

Pacemaker 1.1

Clusters from Scratch

Creare cluster Active/Passive e Active/Active su Fedora



Andrew Beekhof

Pacemaker 1.1 Clusters from Scratch

Creare cluster Active/Passive e Active/Active su Fedora

Edizione 5

Autore	Andrew Beekhof	andrew@beekhof.net
Traduttore	Raoul Scarazzini	rasca@miamammauslinux.org
Traduttore	Dan Frîncu	df.cluster@gmail.com

Copyright © 2009-2012 Andrew Beekhof.

The text of and illustrations in this document are licensed under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA")¹.

In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

In addition to the requirements of this license, the following activities are looked upon favorably:

1. If you are distributing Open Publication works on hardcopy or CD-ROM, you provide email notification to the authors of your intent to redistribute at least thirty days before your manuscript or media freeze, to give the authors time to provide updated documents. This notification should describe modifications, if any, made to the document.
2. All substantive modifications (including deletions) be either clearly marked up in the document or else described in an attachment to the document.
3. Finally, while it is not mandatory under this license, it is considered good form to offer a free copy of any hardcopy or CD-ROM expression of the author(s) work.

Lo scopo di questo documento è di fornire una guida completa per costruire un cluster active/passive con Pacemaker e mostrare come può essere convertito in una configurazione active/active.

Il cluster userà:

1. Fedora 17 come sistema operativo
2. Corosync per fornire i servizi di messaging e membership,
3. Pacemaker per la gestione delle risorse
4. DRBD come alternativa prezzo/prestazioni allo storage condiviso,
5. GFS2 come cluster filesystem (nella modalità active/active)

Per via del processo grafico di installazione di Fedora, diversi screenshot sono inclusi. Ad ogni modo questa guida è composta primariamente dai comandi, dalle ragioni per cui questi vengono eseguiti e l'output da loro prodotto.

¹ An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>

Indice

Prefazione	ix
1. Convenzioni del documento	ix
1.1. Convenzioni tipografiche	ix
1.2. Convenzioni del documento	xi
1.3. Note ed avvertimenti	xi
2. We Need Feedback!	xii
1. Leggimi-Prima	1
1.1. Scopo di questo documento	1
1.2. Cos'è Pacemaker?	1
1.3. Architettura di Pacemaker	2
1.3.1. Componenti interni	3
1.4. Tipologia dei cluster Pacemaker	4
2. Installazione	7
2.1. Installazione del sistema operativo	7
2.2. Post Installation Tasks	9
2.2.1. Networking	9
2.2.2. Leaving the Console	10
2.2.3. Considerazioni sulla sicurezza	11
2.2.4. Nomi di nodi abbreviati	11
2.2.5. NTP	12
2.3. Prima di continuare	12
2.3.1. Definire la rete	12
2.3.2. Configurazione di SSH	13
2.4. Installazione del software cluster	14
2.4.1. Installazione del software cluster	14
2.5. Setup	21
2.5.1. Preparation - Multicast	21
2.5.2. Notes on Multicast Address Assignment	22
2.5.3. Configurare Corosync	23
2.5.4. Propagate the Configuration	24
3. Pacemaker Tools	25
3.1. Utilizzare i tool di Pacemaker	25
4. Verifica dell'installazione del cluster	29
4.1. Verifica dell'installazione di Corosync	29
4.2. Verifica dell'installazione di Pacemaker	30
5. Creare un cluster Active/Passive	33
5.1. Esplorare la configurazione esistente	33
5.2. Aggiungere una risorsa	34
5.3. Effettuare un Failover	36
5.3.1. Quorum e Cluster a due nodi	36
5.3.2. Evitare che le risorse si muovano dopo il recovery	38
6. Apache - Aggiungere ulteriori servizi	39
6.1. Forward	39
6.2. Installazione	39
6.3. Preparazione	40
6.4. Enable the Apache status URL	41
6.5. Aggiornare la configurazione	41
6.6. Assicurarsi che le risorse funzionino sullo stesso host	42

6.7. Controllare l'ordinamento di start e stop di una risorsa	43
6.8. Specificare una Location preferita	44
6.9. Spostamento manuale delle risorse all'interno del cluster	45
6.9.1. Restituire il controllo nuovamente al cluster	46
7. Storage replicato con DRBD	49
7.1. Background	49
7.2. Installare i pacchetti DRBD	49
7.3. Configurare DRBD	50
7.3.1. Creare una partizione per DRBD	50
7.3.2. Scrivere la configurazione DRBD	51
7.3.3. Inizializzare e caricare DRBD	51
7.3.4. Popolare DRBD con i dati	52
7.4. Configurare il cluster per DRBD	53
7.4.1. Testare la migrazione	56
8. Conversione in Active/Active	59
8.1. Requisiti	59
8.1.1. Installing the required Software	59
8.2. Creazione un Filesystem GFS2	60
8.2.1. Preparazione	60
8.2.2. Creazione e popolamento della partizione GFS2	60
8.3. Riconfigurare il cluster per GFS2	63
8.4. Riconfigurare Pacemaker per l'Active/Active	64
8.4.1. Testare il recovery	67
9. Configurare STONITH	69
9.1. What Is STONITH	69
9.2. Quale device STONITH andrebbe utilizzato	69
9.3. Configurare STONITH	69
9.4. Esempio	70
A. Riepilogo delle configurazioni	73
A.1. Configurazione finale del cluster	73
A.2. Lista nodi	74
A.3. Opzioni del cluster	74
A.4. Risorse	74
A.4.1. Opzioni di default	74
A.4.2. Fencing	74
A.4.3. Servizio Address	74
A.4.4. DRBD - Storage condiviso	75
A.4.5. Cluster Filesystem	75
A.4.6. Apache	75
B. Sample Corosync Configuration	77
C. Approfondimenti	79
D. Revision History	81
Indice analitico	83

Lista delle figure

1.1. Panoramica concettuale dello Stack	2
1.2. Lo stack Pacemaker	3
1.3. Componenti interni	3
1.4. Ridondanza Active/Passive	5
1.5. Ridondanza N a N	5

Lista degli esempi

5.1. Questo è l'ultimo XML ad essere utilizzato nel documento.	33
---	----

Prefazione

Indice

1. Convenzioni del documento	ix
1.1. Convenzioni tipografiche	ix
1.2. Convenzioni del documento	xi
1.3. Note ed avvertimenti	xi
2. We Need Feedback!	xii

1. Convenzioni del documento

Questo manuale utilizza numerose convenzioni per evidenziare parole e frasi, ponendo attenzione su informazioni specifiche.

Nelle edizioni PDF e cartacea questo manuale utilizza caratteri presenti nel set [Font Liberation](https://fedorahosted.org/liberation-fonts/)¹. Il set Font Liberation viene anche utilizzato nelle edizioni HTML se il set stesso è stato installato sul vostro sistema. In caso contrario, verranno mostrati caratteri alternativi ma equivalenti. Da notare: Red Hat Enterprise Linux 5 e versioni più recenti, includono per default il set Font Liberation.

1.1. Convenzioni tipografiche

Vengono utilizzate quattro convenzioni tipografiche per richiamare l'attenzione su parole e frasi specifiche. Queste convenzioni, e le circostanze alle quali vengono applicate, sono le seguenti.

Neretto monospazio

Usato per evidenziare l'input del sistema, incluso i comandi della shell, i nomi dei file ed i percorsi. Utilizzato anche per evidenziare tasti e combinazione di tasti. Per esempio:

Per visualizzare i contenuti del file **my_next_bestselling_novel** nella vostra directory di lavoro corrente, inserire il comando **cat my_next_bestselling_novel** al prompt della shell e premere **Invio** per eseguire il comando.

Quanto sopra riportato include il nome del file, un comando della shell ed un tasto, il tutto riportato in neretto monospazio e distinguibile grazie al contesto.

Le combinazioni si distinguono dai tasti singoli tramite l'uso del segno più, il quale viene usato per creare una combinazione di tasti. Per esempio:

Premere **Invio** per eseguire il comando.

Premere **Ctrl+Alt+F2** per usare un terminale virtuale.

Il primo esempio evidenzia il tasto specifico singolo da premere. Il secondo riporta una combinazione di tasti: un insieme di tre tasti premuti contemporaneamente.

Se si discute del codice sorgente, i nomi della classe, i metodi, le funzioni i nomi della variabile ed i valori ritornati indicati all'interno di un paragrafo, essi verranno indicati come sopra, e cioè in **neretto monospazio**. Per esempio:

¹ <https://fedorahosted.org/liberation-fonts/>

Le classi relative ad un file includono **filesystem** per file system, **file** per file, e **dir** per directory. Ogni classe possiede il proprio set associato di permessi.

Proportional Bold

Ciò denota le parole e le frasi incontrate su di un sistema, incluso i nomi delle applicazioni; il testo delle caselle di dialogo; i pulsanti etichettati; le caselle e le etichette per pulsanti di selezione, titoli del menu e dei sottomenu. Per esempio:

Selezionare **Sistema** → **Preferenze** → **Mouse** dalla barra del menu principale per lanciare **Preferenze del Mouse**. Nella scheda **Pulsanti**, fate clic sulla casella di dialogo **mouse per mancini**, e successivamente fate clic su **Chiudi** per cambiare il pulsante primario del mouse da sinistra a destra (rendendo così il mouse idoneo per un utilizzo con la mano sinistra).

Per inserire un carattere speciale in un file **gedit** selezionare **Applicazioni** → **Accessori** → **Mappa del carattere** dalla barra del menu principale. Selezionare successivamente **Cerca** → **Trova...** dal menu **Mappa del carattere**, digitare il nome desiderato nel campo **Cerca** e selezionare **Successivo**. Il carattere desiderato sarà evidenziato nella **Tabella dei caratteri**. Eseguire un doppio clic sul carattere per poterlo posizionare nel campo **Testo da copiare** e successivamente fare clic sul pulsante **Copia**. Ritornare sul documento e selezionare **Modifica** → **Incolla** dalla barra del menu di **gedit**.

Il testo sopra riportato include i nomi delle applicazioni; nomi ed oggetti del menu per l'intero sistema; nomi del menu specifici alle applicazioni; e pulsanti e testo trovati all'interno di una interfaccia GUI, tutti presentati in neretto proporzionale e distinguibili dal contesto.

Corsivo neretto monospazio o Corsivo neretto proporzionale

Sia se si tratta di neretto monospazio o neretto proporzionale, l'aggiunta del carattere corsivo indica un testo variabile o sostituibile. Il carattere corsivo denota un testo che non viene inserito letteralmente, o visualizzato che varia a seconda delle circostanze. Per esempio:

Per collegarsi ad una macchina remota utilizzando ssh, digitare **ssh** **username@domain.name** al prompt della shell. Se la macchina remota è **example.com** ed il nome utente sulla macchina interessata è john, digitare **ssh** **john@example.com**.

Il comando **mount -o remount file-system** rimonta il file system indicato. Per esempio, per rimontare il file system **/home**, il comando è **mount -o remount /home**.

Per visualizzare la versione di un pacchetto attualmente installato, utilizzare il comando **rpm -q package**. Esso ritornerà il seguente risultato: **package-version-release**.

Da notare le parole in corsivo grassetto - username, domain.name, file-system, package, version e release. Ogni parola funge da segnaposto, sia esso un testo inserito per emettere un comando o mostrato dal sistema.

Oltre all'utilizzo normale per la presentazione di un titolo, il carattere Corsivo denota il primo utilizzo di un termine nuovo ed importante. Per esempio:

Publican è un sistema di pubblicazione per *DocBook*.

