
usy  
Mar 24 21:52:08 maelcum.tc.redhat.com kernel: GFS2: fsid=summit:gfs0.2: recover  
jid 0 result busy  
Mar 24 21:52:09 maelcum.tc.redhat.com apache(summit-apache)[30584]: INFO: Succes  
sfully retrieved http header at http://localhost:80  
Mar 24 21:52:09 maelcum.tc.redhat.com crmd[1150]: notice: process_lrm_event: LRM  
operation summit-apache_start_0 (call=77, rc=0, cib-update=61, confirmed=true)  
ok  
Mar 24 21:52:09 maelcum.tc.redhat.com apache(summit-apache)[30651]: INFO: Succes  
sfully retrieved http header at http://localhost:80  
Mar 24 21:52:09 maelcum.tc.redhat.com crmd[1150]: notice: process_lrm_event: LRM  
operation summit-apache_monitor_10000 (call=78, rc=0, cib-update=62, confirmed=  
false) ok  
Mar 24 21:52:09 maelcum.tc.redhat.com kernel: GFS2: fsid=summit:gfs0.2: recover
```

```
root@maelcum:~  
File Edit View Search Terminal Help  
synchronization, ready to provide service.  
Mar 24 21:52:47 maelcum.tc.redhat.com crmd[1150]: notice: crm_update_peer_state:  
pcmk_quorum_notification: Node lady3jane.tc.redhat.com[2] - state is now member  
(was lost)  
Mar 24 21:52:47 maelcum.tc.redhat.com pacemakerd[1140]: notice: crm_update_peer_  
state: pcmk_quorum_notification: Node lady3jane.tc.redhat.com[2] - state is now  
member (was lost)  
Mar 24 21:52:50 maelcum.tc.redhat.com apache(summit-apache)[31177]: INFO: Succes  
sfully retrieved http header at http://localhost:80  
Mar 24 21:52:54 maelcum.tc.redhat.com apache(summit-apache)[31248]: INFO: Attemp  
ting graceful stop of apache PID 30553  
Mar 24 21:52:54 maelcum.tc.redhat.com kernel: dlm: got connection from 2  
Mar 24 21:52:56 maelcum.tc.redhat.com apache(summit-apache)[31298]: INFO: apache  
stopped.  
Mar 24 21:52:56 maelcum.tc.redhat.com crmd[1150]: notice: process_lrm_event: LRM  
operation summit-apache_stop_0 (call=80, rc=0, cib-update=63, confirmed=true) o  
k  
Mar 24 21:52:56 maelcum.tc.redhat.com IPAddr2(summit-ip)[31347]: INFO: IP status  
= ok, IP_CIP=  
Mar 24 21:52:56 maelcum.tc.redhat.com avahi-daemon[737]: Withdrawing address rec  
ord for 172.31.100.250 on pl4p1.  
Mar 24 21:52:56 maelcum.tc.redhat.com crmd[1150]: notice: process_lrm_event: LRM  
operation summit-ip_stop_0 (call=82, rc=0, cib-update=64, confirmed=true) ok  
█
```




10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

NOW GO FORTH AND CLUSTER!

QUESTIONS?



THANK YOU!

- If you liked today's session, please leave feedback!
- Slides available at <http://people.redhat.com/tcameron> and on the Summit 2014 web site!



10 YEARS *and counting*
SAN FRANCISCO | APRIL 14-17, 2014

DIVIDER SLIDE TITLE GOES HERE

Divider subtitle goes here

SLIDE TITLE GOES HERE

Text without bullets

- Bullets layer one
 - Bullets layer two
 - Bullets layer three