1.2. Convenzioni del documento

Gli elenchi originati dal codice sorgente e l'output del terminale vengono evidenziati rispetto al testo circostante.

L'output inviato ad un terminale è impostato su **tondo monospazio** e così presentato:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Gli elenchi del codice sorgente sono impostati in **tondo monospazio** ma vengono presentati ed evidenziati nel modo seguente:

```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo             echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

1.3. Note ed avvertimenti

E per finire, tre stili vengono usati per richiamare l'attenzione su informazioni che in caso contrario potrebbero essere ignorate.



Nota

Una nota è un suggerimento o un approccio alternativo per il compito da svolgere. Non dovrebbe verificarsi alcuna conseguenza negativa se la nota viene ignorata, ma al tempo stesso potreste non usufruire di qualche trucco in grado di facilitarvi il compito.



Importante

Le caselle 'importante' riportano informazioni che potrebbero passare facilmente inosservate: modifiche alla configurazione applicabili solo alla sessione corrente, o servizi i quali necessitano di un riavvio prima di applicare un aggiornamento. Ignorare queste caselle non causa alcuna perdita di dati ma potrebbe causare irritazione e frustrazione da parte dell'utente.



Avvertimento

Un Avvertimento non dovrebbe essere ignorato. Se ignorato, potrebbe verificarsi una perdita di dati.

2. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla² against the product **Pacemaker**.

When submitting a bug report, be sure to mention the manual's identifier: *Clusters_from_Scratch*

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

² <http://bugs.clusterlabs.org>

Leggimi-Prima

Indice

1.1. Scopo di questo documento	1
1.2. Cos'è Pacemaker?	1
1.3. Architettura di Pacemaker	2
1.3.1. Componenti interni	3
1.4. Tipologia dei cluster Pacemaker	4

1.1. Scopo di questo documento

Computer clusters can be used to provide highly available services or resources. The redundancy of multiple machines is used to guard against failures of many types.

This document will walk through the installation and setup of simple clusters using the Fedora distribution, version 17.

The clusters described here will use Pacemaker and Corosync to provide resource management and messaging. Required packages and modifications to their configuration files are described along with the use of the Pacemaker command line tool for generating the XML used for cluster control.

Pacemaker is a central component and provides the resource management required in these systems. This management includes detecting and recovering from the failure of various nodes, resources and services under its control.

When more in depth information is required and for real world usage, please refer to the [Pacemaker Explained](#)¹ manual.

1.2. Cos'è Pacemaker?

Pacemaker is a cluster resource manager. It achieves maximum availability for your cluster services (aka. resources) by detecting and recovering from node and resource-level failures by making use of the messaging and membership capabilities provided by your preferred cluster infrastructure (either Corosync or Heartbeat).

Pacemaker's key features include:

- Rilevazione e ripristino di malfunzionamenti di nodi e servizi
- Storage agnostic, non richiede uno storage condiviso
- Resource agnostic, tutto quello che può essere scriptato può essere clusterizzato
- Supports STONITH for ensuring data integrity
- Supporto a cluster grandi e piccoli
- Supports both quorate and resource driven clusters
- Supports practically any redundancy configuration

¹ <http://www.clusterlabs.org/doc/>

- Configurazione replicata automaticamente che può essere aggiornata da qualsiasi nodo
- Capacità di specificare ordine, collocazione e anti-collocazione per i servizi lato cluster
- Support for advanced service types
 - Cloni: per servizi che necessitano di essere attivi su nodi multipli
 - Multi-state: per servizi con modi multipli (ad esempio master/slave, primary/secondary/)
- Unified, scriptable, cluster management tools.

1.3. Architettura di Pacemaker

Al livello più elevato il cluster è composto da tre componenti:

- Non-cluster aware components (illustrated in green). These pieces include the resources themselves, scripts that start, stop and monitor them, and also a local daemon that masks the differences between the different standards these scripts implement.
- 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; and other administrative actions. 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.
- Low level infrastructure Corosync provides reliable messaging, membership and quorum information about the cluster (illustrated in red).



Figura 1.1. Panoramica concettuale dello Stack

When combined with Corosync, Pacemaker also supports popular open source cluster filesystems.²

² Even though Pacemaker also supports Heartbeat, the filesystems need to use the stack for messaging and membership and Corosync seems to be what they're standardizing on. Technically it would be possible for them to support Heartbeat as well, however there seems little interest in this.

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.




Figura 1.2. Lo stack Pacemaker

1.3.1. Componenti interni

Pacemaker stesso è composto da quattro componenti chiave (illustrati sotto nello stesso schema di colori del diagramma precedente):

- CIB (acronimo di come Cluster Information Base)
- CRMD (acronimo di Cluster Resource Management daemon)
- PEngine (acronimo di Policy Engine)
- STONITHd




Figura 1.3. Componenti interni

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.

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).

Gli altri nodi riferiscono i risultati delle loro operazioni al DC. Attraverso l'analisi dei risultati aspettati e di quelli attuali, i nodi eseguiranno qualsiasi azione necessaria per attendere il completamento della precedente oppure interromperanno il processo, richiedendo al PEngine di calcolare nuovamente lo stato ideale del cluster basandosi sui risultati inaspettati.

In alcuni casi, potrebbe essere necessario spegnere i nodi per preservare dati condivisi o completare il ripristino di una risorsa. Per questo in Pacemaker esiste STONITHd. STONITH è un acronimo per Shoot-The-Other-Node-In-The-Head e viene implementato tipicamente con un switch di potenza remoto. In Pacemaker i dispositivi STONITH sono modellati come risorse (e configurati all'interno del CIB) per facilitare il monitoraggio delle anomalie. STONITHd si prende cura di capire la topologia STONITH così che i suoi client debbano unicamente richiedere unicamente la morte di un nodo ed esso si preoccupi del resto.

1.4. Tipologia dei cluster Pacemaker

Pacemaker non fa alcuna ipotesi in merito all'ambiente operativo, questo consente di supportare praticamente qualsiasi *configurazione ridondata*³ come Active/Active, Active/Passive, N+1, N+M, N-to-1 e N-to-N.

In this document we will focus on the setup of a highly available Apache web server with an Active/Passive cluster using DRBD and Ext4 to store data. Then, we will upgrade this cluster to Active/Active using GFS2.

³ http://en.wikipedia.org/wiki/High-availability_cluster#Node_configurations




Figura 1.4. Ridondanza Active/Passive




Figura 1.5. Ridondanza N a N

Installazione

Indice

2.1. Installazione del sistema operativo	7
2.2. Post Installation Tasks	9
2.2.1. Networking	9
2.2.2. Leaving the Console	10
2.2.3. Considerazioni sulla sicurezza	11
2.2.4. Nomi di nodi abbreviati	11
2.2.5. NTP	12
2.3. Prima di continuare	12
2.3.1. Definire la rete	12
2.3.2. Configurazione di SSH	13
2.4. Installazione del software cluster	14
2.4.1. Installazione del software cluster	14
2.5. Setup	21
2.5.1. Preparation - Multicast	21
2.5.2. Notes on Multicast Address Assignment	22
2.5.3. Configurare Corosync	23
2.5.4. Propagate the Configuration	24

2.1. Installazione del sistema operativo

Detailed instructions for installing Fedora are available at http://docs.fedoraproject.org/en-US/Fedora/17/html/Installation_Guide/ in a number of languages. The abbreviated version is as follows...

Point your browser to <http://fedoraproject.org/en/get-fedora-all>, locate the **Install Media** section and download the install DVD that matches your hardware.

Burn the disk image to a DVD ¹ and boot from it, or use the image to boot a virtual machine.

After clicking through the welcome screen, select your language, keyboard layout ² and storage type ³

Assign your machine a host name. ⁴ I happen to control the clusterlabs.org domain name, so I will use that here.

¹ http://docs.fedoraproject.org/en-US/Fedora/16/html/Burning_ISO_images_to_disc/index.html

² http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/sn-keyboard-x86.html

³ http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/Storage_Devices-x86.html


⁴ http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/sn-Netconfig-x86.html



Importante

Do not accept the default network settings. Cluster machines should never obtain an IP address via DHCP.

When you are presented with the **Configure Network** advanced option, select that option before continuing with the installation process to specify a fixed IPv4 address for **System eth0**. Be sure to also enter the **Routes** section and add an entry for your default gateway.



If you miss this step, this can easily be configured after installation. You will have to navigate to **system settings** and select **network**. From there you can select what device to configure.

You will then be prompted to indicate the machine's physical location⁵ and to supply a root password.⁶

Now select where you want Fedora installed.⁷ As I don't care about any existing data, I will accept the default and allow Fedora to use the complete drive.

⁵ http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/s1-timezone-x86.html

⁶ http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/sn-account_configuration-x86.html

⁷ http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/s1-diskpartsetup-x86.html



Importante

By default Fedora uses LVM for partitioning which allows us to dynamically change the amount of space allocated to a given partition.

However, by default it also allocates all free space to the / (aka. **root**) partition which cannot be dynamically *reduced* in size (dynamic increases are fine by-the-way).

So if you plan on following the DRBD or GFS2 portions of this guide, you should reserve at least 1Gb of space on each machine from which to create a shared volume. To do so select the **Review and modify partitioning layout** checkbox before clicking **Next**. You will then be given an opportunity to reduce the size of the **root** partition.

Next choose which software should be installed.⁸ Change the selection to Minimal so that we see everything that gets installed. Don't enable updates yet, we'll do that (and install any extra software we need) later. After you click next, Fedora will begin installing.

Go grab something to drink, this may take a while.

Once the node reboots, you'll see a (possibly mangled) login prompt on the console. Login using **root** and the password you created earlier.

```
Fedora reStarted Sendmail Mail Transport Agent
[ OK ]
lease 17 (BeefyStarting Sendmail Mail Transport Client...

Miracle)
Kernel 3.3.0-0.rc3.git7.2.fc17.x86_64 on an x86_64 (tty1)

pcmk-1 login: Started Sendmail Mail Transport Client
[ OK ]
Started SSH server keys generation. [ OK ]
Starting OpenSSH server daemon... [ OK ]
Started OpenSSH server daemon
-
```



Nota

From here on in we're going to be working exclusively from the terminal.

2.2. Post Installation Tasks

2.2.1. Networking

Bring up the network and ensure it starts at boot

```
# service network start
```

⁸ http://docs.fedoraproject.org/en-US/Fedora/16/html/Installation_Guide/s1-pkgselection-x86.html

Capitolo 2. Installazione

```
# chkconfig network on
```

Check the machine has the static IP address you configured earlier

```
# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 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
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 52:54:00:d7:d6:08 brd ff:ff:ff:ff:ff:ff
    inet 192.168.122.101/24 scope global eth0
    inet6 fe80::5054:ff:fed7:d608/64 scope link
        valid_lft forever preferred_lft forever
```

Now check the default route setting:

```
[root@pcmk-1 ~]# ip route
default via 192.168.122.1 dev eth0
192.168.122.0/24 dev eth0 proto kernel scope link src 192.168.122.101
```

If there is no line beginning with **default via**, then you may need to add a line such as

```
GATEWAY=192.168.122.1
```

to `/etc/sysconfig/network` and restart the network.

Now check for connectivity to the outside world. Start small by testing if we can reach the gateway we configured.

```
# ping -c 1 192.168.122.1
PING 192.168.122.1 (192.168.122.1) 56(84) bytes of data.
64 bytes from 192.168.122.1: icmp_req=1 ttl=64 time=0.249 ms

--- 192.168.122.1 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.249/0.249/0.249/0.000 ms
```

Now try something external, choose a location you know will be available.

```
# ping -c 1 www.google.com
PING www.l.google.com (173.194.72.106) 56(84) bytes of data.
64 bytes from tf-in-f106.1e100.net (173.194.72.106): icmp_req=1 ttl=41 time=167 ms

--- www.l.google.com ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 167.618/167.618/167.618/0.000 ms
```

2.2.2. Leaving the Console

The console isn't a very friendly place to work from, we will now switch to accessing the machine remotely via SSH where we can use copy&paste etc.

First we check we can see the newly installed at all:

```
beekhof@f16 ~ # ping -c 1 192.168.122.101
PING 192.168.122.101 (192.168.122.101) 56(84) bytes of data.
64 bytes from 192.168.122.101: icmp_req=1 ttl=64 time=1.01 ms
```

```
--- 192.168.122.101 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 1.012/1.012/1.012/0.000 ms
```

Next we login via SSH

```
beekhof@f16 ~ # ssh -l root 192.168.122.11
root@192.168.122.11's password:
Last login: Fri Mar 30 19:41:19 2012 from 192.168.122.1
[root@pcmk-1 ~]#
```

2.2.3. Considerazioni sulla sicurezza

To simplify this guide and focus on the aspects directly connected to clustering, we will now disable the machine's firewall and SELinux installation.



Avvertimento

Both of these actions create significant security issues and should not be performed on machines that will be exposed to the outside world.



Importante

TODO: Creare un Appendice che valuti (almeno) la riabilitazione del firewall.

```
# setenforce 0
# sed -i.bak "s/SELINUX=enforcing/SELINUX=permissive/g" /etc/selinux/config
# systemctl disable iptables.service
# rm '/etc/systemd/system/basic.target.wants/iptables.service'
# systemctl stop iptables.service
```

2.2.4. Nomi di nodi abbreviati

During installation, we filled in the machine's fully qualified domain name (FQDN) which can be rather long when it appears in cluster logs and status output. See for yourself how the machine identifies itself:

```
# uname -n
pcmk-1.clusterlabs.org
# dnsdomainname
clusterlabs.org
```

L'output del secondo comando è corretto, ma non è necessario il nome del dominio incluso nei dettagli essenziali dell'host. Per risolvere la situazione è necessario modificare il file `/etc/sysconfig/network`. Ecco come dovrà apparire.

```
# cat /etc/sysconfig/network
NETWORKING=yes
HOSTNAME=pcmk-1.clusterlabs.org
GATEWAY=192.168.122.1
```

L'unica cosa da fare sarà di rimuovere la parte relativa al nome del dominio, che rimarrà comunque registrata altrove.

```
# sed -i.sed 's/\.[a-z].*//g' /etc/sysconfig/network
```

Per verificare l'efficacia dei cambiamenti apportati il contenuto del file dovrà essere simile a quanto riportato.

```
# cat /etc/sysconfig/network
NETWORKING=yes
HOSTNAME=pcmk-1
GATEWAY=192.168.122.1
```

Ma non è ancora finita. La macchina non sarà allineata con i nomi abbreviati finché non verrà riavviata, ma è possibile forzare l'aggiornamento.

```
# source /etc/sysconfig/network
# hostname $HOSTNAME
```

E' possibile quindi verificare che la macchina utilizzi il nome corretto

```
# uname -n
pcmk-1
# dnsdomainname
clusterlabs.org
```

2.2.5. NTP

It is highly recommended to enable NTP on your cluster nodes. Doing so ensures all nodes agree on the current time and makes reading log files significantly easier.⁹

2.3. Prima di continuare

Repeat the Installation steps so far, so that you have two Fedora nodes ready to have the cluster software installed.

For the purposes of this document, the additional node is called pcmk-2 with address 192.168.122.102.

2.3.1. Definire la rete

Confirm that you can communicate between the two new nodes:

```
# ping -c 3 192.168.122.102
PING 192.168.122.102 (192.168.122.102) 56(84) bytes of data.
```

⁹ http://docs.fedoraproject.org/en-US/Fedora/17/html-single/System_Administrators_Guide/index.html#ch-Configuring_the_Date_and_Time


```
64 bytes from 192.168.122.102: icmp_seq=1 ttl=64 time=0.343 ms
64 bytes from 192.168.122.102: icmp_seq=2 ttl=64 time=0.402 ms
64 bytes from 192.168.122.102: icmp_seq=3 ttl=64 time=0.558 ms

--- 192.168.122.102 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2000ms
rtt min/avg/max/mdev = 0.343/0.434/0.558/0.092 ms
```

Now we need to make sure we can communicate with the machines by their name. If you have a DNS server, add additional entries for the two machines. Otherwise, you'll need to add the machines to `/etc/hosts`. Below are the entries for my cluster nodes:

```
# grep pcmk /etc/hosts
192.168.122.101 pcmk-1.clusterlabs.org pcmk-1
192.168.122.102 pcmk-2.clusterlabs.org pcmk-2
```

Ora è possibile verificare il setup usando nuovamente ping:

```
# ping -c 3 pcmk-2
PING pcmk-2.clusterlabs.org (192.168.122.101) 56(84) bytes of data.
64 bytes from pcmk-1.clusterlabs.org (192.168.122.101): icmp_seq=1 ttl=64 time=0.164 ms
64 bytes from pcmk-1.clusterlabs.org (192.168.122.101): icmp_seq=2 ttl=64 time=0.475 ms
64 bytes from pcmk-1.clusterlabs.org (192.168.122.101): icmp_seq=3 ttl=64 time=0.186 ms

--- pcmk-2.clusterlabs.org ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2001ms
rtt min/avg/max/mdev = 0.164/0.275/0.475/0.141 ms
```

2.3.2. Configurazione di SSH

SSH is a convenient and secure way to copy files and perform commands remotely. For the purposes of this guide, we will create a key without a password (using the `-N` option) so that we can perform remote actions without being prompted.



Avvertimento

Unprotected SSH keys, those without a password, are not recommended for servers exposed to the outside world. We use them here only to simplify the demo.

Creazione di una nuova chiave ed abilitazione al proprietario della chiave di effettuare login:

Creazione ed attivazione di una nuova chiave SSH

```
# ssh-keygen -t dsa -f ~/.ssh/id_dsa -N ""
Generating public/private dsa key pair.
Your identification has been saved in /root/.ssh/id_dsa.
Your public key has been saved in /root/.ssh/id_dsa.pub.
The key fingerprint is:
91:09:5c:82:5a:6a:50:08:4e:b2:0c:62:de:cc:74:44 root@pcmk-1.clusterlabs.org

The key's randomart image is:
+--[ DSA 1024]-----+
|==.ooEo..          |
```

```
|X 0 + .o o |
| * A + |
| + . |
| . S |
| |
| |
| |
+-----+

# cp .ssh/id_dsa.pub .ssh/authorized_keys
```

Installare la chiave sull'altro nodo e verificare come sia possibile lanciare comandi remotamente, senza ricevere prompt

Installare la chiave SSH su un altro host

```
# scp -r .ssh pcmk-2:
The authenticity of host 'pcmk-2 (192.168.122.102)' can't be established.
RSA key fingerprint is b1:2b:55:93:f1:d9:52:2b:0f:f2:8a:4e:ae:c6:7c:9a.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'pcmk-2,192.168.122.102' (RSA) to the list of known
hosts.root@pcmk-2's password:
id_dsa.pub                100% 616      0.6KB/s   00:00
id_dsa                    100% 672      0.7KB/s   00:00
known_hosts               100% 400      0.4KB/s   00:00
authorized_keys           100% 616      0.6KB/s   00:00
# ssh pcmk-2 -- uname -n
pcmk-2
#
```

2.4. Installazione del software cluster

2.4.1. Installazione del software cluster

Since version 12, Fedora comes with recent versions of everything you need, so simply fire up a shell on all your nodes and run:

```
[ALL] # yum install -y pacemaker corosync
```

```
fedora/metalink           | 38 kB    00:00
fedora                    | 4.2 kB    00:00
fedora/primary_db         | 14 MB    00:21
updates/metalink          | 2.7 kB    00:00
updates                   | 2.6 kB    00:00
updates/primary_db        | 1.2 kB    00:00
updates-testing/metalink  | 28 kB    00:00
updates-testing           | 4.5 kB    00:00
updates-testing/primary_db | 4.5 MB    00:12
Setting up Install Process
Resolving Dependencies
--> Running transaction check
---> Package corosync.x86_64 0:1.99.9-1.fc17 will be installed
--> Processing Dependency: corosynclib = 1.99.9-1.fc17 for package:
    corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libxslt for package: corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libvotequorum.so.5(COROSYNC_VOTEQUORUM_1.0)(64bit) for package:
    corosync-1.99.9-1.fc17.x86_64
```

```
--> Processing Dependency: libquorum.so.5(COROSYNC_QUORUM_1.0)(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcpq.so.4(COROSYNC_CPG_1.0)(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcmmap.so.4(COROSYNC_CMAP_1.0)(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcfg.so.6(COROSYNC_CFG_0.82)(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libvotequorum.so.5()(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libtotem_pg.so.5()(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libquorum.so.5()(64bit) for package: corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libqb.so.0()(64bit) for package: corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libnetsnmp.so.30()(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcpq.so.4()(64bit) for package: corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcorosync_common.so.4()(64bit) for package:
corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcmmap.so.4()(64bit) for package: corosync-1.99.9-1.fc17.x86_64
--> Processing Dependency: libcfg.so.6()(64bit) for package: corosync-1.99.9-1.fc17.x86_64
---> Package pacemaker.x86_64 0:1.1.7-2.fc17 will be installed
--> Processing Dependency: pacemaker-libs = 1.1.7-2.fc17 for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: pacemaker-cluster-libs = 1.1.7-2.fc17 for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: pacemaker-cli = 1.1.7-2.fc17 for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: resource-agents for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: perl(Getopt::Long) for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libgnutls.so.26(GNUTLS_1_4)(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: cluster-glue for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: /usr/bin/perl for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libtransitioner.so.1()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libstonithd.so.1()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libstonith.so.1()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libplumb.so.2()(64bit) for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libpils.so.2()(64bit) for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libpengine.so.3()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libpe_status.so.3()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libpe_rules.so.2()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libltdl.so.7()(64bit) for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: liblrm.so.2()(64bit) for package: pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libgnutls.so.26()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libcrmcommon.so.2()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libcrmcluster.so.1()(64bit) for package:
pacemaker-1.1.7-2.fc17.x86_64
--> Processing Dependency: libcib.so.1()(64bit) for package: pacemaker-1.1.7-2.fc17.x86_64
--> Running transaction check
---> Package cluster-glue.x86_64 0:1.0.6-9.fc17.1 will be installed
--> Processing Dependency: perl-TimeDate for package: cluster-glue-1.0.6-9.fc17.1.x86_64
--> Processing Dependency: libOpenIPMIutils.so.0()(64bit) for package: cluster-
glue-1.0.6-9.fc17.1.x86_64
--> Processing Dependency: libOpenIPMIposix.so.0()(64bit) for package: cluster-
glue-1.0.6-9.fc17.1.x86_64
--> Processing Dependency: libOpenIPMI.so.0()(64bit) for package: cluster-
glue-1.0.6-9.fc17.1.x86_64
---> Package cluster-glue-libs.x86_64 0:1.0.6-9.fc17.1 will be installed
```

```
--> Package corosynclib.x86_64 0:1.99.9-1.fc17 will be installed
--> Processing Dependency: librdmacm.so.1(RDMACM_1.0)(64bit) for package:
corosynclib-1.99.9-1.fc17.x86_64
--> Processing Dependency: libibverbs.so.1(IBVERBS_1.1)(64bit) for package:
corosynclib-1.99.9-1.fc17.x86_64
--> Processing Dependency: libibverbs.so.1(IBVERBS_1.0)(64bit) for package:
corosynclib-1.99.9-1.fc17.x86_64
--> Processing Dependency: librdmacm.so.1()(64bit) for package:
corosynclib-1.99.9-1.fc17.x86_64
--> Processing Dependency: libibverbs.so.1()(64bit) for package:
corosynclib-1.99.9-1.fc17.x86_64
--> Package gnutls.x86_64 0:2.12.17-1.fc17 will be installed
--> Processing Dependency: libtasn1.so.3(LIBTASN1_0_3)(64bit) for package:
gnutls-2.12.17-1.fc17.x86_64
--> Processing Dependency: libtasn1.so.3()(64bit) for package: gnutls-2.12.17-1.fc17.x86_64
--> Processing Dependency: libp11-kit.so.0()(64bit) for package: gnutls-2.12.17-1.fc17.x86_64
--> Package libqb.x86_64 0:0.11.1-1.fc17 will be installed
--> Package libtool-ltdl.x86_64 0:2.4.2-3.fc17 will be installed
--> Package libxslt.x86_64 0:1.1.26-9.fc17 will be installed
--> Package net-snmp-libs.x86_64 1:5.7.1-4.fc17 will be installed
--> Package pacemaker-cli.x86_64 0:1.1.7-2.fc17 will be installed
--> Package pacemaker-cluster-libs.x86_64 0:1.1.7-2.fc17 will be installed
--> Package pacemaker-libs.x86_64 0:1.1.7-2.fc17 will be installed
--> Package perl.x86_64 4:5.14.2-211.fc17 will be installed
--> Processing Dependency: perl-libs = 4:5.14.2-211.fc17 for package:
4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(threads::shared) >= 1.21 for package:
4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Socket) >= 1.3 for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Scalar::Util) >= 1.10 for package:
4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(File::Spec) >= 0.8 for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl-macros for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl-libs for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(threads::shared) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(threads) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Socket) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Scalar::Util) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Pod::Simple) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Module::Pluggable) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(List::Util) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(File::Spec::Unix) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(File::Spec::Functions) for package:
4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(File::Spec) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Cwd) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: perl(Carp) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Processing Dependency: libperl.so()(64bit) for package: 4:perl-5.14.2-211.fc17.x86_64
--> Package resource-agents.x86_64 0:3.9.2-2.fc17.1 will be installed
--> Processing Dependency: /usr/sbin/rpc.nfsd for package: resource-
agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /usr/sbin/rpc.mountd for package: resource-
agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /usr/sbin/ethtool for package: resource-
agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/rpc.statd for package: resource-agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/quotad for package: resource-agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/quotacheck for package: resource-
agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/mount.nfs4 for package: resource-
agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/mount.nfs for package: resource-agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/mount.cifs for package: resource-
agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: /sbin/fsck.xfs for package: resource-agents-3.9.2-2.fc17.1.x86_64
--> Processing Dependency: libnet.so.1()(64bit) for package: resource-
agents-3.9.2-2.fc17.1.x86_64
```

```

--> Running transaction check
---> Package OpenIPMI-libs.x86_64 0:2.0.18-13.fc17 will be installed
---> Package cifs-utils.x86_64 0:5.3-2.fc17 will be installed
--> Processing Dependency: libtalloc.so.2(TALLOC_2.0.2)(64bit) for package: cifs-
utils-5.3-2.fc17.x86_64
--> Processing Dependency: keyutils for package: cifs-utils-5.3-2.fc17.x86_64
--> Processing Dependency: libwbclient.so.0()(64bit) for package: cifs-
utils-5.3-2.fc17.x86_64
--> Processing Dependency: libtalloc.so.2()(64bit) for package: cifs-utils-5.3-2.fc17.x86_64
---> Package ethtool.x86_64 2:3.2-2.fc17 will be installed
---> Package libibverbs.x86_64 0:1.1.6-2.fc17 will be installed
---> Package libnet.x86_64 0:1.1.5-3.fc17 will be installed
---> Package librdmacm.x86_64 0:1.0.15-1.fc17 will be installed
---> Package libtasn1.x86_64 0:2.12-1.fc17 will be installed
---> Package nfs-utils.x86_64 1:1.2.5-12.fc17 will be installed
--> Processing Dependency: rpcbind for package: 1:nfs-utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libtirpc for package: 1:nfs-utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libnfsidmap for package: 1:nfs-utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libgssglue.so.1(libgssapi_CITI_2)(64bit) for package: 1:nfs-
utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libgssglue for package: 1:nfs-utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libevent for package: 1:nfs-utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libtirpc.so.1()(64bit) for package: 1:nfs-
utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libnfsidmap.so.0()(64bit) for package: 1:nfs-
utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libgssglue.so.1()(64bit) for package: 1:nfs-
utils-1.2.5-12.fc17.x86_64
--> Processing Dependency: libevent-2.0.so.5()(64bit) for package: 1:nfs-
utils-1.2.5-12.fc17.x86_64
---> Package p11-kit.x86_64 0:0.12-1.fc17 will be installed
---> Package perl-Carp.noarch 0:1.22-2.fc17 will be installed
---> Package perl-Module-Pluggable.noarch 1:3.90-211.fc17 will be installed
---> Package perl-PathTools.x86_64 0:3.33-211.fc17 will be installed
---> Package perl-Pod-Simple.noarch 1:3.16-211.fc17 will be installed
--> Processing Dependency: perl(Pod::Escapes) >= 1.04 for package: 1:perl-Pod-
Simple-3.16-211.fc17.noarch
---> Package perl-Scalar-List-Utils.x86_64 0:1.25-1.fc17 will be installed
---> Package perl-Socket.x86_64 0:2.001-1.fc17 will be installed
---> Package perl-TimeDate.noarch 1:1.20-6.fc17 will be installed
---> Package perl-libs.x86_64 4:5.14.2-211.fc17 will be installed
---> Package perl-macros.x86_64 4:5.14.2-211.fc17 will be installed
---> Package perl-threads.x86_64 0:1.86-2.fc17 will be installed
---> Package perl-threads-shared.x86_64 0:1.40-2.fc17 will be installed
---> Package quota.x86_64 1:4.00-3.fc17 will be installed
--> Processing Dependency: quota-nls = 1:4.00-3.fc17 for package: 1:quota-4.00-3.fc17.x86_64
--> Processing Dependency: tcp_wrappers for package: 1:quota-4.00-3.fc17.x86_64
---> Package xfsprogs.x86_64 0:3.1.8-1.fc17 will be installed
--> Running transaction check
---> Package keyutils.x86_64 0:1.5.5-2.fc17 will be installed
---> Package libevent.x86_64 0:2.0.14-2.fc17 will be installed
---> Package libgssglue.x86_64 0:0.3-1.fc17 will be installed
---> Package libnfsidmap.x86_64 0:0.25-1.fc17 will be installed
---> Package libtalloc.x86_64 0:2.0.7-4.fc17 will be installed
---> Package libtirpc.x86_64 0:0.2.2-2.1.fc17 will be installed
---> Package libwbclient.x86_64 1:3.6.3-81.fc17.1 will be installed
---> Package perl-Pod-Escapes.noarch 1:1.04-211.fc17 will be installed
---> Package quota-nls.noarch 1:4.00-3.fc17 will be installed
---> Package rpcbind.x86_64 0:0.2.0-16.fc17 will be installed
---> Package tcp_wrappers.x86_64 0:7.6-69.fc17 will be installed
--> Finished Dependency Resolution

```

Dependencies Resolved

```

=====
Package                                Arch      Version                               Repository                               Size
=====

```

```
Installing:
corosync                x86_64      1.99.9-1.fc17      updates-testing    159 k
pacemaker                x86_64      1.1.7-2.fc17      updates-testing    362 k
Installing for dependencies:
OpenIPMI-libs            x86_64      2.0.18-13.fc17     fedora              466 k
cifs-utils               x86_64      5.3-2.fc17         updates-testing     66 k
cluster-glue             x86_64      1.0.6-9.fc17.1     fedora              229 k
cluster-glue-libs        x86_64      1.0.6-9.fc17.1     fedora              121 k
corosynclib              x86_64      1.99.9-1.fc17      updates-testing     96 k
ethtool                  x86_64      2:3.2-2.fc17       fedora              94 k
gnutls                   x86_64      2.12.17-1.fc17     fedora              385 k
keyutils                  x86_64      1.5.5-2.fc17       fedora              49 k
libevent                  x86_64      2.0.14-2.fc17      fedora              160 k
libgssglue                x86_64      0.3-1.fc17         fedora              24 k
libibverbs                x86_64      1.1.6-2.fc17       fedora              44 k
libnet                    x86_64      1.1.5-3.fc17       fedora              54 k
libnfsidmap               x86_64      0.25-1.fc17        fedora              34 k
libqb                     x86_64      0.11.1-1.fc17      updates-testing     68 k
librdmacm                 x86_64      1.0.15-1.fc17      fedora              27 k
libtalloc                 x86_64      2.0.7-4.fc17       fedora              22 k
libtasn1                  x86_64      2.12-1.fc17        updates-testing     319 k
libtirpc                  x86_64      0.2.2-2.1.fc17     fedora              78 k
libtool-ltdl              x86_64      2.4.2-3.fc17       fedora              45 k
libwbclient               x86_64      1:3.6.3-81.fc17.1  updates-testing     68 k
libxslt                   x86_64      1.1.26-9.fc17      fedora              416 k
net-snmp-libs             x86_64      1:5.7.1-4.fc17     fedora              713 k
nfs-utils                 x86_64      1:1.2.5-12.fc17    fedora              311 k
p11-kit                   x86_64      0.12-1.fc17        updates-testing     36 k
pacemaker-cli             x86_64      1.1.7-2.fc17      updates-testing     368 k
pacemaker-cluster-libs    x86_64      1.1.7-2.fc17      updates-testing     77 k
pacemaker-libs            x86_64      1.1.7-2.fc17      updates-testing     322 k
perl                      x86_64      4:5.14.2-211.fc17  fedora              10 M
perl-Carp                 noarch      1.22-2.fc17        fedora              17 k
perl-Module-Pluggable     noarch      1:3.90-211.fc17    fedora              47 k
perl-PathTools            x86_64      3.33-211.fc17      fedora              105 k
perl-Pod-Escapes          noarch      1:1.04-211.fc17    fedora              40 k
perl-Pod-Simple           noarch      1:3.16-211.fc17    fedora              223 k
perl-Scalar-List-Utils    x86_64      1.25-1.fc17        updates-testing     33 k
perl-Socket               x86_64      2.001-1.fc17       updates-testing     44 k
perl-TimeDate             noarch      1:1.20-6.fc17      fedora              43 k
perl-libs                 x86_64      4:5.14.2-211.fc17  fedora              628 k
perl-macros               x86_64      4:5.14.2-211.fc17  fedora              32 k
perl-threads              x86_64      1.86-2.fc17        fedora              47 k
perl-threads-shared       x86_64      1.40-2.fc17        fedora              36 k
quota                     x86_64      1:4.00-3.fc17      fedora              160 k
quota-nls                 noarch      1:4.00-3.fc17      fedora              74 k
resource-agents           x86_64      3.9.2-2.fc17.1     fedora              466 k
rpcbind                   x86_64      0.2.0-16.fc17      fedora              52 k
tcp_wrappers              x86_64      7.6-69.fc17        fedora              72 k
xfsprogs                  x86_64      3.1.8-1.fc17       updates-testing     715 k
```

Transaction Summary

```
=====
Install 2 Packages (+46 Dependent packages)
```

Total download size: 18 M

Installed size: 59 M

Downloading Packages:

```
(1/48): OpenIPMI-libs-2.0.18-13.fc17.x86_64.rpm | 466 kB 00:00
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID 1aca3465: NOKEY
Public key for OpenIPMI-libs-2.0.18-13.fc17.x86_64.rpm is not installed
(2/48): cifs-utils-5.3-2.fc17.x86_64.rpm | 66 kB 00:01
Public key for cifs-utils-5.3-2.fc17.x86_64.rpm is not installed
(3/48): cluster-glue-1.0.6-9.fc17.1.x86_64.rpm | 229 kB 00:00
(4/48): cluster-glue-libs-1.0.6-9.fc17.1.x86_64.rpm | 121 kB 00:00
(5/48): corosync-1.99.9-1.fc17.x86_64.rpm | 159 kB 00:01
(6/48): corosynclib-1.99.9-1.fc17.x86_64.rpm | 96 kB 00:00
```

```

(7/48): ethtool-3.2-2.fc17.x86_64.rpm | 94 kB 00:00
(8/48): gnutls-2.12.17-1.fc17.x86_64.rpm | 385 kB 00:00
(9/48): keyutils-1.5.5-2.fc17.x86_64.rpm | 49 kB 00:00
(10/48): libevent-2.0.14-2.fc17.x86_64.rpm | 160 kB 00:00
(11/48): libgssglue-0.3-1.fc17.x86_64.rpm | 24 kB 00:00
(12/48): libibverbs-1.1.6-2.fc17.x86_64.rpm | 44 kB 00:00
(13/48): libnet-1.1.5-3.fc17.x86_64.rpm | 54 kB 00:00
(14/48): libnfsidmap-0.25-1.fc17.x86_64.rpm | 34 kB 00:00
(15/48): libqb-0.11.1-1.fc17.x86_64.rpm | 68 kB 00:01
(16/48): librdmacm-1.0.15-1.fc17.x86_64.rpm | 27 kB 00:00
(17/48): libtalloc-2.0.7-4.fc17.x86_64.rpm | 22 kB 00:00
(18/48): libtasn1-2.12-1.fc17.x86_64.rpm | 319 kB 00:02
(19/48): libtirpc-0.2.2-2.1.fc17.x86_64.rpm | 78 kB 00:00
(20/48): libtool-ltdl-2.4.2-3.fc17.x86_64.rpm | 45 kB 00:00
(21/48): libwbclient-3.6.3-81.fc17.1.x86_64.rpm | 68 kB 00:00
(22/48): libxslt-1.1.26-9.fc17.x86_64.rpm | 416 kB 00:00
(23/48): net-snmp-libs-5.7.1-4.fc17.x86_64.rpm | 713 kB 00:01
(24/48): nfs-utils-1.2.5-12.fc17.x86_64.rpm | 311 kB 00:00
(25/48): p11-kit-0.12-1.fc17.x86_64.rpm | 36 kB 00:01
(26/48): pacemaker-1.1.7-2.fc17.x86_64.rpm | 362 kB 00:02
(27/48): pacemaker-cli-1.1.7-2.fc17.x86_64.rpm | 368 kB 00:02
(28/48): pacemaker-cluster-libs-1.1.7-2.fc17.x86_64.rpm | 77 kB 00:00
(29/48): pacemaker-libs-1.1.7-2.fc17.x86_64.rpm | 322 kB 00:01
(30/48): perl-5.14.2-211.fc17.x86_64.rpm | 10 MB 00:15
(31/48): perl-Carp-1.22-2.fc17.noarch.rpm | 17 kB 00:00
(32/48): perl-Module-Pluggable-3.90-211.fc17.noarch.rpm | 47 kB 00:00
(33/48): perl-PathTools-3.33-211.fc17.x86_64.rpm | 105 kB 00:00
(34/48): perl-Pod-Escapes-1.04-211.fc17.noarch.rpm | 40 kB 00:00
(35/48): perl-Pod-Simple-3.16-211.fc17.noarch.rpm | 223 kB 00:00
(36/48): perl-Scalar-List-Utils-1.25-1.fc17.x86_64.rpm | 33 kB 00:01
(37/48): perl-Socket-2.001-1.fc17.x86_64.rpm | 44 kB 00:00
(38/48): perl-TimeDate-1.20-6.fc17.noarch.rpm | 43 kB 00:00
(39/48): perl-libs-5.14.2-211.fc17.x86_64.rpm | 628 kB 00:00
(40/48): perl-macros-5.14.2-211.fc17.x86_64.rpm | 32 kB 00:00
(41/48): perl-threads-1.86-2.fc17.x86_64.rpm | 47 kB 00:00
(42/48): perl-threads-shared-1.40-2.fc17.x86_64.rpm | 36 kB 00:00
(43/48): quota-4.00-3.fc17.x86_64.rpm | 160 kB 00:00
(44/48): quota-nls-4.00-3.fc17.noarch.rpm | 74 kB 00:00
(45/48): resource-agents-3.9.2-2.fc17.1.x86_64.rpm | 466 kB 00:00
(46/48): rpcbind-0.2.0-16.fc17.x86_64.rpm | 52 kB 00:00
(47/48): tcp_wrappers-7.6-69.fc17.x86_64.rpm | 72 kB 00:00
(48/48): xfsprogs-3.1.8-1.fc17.x86_64.rpm | 715 kB 00:03
-----
Total 333 kB/s | 18 MB 00:55
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-fedora-x86_64
Importing GPG key 0x1ACA3465:
  Userid : "Fedora (17) <fedora@fedoraproject.org>"
  Fingerprint: cac4 3fb7 74a4 a673 d81c 5de7 50e9 4c99 1aca 3465
  Package : fedora-release-17-0.8.noarch (@anaconda-0)
  From : /etc/pki/rpm-gpg/RPM-GPG-KEY-fedora-x86_64
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : libqb-0.11.1-1.fc17.x86_64 1/48
  Installing : libtool-ltdl-2.4.2-3.fc17.x86_64 2/48
  Installing : cluster-glue-libs-1.0.6-9.fc17.1.x86_64 3/48
  Installing : libxslt-1.1.26-9.fc17.x86_64 4/48
  Installing : 1:perl-Pod-Escapes-1.04-211.fc17.noarch 5/48
  Installing : perl-threads-1.86-2.fc17.x86_64 6/48
  Installing : 4:perl-macros-5.14.2-211.fc17.x86_64 7/48
  Installing : 1:perl-Pod-Simple-3.16-211.fc17.noarch 8/48
  Installing : perl-Socket-2.001-1.fc17.x86_64 9/48
  Installing : perl-Carp-1.22-2.fc17.noarch 10/48
  Installing : 4:perl-libs-5.14.2-211.fc17.x86_64 11/48
  Installing : perl-threads-shared-1.40-2.fc17.x86_64 12/48
  Installing : perl-Scalar-List-Utils-1.25-1.fc17.x86_64 13/48

```

Capitolo 2. Installazione

Installing	: 1:perl-Module-Pluggable-3.90-211.fc17.noarch	14/48
Installing	: perl-PathTools-3.33-211.fc17.x86_64	15/48
Installing	: 4:perl-5.14.2-211.fc17.x86_64	16/48
Installing	: libibverbs-1.1.6-2.fc17.x86_64	17/48
Installing	: keyutils-1.5.5-2.fc17.x86_64	18/48
Installing	: libgssglue-0.3-1.fc17.x86_64	19/48
Installing	: libtirpc-0.2.2-2.1.fc17.x86_64	20/48
Installing	: 1:net-snmp-libs-5.7.1-4.fc17.x86_64	21/48
Installing	: rpcbind-0.2.0-16.fc17.x86_64	22/48
Installing	: librdmacm-1.0.15-1.fc17.x86_64	23/48
Installing	: corosynclib-1.99.9-1.fc17.x86_64	24/48
Installing	: corosync-1.99.9-1.fc17.x86_64	25/48
error reading information on service corosync: No such file or directory		
Installing	: 1:perl-TimeDate-1.20-6.fc17.noarch	26/48
Installing	: 1:quota-nls-4.00-3.fc17.noarch	27/48
Installing	: tcp_wrappers-7.6-69.fc17.x86_64	28/48
Installing	: 1:quota-4.00-3.fc17.x86_64	29/48
Installing	: libnfsidmap-0.25-1.fc17.x86_64	30/48
Installing	: 1:libwbclient-3.6.3-81.fc17.1.x86_64	31/48
Installing	: libnet-1.1.5-3.fc17.x86_64	32/48
Installing	: 2:ethtool-3.2-2.fc17.x86_64	33/48
Installing	: libevent-2.0.14-2.fc17.x86_64	34/48
Installing	: 1:nfs-utils-1.2.5-12.fc17.x86_64	35/48
Installing	: libtalloc-2.0.7-4.fc17.x86_64	36/48
Installing	: cifs-utils-5.3-2.fc17.x86_64	37/48
Installing	: libtasn1-2.12-1.fc17.x86_64	38/48
Installing	: OpenIPMI-libs-2.0.18-13.fc17.x86_64	39/48
Installing	: cluster-glue-1.0.6-9.fc17.1.x86_64	40/48
Installing	: p11-kit-0.12-1.fc17.x86_64	41/48
Installing	: gnutls-2.12.17-1.fc17.x86_64	42/48
Installing	: pacemaker-libs-1.1.7-2.fc17.x86_64	43/48
Installing	: pacemaker-cluster-libs-1.1.7-2.fc17.x86_64	44/48
Installing	: pacemaker-cli-1.1.7-2.fc17.x86_64	45/48
Installing	: xfsprogs-3.1.8-1.fc17.x86_64	46/48
Installing	: resource-agents-3.9.2-2.fc17.1.x86_64	47/48
Installing	: pacemaker-1.1.7-2.fc17.x86_64	48/48
Verifying	: xfsprogs-3.1.8-1.fc17.x86_64	1/48
Verifying	: 1:net-snmp-libs-5.7.1-4.fc17.x86_64	2/48
Verifying	: corosync-1.99.9-1.fc17.x86_64	3/48
Verifying	: cluster-glue-1.0.6-9.fc17.1.x86_64	4/48
Verifying	: perl-PathTools-3.33-211.fc17.x86_64	5/48
Verifying	: p11-kit-0.12-1.fc17.x86_64	6/48
Verifying	: 1:perl-Pod-Simple-3.16-211.fc17.noarch	7/48
Verifying	: OpenIPMI-libs-2.0.18-13.fc17.x86_64	8/48
Verifying	: libtasn1-2.12-1.fc17.x86_64	9/48
Verifying	: perl-threads-1.86-2.fc17.x86_64	10/48
Verifying	: 1:perl-Pod-Escapes-1.04-211.fc17.noarch	11/48
Verifying	: pacemaker-1.1.7-2.fc17.x86_64	12/48
Verifying	: 4:perl-5.14.2-211.fc17.x86_64	13/48
Verifying	: gnutls-2.12.17-1.fc17.x86_64	14/48
Verifying	: perl-threads-shared-1.40-2.fc17.x86_64	15/48
Verifying	: 4:perl-macros-5.14.2-211.fc17.x86_64	16/48
Verifying	: 1:perl-Module-Pluggable-3.90-211.fc17.noarch	17/48
Verifying	: 1:nfs-utils-1.2.5-12.fc17.x86_64	18/48
Verifying	: cluster-glue-libs-1.0.6-9.fc17.1.x86_64	19/48
Verifying	: pacemaker-libs-1.1.7-2.fc17.x86_64	20/48
Verifying	: libtalloc-2.0.7-4.fc17.x86_64	21/48
Verifying	: libevent-2.0.14-2.fc17.x86_64	22/48
Verifying	: perl-Socket-2.001-1.fc17.x86_64	23/48
Verifying	: libgssglue-0.3-1.fc17.x86_64	24/48
Verifying	: perl-Carp-1.22-2.fc17.noarch	25/48
Verifying	: libtirpc-0.2.2-2.1.fc17.x86_64	26/48
Verifying	: 2:ethtool-3.2-2.fc17.x86_64	27/48
Verifying	: 4:perl-libs-5.14.2-211.fc17.x86_64	28/48
Verifying	: libxslt-1.1.26-9.fc17.x86_64	29/48
Verifying	: rpcbind-0.2.0-16.fc17.x86_64	30/48
Verifying	: librdmacm-1.0.15-1.fc17.x86_64	31/48


```

Verifying : resource-agents-3.9.2-2.fc17.1.x86_64 32/48
Verifying : 1:quota-4.00-3.fc17.x86_64 33/48
Verifying : 1:perl-TimeDate-1.20-6.fc17.noarch 34/48
Verifying : perl-Scalar-List-Utils-1.25-1.fc17.x86_64 35/48
Verifying : libtool-ltdl-2.4.2-3.fc17.x86_64 36/48
Verifying : pacemaker-cluster-libs-1.1.7-2.fc17.x86_64 37/48
Verifying : cifs-utils-5.3-2.fc17.x86_64 38/48
Verifying : libnet-1.1.5-3.fc17.x86_64 39/48
Verifying : corosync-1.99.9-1.fc17.x86_64 40/48
Verifying : libqb-0.11.1-1.fc17.x86_64 41/48
Verifying : 1:libwbclient-3.6.3-81.fc17.1.x86_64 42/48
Verifying : libnfsidmap-0.25-1.fc17.x86_64 43/48
Verifying : tcp_wrappers-7.6-69.fc17.x86_64 44/48
Verifying : keyutils-1.5.5-2.fc17.x86_64 45/48
Verifying : libibverbs-1.1.6-2.fc17.x86_64 46/48
Verifying : 1:quota-nls-4.00-3.fc17.noarch 47/48
Verifying : pacemaker-cli-1.1.7-2.fc17.x86_64 48/48

Installed:
corosync.x86_64 0:1.99.9-1.fc17 pacemaker.x86_64 0:1.1.7-2.fc17

Dependency Installed:
OpenIPMI-libs.x86_64 0:2.0.18-13.fc17 cifs-utils.x86_64 0:5.3-2.fc17
cluster-glue.x86_64 0:1.0.6-9.fc17.1 cluster-glue-libs.x86_64 0:1.0.6-9.fc17.1
corosync-libs.x86_64 0:1.99.9-1.fc17 ethtool.x86_64 2:3.2-2.fc17
gnutls.x86_64 0:2.12.17-1.fc17 keyutils.x86_64 0:1.5.5-2.fc17
libevent.x86_64 0:2.0.14-2.fc17 libgssglue.x86_64 0:0.3-1.fc17
libibverbs.x86_64 0:1.1.6-2.fc17 libnet.x86_64 0:1.1.5-3.fc17
libnfsidmap.x86_64 0:0.25-1.fc17 libqb.x86_64 0:0.11.1-1.fc17
librdmacm.x86_64 0:1.0.15-1.fc17 libtalloc.x86_64 0:2.0.7-4.fc17
libtasn1.x86_64 0:2.12-1.fc17 libtirpc.x86_64 0:0.2.2-2.1.fc17
libtool-ltdl.x86_64 0:2.4.2-3.fc17 libwbclient.x86_64 1:3.6.3-81.fc17.1
libxslt.x86_64 0:1.1.26-9.fc17 net-snmp-libs.x86_64 1:5.7.1-4.fc17
nfs-utils.x86_64 1:1.2.5-12.fc17 p11-kit.x86_64 0:0.12-1.fc17
pacemaker-cli.x86_64 0:1.1.7-2.fc17 pacemaker-cluster-libs.x86_64 0:1.1.7-2.fc17
pacemaker-libs.x86_64 0:1.1.7-2.fc17 perl.x86_64 4:5.14.2-211.fc17
perl-Carp.noarch 0:1.22-2.fc17 perl-Module-Pluggable.noarch 1:3.90-211.fc17
perl-PathTools.x86_64 0:3.33-211.fc17 perl-Pod-Escapes.noarch 1:1.04-211.fc17
perl-Pod-Simple.noarch 1:3.16-211.fc17 perl-Scalar-List-Utils.x86_64 0:1.25-1.fc17
perl-Socket.x86_64 0:2.001-1.fc17 perl-TimeDate.noarch 1:1.20-6.fc17
perl-libs.x86_64 4:5.14.2-211.fc17 perl-macros.x86_64 4:5.14.2-211.fc17
perl-threads.x86_64 0:1.86-2.fc17 perl-threads-shared.x86_64 0:1.40-2.fc17
quota.x86_64 1:4.00-3.fc17 quota-nls.noarch 1:4.00-3.fc17
resource-agents.x86_64 0:3.9.2-2.fc17.1 rpcbind.x86_64 0:0.2.0-16.fc17
tcp_wrappers.x86_64 0:7.6-69.fc17 xfsprogs.x86_64 0:3.1.8-1.fc17

Complete!
[root@pcmk-1 ~]#

```

Now install the cluster software on the second node.

2.5. Setup

2.5.1. Preparation - Multicast

Choose a port number and [multi-cast](http://en.wikipedia.org/wiki/Multicast_address)¹⁰ address. http://en.wikipedia.org/wiki/Multicast_address

Be sure that the values you chose do not conflict with any existing clusters you might have. For this document, I have chosen port 4000 and used 239.255.1.1 as the multi-cast address.

¹⁰ <http://en.wikipedia.org/wiki/Multicast>

2.5.2. Notes on Multicast Address Assignment

There are several subtle points that often deserve consideration when choosing/assigning multicast addresses for corosync.¹¹

1. Avoid 224.0.0.x

Traffic to addresses of the form 224.0.0.x is often flooded to all switch ports. This address range is reserved for link-local uses. Many routing protocols assume that all traffic within this range will be received by all routers on the network. Hence (at least all Cisco) switches flood traffic within this range. The flooding behavior overrides the normal selective forwarding behavior of a multicast-aware switch (e.g. IGMP snooping, CGMP, etc.).

2. Watch for 32:1 overlap

32 non-contiguous IP multicast addresses are mapped onto each Ethernet multicast address. A receiver that joins a single IP multicast group implicitly joins 31 others due to this overlap. Of course, filtering in the operating system discards undesired multicast traffic from applications, but NIC bandwidth and CPU resources are nonetheless consumed discarding it. The overlap occurs in the 5 high-order bits, so it's best to use the 23 low-order bits to make distinct multicast streams unique. For example, IP multicast addresses in the range 239.0.0.0 to 239.127.255.255 all map to unique Ethernet multicast addresses. However, IP multicast address 239.128.0.0 maps to the same Ethernet multicast address as 239.0.0.0, 239.128.0.1 maps to the same Ethernet multicast address as 239.0.0.1, etc.

3. Avoid x.0.0.y and x.128.0.y

Combining the above two considerations, it's best to avoid using IP multicast addresses of the form x.0.0.y and x.128.0.y since they all map onto the range of Ethernet multicast addresses that are flooded to all switch ports.

4. Watch for address assignment conflicts

IANA¹² administers *Internet multicast addresses*¹³. Potential conflicts with Internet multicast address assignments can be avoided by using *GLOP addressing*¹⁴ (AS¹⁵ required) or *administratively scoped*¹⁶ addresses. Such addresses can be safely used on a network connected to the Internet without fear of conflict with multicast sources originating on the Internet. Administratively scoped addresses are roughly analogous to the unicast address space for *private internets*¹⁷. Site-local multicast addresses are of the form 239.255.x.y, but can grow down to 239.252.x.y if needed. Organization-local multicast addresses are of the form 239.192-251.x.y, but can grow down to 239.x.y.z if needed.

For a more detailed treatment (57 pages!), see *Cisco's Guidelines for Enterprise IP Multicast Address Allocation*¹⁸ paper.

¹¹ This information is borrowed from, the now defunct, <http://web.archive.org/web/20101211210054/http://29west.com/docs/THPM/multicast-address-assignment.html>

¹² <http://www.iana.org/>

¹³ <http://www.iana.org/assignments/multicast-addresses>

¹⁴ <http://www.ietf.org/rfc/rfc3180.txt>

¹⁵ http://en.wikipedia.org/wiki/Autonomous_system_%28Internet%29

¹⁶ <http://www.ietf.org/rfc/rfc2365.txt>

¹⁷ <http://www.ietf.org/rfc/rfc1918.txt>

¹⁸ http://www.cisco.com/en/US/tech/tk828/technologies_white_paper09186a00802d4643.shtml

2.5.3. Configurare Corosync



Importante

The instructions below only apply for a machine with a single NIC. If you have a more complicated setup, you should edit the configuration manually.

```
# export ais_port=4000
# export ais_mcast=239.255.1.1
```

Next we automatically determine the hosts address. By not using the full address, we make the configuration suitable to be copied to other nodes.

```
export ais_addr=`ip addr | grep "inet " | tail -n 1 | awk '{print $4}' | sed s/255/0/g`
```

Display and verify the configuration options

```
# env | grep ais_
ais_mcast=239.255.1.1
ais_port=4000
ais_addr=192.168.122.0
```

Once you're happy with the chosen values, update the Corosync configuration

```
# cp /etc/corosync/corosync.conf.example /etc/corosync/corosync.conf
# sed -i.bak "s/.*mcastaddr:./mcastaddr:\ $ais_mcast/g" /etc/corosync/corosync.conf
# sed -i.bak "s/.*mcastport:./mcastport:\ $ais_port/g" /etc/corosync/corosync.conf
# sed -i.bak "s/.*\tbindnetaddr:./bindnetaddr:\ $ais_addr/g" /etc/corosync/corosync.conf
```

Lastly, you'll need to enable quorum

```
cat << END >> /etc/corosync/corosync.conf
quorum {
    provider: corosync_votequorum
    expected_votes: 2
}
END
```

The final /etc/corosync.conf configuration on each node should look something like the sample in Appendix B, Sample Corosync Configuration.



Importante

Pacemaker used to obtain membership and quorum from a custom Corosync plugin. This plugin also had the capability to start Pacemaker automatically when Corosync was started.

Neither behavior is possible with Corosync 2.0 and beyond as support for plugins was removed.

Instead, Pacemaker must be started as a separate service. Also, since Pacemaker made use of the plugin for message routing, a node using the plugin (Corosync prior to 2.0) cannot talk to one that isn't (Corosync 2.0+).

Rolling upgrades between these versions are therefore not possible and an alternate strategy¹⁹ must be used.

2.5.4. Propagate the Configuration

Now we need to copy the changes so far to the other node:

```
# for f in /etc/corosync/corosync.conf /etc/hosts; do scp $f pcmk-2:$f ; done
corosync.conf          100% 1528    1.5KB/s   00:00
hosts                  100% 281     0.3KB/s   00:00
#
```

¹⁹ http://www.clusterlabs.org/doc/en-US/Pacemaker/1.1/html/Pacemaker_Explained/ap-upgrade.html

Pacemaker Tools

Indice

3.1. Utilizzare i tool di Pacemaker	25
---	----

3.1. Utilizzare i tool di Pacemaker

Nell'oscuro passato configurare Pacemaker richiedeva all'amministratore di leggere e scrivere XML. In vero stile UNIX c'erano inoltre parecchi comandi differendi, specializzati in aspetti differenti di interrogazione ed aggiornamento del cluster.

All of that has been greatly simplified with the creation of unified command-line shells (and GUIs) that hide all the messy XML scaffolding.

These shells take all the individual aspects required for managing and configuring a cluster, and packs them into one simple to use command line tool.

They even allow you to queue up several changes at once and commit them atomically.

There are currently two command-line shells that people use, **pcs** and **crmsh**. This edition of Clusters from Scratch is based on **crmsh**. Start by taking some time to familiarize yourself with what it can do.



Nota

The two shells share many concepts but the scope, layout and syntax does differ, so make sure you read the version of this guide that corresponds to the software installed on your system.

crm --help

The primary tool for monitoring the status of the cluster is `crm_mon` (also available as `crm status`). It can be run in a variety of modes and has a number of output options. To find out about any of the tools that come with Pacemaker, simply invoke them with the `--help` option or consult the included man pages. Both sets of output are created from the tool, and so will always be in sync with each other and the tool itself.

Additionally, the Pacemaker version and supported cluster stack(s) are available via the `--feature` option to `pacemakerd`.

```
# pacemakerd --features
```

```
Pacemaker 1.1.9-3.fc20.2 (Build: 781a388)
Supporting v3.0.7: generated-manpages agent-manpages ncurses libqb-logging libqb-ipc
upstart systemd nagios corosync-native
```

```
# crm_mon --help
```

```
crm_mon - Provides a summary of cluster's current state.

Outputs varying levels of detail in a number of different formats.
```

```
Usage: crm_mon mode [options]
Options:
-?, --help      This text
-$, --version   Version information
-V, --verbose   Increase debug output
-Q, --quiet     Display only essential output

Modes:
-h, --as-html=value Write cluster status to the named html file
-X, --as-xml       Write cluster status as xml to stdout. This will enable one-shot mode.
-w, --web-cgi      Web mode with output suitable for cgi
-s, --simple-status Display the cluster status once as a simple one line output (suitable
for nagios)

Display Options:
-n, --group-by-node  Group resources by node
-r, --inactive      Display inactive resources
-f, --failcounts     Display resource fail counts
-o, --operations     Display resource operation history
-t, --timing-details  Display resource operation history with timing details
-c, --tickets        Display cluster tickets
-W, --watch-fencing  Listen for fencing events. For use with --external-agent, --mail-to
and/or --snmp-traps where supported
-A, --show-node-attributes Display node attributes

Additional Options:
-i, --interval=value Update frequency in seconds
-1, --one-shot       Display the cluster status once on the console and exit
-N, --disable-ncurses Disable the use of ncurses
-d, --daemonize      Run in the background as a daemon
-p, --pid-file=value (Advanced) Daemon pid file location
-E, --external-agent=value A program to run when resource operations take place.
-e, --external-recipient=value A recipient for your program (assuming you want the program
to send something to someone).

Examples:

Display the cluster status on the console with updates as they occur:

# crm_mon

Display the cluster status on the console just once then exit:

# crm_mon -1

Display your cluster status, group resources by node, and include inactive resources in the
list:

# crm_mon --group-by-node --inactive

Start crm_mon as a background daemon and have it write the cluster status to an HTML file:

# crm_mon --daemonize --as-html /path/to/docroot/filename.html

Start crm_mon and export the current cluster status as xml to stdout, then exit.:

# crm_mon --as-xml

Report bugs to pacemaker@oss.clusterlabs.org
```

**Nota**

Se SNMP e/o le opzioni mail non sono listate allora Pacemaker non è stato compilato per supportare queste funzionalità. Potrebbe essere dovuto ad una scelta della distribuzione che si sta utilizzando oppure al fatto che le librerie necessarie non sono disponibili. In questo caso è bene contattare il fornitore dei pacchetti per avere maggiori dettagli.

Verifica dell'installazione del cluster

Indice

4.1. Verifica dell'installazione di Corosync	29
4.2. Verifica dell'installazione di Pacemaker	30

4.1. Verifica dell'installazione di Corosync

Start Corosync on the first node

```
# systemctl start corosync.service
```

The first thing to check is if cluster communication is happy, for that we use **corosync-cfgtool**.

```
# corosync-cfgtool -s
Printing ring status.
Local node ID 1702537408
RING ID 0
      id      = 192.168.122.101
      status   = ring 0 active with no faults
```

We can see here that everything appears normal with our fixed IP address, not a 127.0.0.x loopback address, listed as the **id** and **no faults** for the status.

If you see something different, you might want to start by checking the node's network, firewall and selinux configurations.

Next we check the membership and quorum APIs:

```
# corosync-cmapctl | grep members
runtime.totem.pg.mrp.srp.members.1702537408.ip (str) = r(0) ip(192.168.122.101)
runtime.totem.pg.mrp.srp.members.1702537408.join_count (u32) = 1
runtime.totem.pg.mrp.srp.members.1702537408.status (str) = joined

# corosync-quorumtool -l
Membership information
-----
      Nodeid      Votes Name
1702537408          1 pcmk-1
```

The node see's itself in both locations which is a good sign.

If the node list is empty when you call **corosync-quorumtool**, then you've not correctly quorum in *corosync.conf*.

With everything looking healthy, we start Corosync on the second node and run the same communications check.

```
# ssh pcmk-2 -- systemctl start corosync.service
# ssh pcmk-2 -- corosync-cfgtool -s
Printing ring status.
Local node ID 1719314624
RING ID 0
      id      = 192.168.122.102
```

```
status = ring 0 active with no faults
```

Everything appears to look ok from **pcmk-2**, time to re-run the membership and quorum checks to see if it shows up there too.

Again, if you see something different to the above, check for the usual suspects: network, firewall and selinux.

```
# corosync-cmapctl | grep members
runtime.totem.pg.mrp.srp.members.1702537408.ip (str) = r(0) ip(192.168.122.101)
runtime.totem.pg.mrp.srp.members.1702537408.join_count (u32) = 1
runtime.totem.pg.mrp.srp.members.1702537408.status (str) = joined
runtime.totem.pg.mrp.srp.members.1719314624.ip (str) = r(0) ip(192.168.122.102)
runtime.totem.pg.mrp.srp.members.1719314624.join_count (u32) = 1
runtime.totem.pg.mrp.srp.members.1719314624.status (str) = joined

# corosync-quorumtool -l

Membership information
-----
Nodeid      Votes Name
1702537408   1 pcmk-1
1719314624   1 pcmk-2
```

All good!

4.2. Verifica dell'installazione di Pacemaker

Now that we have confirmed that Corosync is functional we can check the rest of the stack. Start Pacemaker and check the necessary processes have been started.

```
# systemctl start pacemaker.service
# ps axf
PID TTY      STAT   TIME COMMAND
  2 ?        S      0:00 [kthreadd]
...lots of processes...
28019 ?        Ssl    0:03 /usr/sbin/corosync
28047 ?        Ss     0:00 /usr/sbin/pacemakerd -f
28048 ?        Ss     0:00 \_ /usr/libexec/pacemaker/cib
28049 ?        Ss     0:00 \_ /usr/libexec/pacemaker/stonithd
28050 ?        Ss     0:00 \_ /usr/lib64/heartbeat/lrmd
28051 ?        Ss     0:00 \_ /usr/libexec/pacemaker/attrd
28052 ?        Ss     0:00 \_ /usr/libexec/pacemaker/pengine
28053 ?        Ss     0:00 \_ /usr/libexec/pacemaker/crmd
```

If that looks ok, check the logs and crm_mon.

```
# grep pacemakerd /var/log/messages | grep -e get_cluster_type -e read_config
Apr  3 09:19:32 pcmk-1 pacemakerd[28047]: info: get_cluster_type: Detected an active
'corosync' cluster
Apr  3 09:19:32 pcmk-1 pacemakerd[28047]: info: read_config: Reading configure for stack:
corosync
# crm_mon -l
=====
Last updated: Tue Apr  3 09:21:37 2012
Last change: Tue Apr  3 09:19:54 2012 via crmd on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
 1 Nodes configured, unknown expected votes
 0 Resources configured.
```

```
=====
Online: [ pcmk-1 ]
```

Next, check for any ERRORS during startup - there shouldn't be any.

```
# grep -i error /var/log/messages
```

Repeat on the other node and display the cluster's status.

```
# ssh pcmk-2 -- systemctl start pacemaker.service
# crm_mon -1
=====
Last updated: Tue Apr  3 09:26:23 2012
Last change: Tue Apr  3 09:26:21 2012 via crmd on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
0 Resources configured.
=====
Online: [ pcmk-1 pcmk-2 ]
```


Creare un cluster Active/Passive

Indice

5.1. Esplorare la configurazione esistente	33
5.2. Aggiungere una risorsa	34
5.3. Effettuare un Failover	36
5.3.1. Quorum e Cluster a due nodi	36
5.3.2. Evitare che le risorse si muovano dopo il recovery	38

5.1. Esplorare la configurazione esistente

Quando Pacemaker viene avviato automatica registra il numero ed i dettagli dei nodi nel cluster, così come lo stack è utilizzato e la versione di Pacemaker utilizzata.

Ecco come dovrebbe apparire la configurazione base.

```
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync"
```

For those that are not of afraid of XML, you can see the raw configuration by appending "xml" to the previous command.

Esempio 5.1. Questo è l'ultimo XML ad essere utilizzato nel documento.

```
# crm configure show xml
```

```
<?xml version="1.0" ?>
<cib admin_epoch="0" cib-last-written="Tue Apr  3 09:26:21 2012" crm_feature_set="3.0.6"
dc-uuid="1702537408" epoch="4" have-quorum="1" num_updates="14" update-client="crmd"
update-origin="pcmk-1" validate-with="pacemaker-1.2">
  <configuration>
    <crm_config>
      <cluster_property_set id="cib-bootstrap-options">
        <nvpair id="cib-bootstrap-options-dc-version" name="dc-version"
value="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff"/>
        <nvpair id="cib-bootstrap-options-cluster-infrastructure" name="cluster-
infrastructure" value="corosync"/>
      </cluster_property_set>
    </crm_config>
    <nodes>
      <node id="1719314624" type="normal" uname="pcmk-2"/>
      <node id="1702537408" type="normal" uname="pcmk-1"/>
    </nodes>
    <resources/>
    <constraints/>
  </configuration>
</cib>
```

Prima di effettuare qualsiasi cambiamento è buona norma controllare la validità della configurazione.

```
# crm_verify -L -V
error: unpack_resources: Resource start-up disabled since no STONITH resources have been
defined
error: unpack_resources: Either configure some or disable STONITH with the stonith-enabled
option
error: unpack_resources: NOTE: Clusters with shared data need STONITH to ensure data
integrity
Errors found during check: config not valid
-V may provide more details
```

Come si può notare il tool ha trovato qualche errore.

In order to guarantee the safety of your data ¹, the default for STONITH ² in Pacemaker is **enabled**. However it also knows when no STONITH configuration has been supplied and reports this as a problem (since the cluster would not be able to make progress if a situation requiring node fencing arose).

Per adesso la funzionalità verrà disabilitata e configurata in seguito nella sezione Configurare STONITH. E' importante notare che l'uso di STONITH è altamente consigliato, disabilitarlo indica al cluster di dare per scontato che i nodi falliti vengano spenti. Alcuni rivenditori potrebbero rifiutarsi di supportare cluster che hanno STONITH disabilitato.

To disable STONITH, we set the *stonith-enabled* cluster option to false.

```
# crm configure property stonith-enabled=false
# crm_verify -L
```

Con la nuova opzione impostata la configurazione del cluster è ora valida.



Avvertimento

The use of `stonith-enabled=false` is completely inappropriate for a production cluster. We use it here to defer the discussion of its configuration which can differ widely from one installation to the next. See [Sezione 9.1, «What Is STONITH»](#) for information on why STONITH is important and details on how to configure it.

5.2. Aggiungere una risorsa

The first thing we should do is configure an IP address. Regardless of where the cluster service(s) are running, we need a consistent address to contact them on. Here I will choose and add 192.168.122.120 as the floating address, give it the imaginative name ClusterIP and tell the cluster to check that its running every 30 seconds.

¹ If the data is corrupt, there is little point in continuing to make it available

² A common node fencing mechanism. Used to ensure data integrity by powering off "bad" nodes



Importante

L'indirizzo scelto non dovrà essere già associato ad un nodo fisico

```
# crm configure primitive ClusterIP ocf:heartbeat:IPaddr2 \
  params ip=192.168.122.120 cidr_netmask=32 \
  op monitor interval=30s
```

The other important piece of information here is ocf:heartbeat:IPaddr2.

This tells Pacemaker three things about the resource you want to add. The first field, ocf, is the standard to which the resource script conforms to and where to find it. The second field is specific to OCF resources and tells the cluster which namespace to find the resource script in, in this case heartbeat. The last field indicates the name of the resource script.

To obtain a list of the available resource classes, run

```
# crm ra classes
heartbeat
lsb
ocf / heartbeat pacemaker
stonith
```

To then find all the OCF resource agents provided by Pacemaker and Heartbeat, run

```
# crm ra list ocf pacemaker
ClusterMon      Dummy          HealthCPU      HealthSMART    Stateful       SysInfo
SystemHealth    controld       o2cb           ping           pingd
# crm ra list ocf heartbeat
AoEtarget        AudibleAlarm   CTDB            EvmSSCC         ClusterMon
Delay            Dummy          EvmSSCC         IPaddr          Evmsd
Filesystem       ICP            IPaddr          LinuxSCSI       IPaddr2
IPsrcaddr        IPv6addr       LVM             Pure-FTPd       LinuxSCSI
MailTo           ManageRAID     ManageVE        SAPDatabase     Pure-FTPd
Raid1            Route          SphinxSearchDaemon  Squid           SAPInstance
SendArp          ServerAID      VIPArip         VirtualDomain
Stateful          SysInfo       WinPopup        Xen
WAS              WAS6           apache          connttrackd
Xinetd           anything       eDir88          ethmonitor
db2              drbd           iSCSILogicalUnit  iSCSITarget
exportfs         fio            jboss           ldirectord
ids              iscsi          mysql-proxy     nfsserver
lxc              mysql          oralsnr         pgsql
nginx            oracle         postfix         proftpd
pingd            portblock     sfex            symlink
rsyncd           scsi2reservation
syslog-ng        tomcat        vmware
```

A questo punto va verificato come la risorsa IP sia stata aggiunta e visualizzato lo stato del cluster per vedere che ora è attiva.

```
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
```

```
params ip="192.168.122.120" cidr_netmask="32" \  
op monitor interval="30s"\  
property $id="cib-bootstrap-options" \  
dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \  
cluster-infrastructure="corosync" \  
stonith-enabled="false"\  
# crm_mon -1\  
=====\  
Last updated: Tue Apr  3 09:56:50 2012\  
Last change: Tue Apr  3 09:54:37 2012 via cibadmin on pcmk-1\  
Stack: corosync\  
Current DC: pcmk-1 (1702537408) - partition with quorum\  
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff\  
2 Nodes configured, unknown expected votes\  
1 Resources configured.  
=====\  
  
Online: [ pcmk-1 pcmk-2 ]  
  
ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-1
```

5.3. Effettuare un Failover

Trattandosi di un cluster ad alta-affidabilità, è necessario testare il failover della nostra nuova risorsa prima di proseguire.

Per prima cosa va identificato da quale nodo l'indirizzo IP è erogato

```
# crm resource status ClusterIP\  
resource ClusterIP is running on: pcmk-1
```

Shut down Pacemaker and Corosync on that machine.

```
# ssh pcmk-1 -- service pacemaker stop\  
# ssh pcmk-1 -- service corosync stop
```

Once Corosync is no longer running, go to the other node and check the cluster status with `crm_mon`.

```
# crm_mon -1\  
=====\  
Last updated: Tue Apr  3 10:01:28 2012\  
Last change: Tue Apr  3 09:54:39 2012 via cibadmin on pcmk-1\  
Stack: corosync\  
Current DC: pcmk-2 (1719314624) - partition WITHOUT quorum\  
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff\  
2 Nodes configured, unknown expected votes\  
1 Resources configured.  
=====\  
  
Online: [ pcmk-2 ]  
OFFLINE: [ pcmk-1 ]
```

There are three things to notice about the cluster's current state. The first is that, as expected, **pcmk-1** is now offline. However we can also see that **ClusterIP** isn't running anywhere!

5.3.1. Quorum e Cluster a due nodi

This is because the cluster no longer has quorum, as can be seen by the text "partition WITHOUT quorum" in the status output. In order to reduce the possibility of data corruption, Pacemaker's default behavior is to stop all resources if the cluster does not have quorum.

Un cluster viene definito con quorum quando più della metà dei nodi conosciuti o aspettati sono online o, attraverso la matematica, quando la seguente equazione è vera:

$$\text{total_nodes} < 2 * \text{active_nodes}$$

Therefore a two-node cluster only has quorum when both nodes are running, which is no longer the case for our cluster. This would normally make the creation of a two-node cluster pointless³, however it is possible to control how Pacemaker behaves when quorum is lost. In particular, we can tell the cluster to simply ignore quorum altogether.

```
# crm configure property no-quorum-policy=ignore
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore"
```

Dopo alcuni istanti il cluster avvierà l'indirizzo IP sui nodi rimanenti. E' da notare che il cluster non ha comunque il quorum.

```
# crm_mon -1
=====
Last updated: Tue Apr  3 10:02:46 2012
Last change: Tue Apr  3 10:02:08 2012 via cibadmin on pcmk-2
Stack: corosync
Current DC: pcmk-2 (1719314624) - partition WITHOUT quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
1 Resources configured.
=====

Online: [ pcmk-2 ]
Offline: [ pcmk-1 ]

ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-2
```

Now simulate node recovery by restarting the cluster stack on **pcmk-1** and check the cluster's status. Note, if you get an authentication error with the *pcs cluster start pcmk-1* command, you must authenticate on the node using the *pcs cluster auth pcmk pcmk-1 pcmk-2* command discussed earlier.

```
# service corosync start
Starting Corosync Cluster Engine (corosync): [ OK ]
# service pacemaker start
Starting Pacemaker Cluster Manager: [ OK ]
# crm_mon
=====
Last updated: Fri Aug 28 15:32:13 2009
Stack: openais
Current DC: pcmk-2 - partition with quorum
Version: 1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f
2 Nodes configured, 2 expected votes
```

³ Actually some would argue that two-node clusters are always pointless, but that is an argument for another time

```
1 Resources configured.
=====
Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf::heartbeat:IPaddr):   Started pcmk-2
```



Nota

In the dark days, the cluster may have moved the IP back to its original location (**pcmk-1**). Usually this is no longer the case.

5.3.2. Evitare che le risorse si muovano dopo il recovery

In most circumstances, it is highly desirable to prevent healthy resources from being moved around the cluster. Moving resources almost always requires a period of downtime. For complex services like Oracle databases, this period can be quite long.

To address this, Pacemaker has the concept of resource stickiness which controls how much a service prefers to stay running where it is. You may like to think of it as the "cost" of any downtime. By default, Pacemaker assumes there is zero cost associated with moving resources and will do so to achieve "optimal"⁴ resource placement. We can specify a different stickiness for every resource, but it is often sufficient to change the default.

```
# crm configure rsc_defaults resource-stickiness=100
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
```

⁴ It should be noted that Pacemaker's definition of optimal may not always agree with that of a human's. The order in which Pacemaker processes lists of resources and nodes creates implicit preferences in situations where the administrator has not explicitly specified them

Apache - Aggiungere ulteriori servizi

Indice

6.1. Forward	39
6.2. Installazione	39
6.3. Preparazione	40
6.4. Enable the Apache status URL	41
6.5. Aggiornare la configurazione	41
6.6. Assicurarsi che le risorse funzionino sullo stesso host	42
6.7. Controllare l'ordinamento di start e stop di una risorsa	43
6.8. Specificare una Location preferita	44
6.9. Spostamento manuale delle risorse all'interno del cluster	45
6.9.1. Restituire il controllo nuovamente al cluster	46

6.1. Forward

Ora che è stato realizzato un semplice, ma funzionale cluster a due nodi active/passive, è tempo di aggiungere servizi reali. Si partirà con Apache poiché molti cluster lo includono ed è relativamente semplice da configurare.

6.2. Installazione

Before continuing, we need to make sure Apache is installed on both hosts. We also need the wget tool in order for the cluster to be able to check the status of the Apache server.

```
# yum install -y httpd wget
```

```
Loaded plugins: langpacks, presto, refresh-packagekit
fedora/metalink                | 2.6 kB    00:00
updates/metalink               | 3.2 kB    00:00
updates-testing/metalink       | 41 kB     00:00
Resolving Dependencies
--> Running transaction check
---> Package httpd.x86_64 0:2.2.22-3.fc17 will be installed
--> Processing Dependency: httpd-tools = 2.2.22-3.fc17 for package:
    httpd-2.2.22-3.fc17.x86_64
--> Processing Dependency: apr-util-ldap for package: httpd-2.2.22-3.fc17.x86_64
--> Processing Dependency: libaprutil-1.so.0()(64bit) for package: httpd-2.2.22-3.fc17.x86_64
--> Processing Dependency: libapr-1.so.0()(64bit) for package: httpd-2.2.22-3.fc17.x86_64
--> Running transaction check
---> Package apr.x86_64 0:1.4.6-1.fc17 will be installed
---> Package apr-util.x86_64 0:1.4.1-2.fc17 will be installed
---> Package apr-util-ldap.x86_64 0:1.4.1-2.fc17 will be installed
---> Package httpd-tools.x86_64 0:2.2.22-3.fc17 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package                Arch          Version              Repository            Size
=====
Installing:
httpd                  x86_64        2.2.22-3.fc17        updates-testing       823 k
wget                   x86_64        1.13.4-2.fc17        fedora                 495 k
=====
```

```
Installing for dependencies:
apr                x86_64          1.4.6-1.fc17          fedora                99 k
apr-util           x86_64          1.4.1-2.fc17          fedora                78 k
apr-util-ldap      x86_64          1.4.1-2.fc17          fedora                17 k
httpd-tools        x86_64          2.2.22-3.fc17         updates-testing       74 k

Transaction Summary
=====
Install 1 Package (+4 Dependent packages)

Total download size: 1.1 M
Installed size: 3.5 M
Downloading Packages:
(1/6): apr-1.4.6-1.fc17.x86_64.rpm           | 99 kB    00:00
(2/6): apr-util-1.4.1-2.fc17.x86_64.rpm      | 78 kB    00:00
(3/6): apr-util-ldap-1.4.1-2.fc17.x86_64.rpm | 17 kB    00:00
(4/6): httpd-2.2.22-3.fc17.x86_64.rpm        | 823 kB   00:01
(5/6): httpd-tools-2.2.22-3.fc17.x86_64.rpm  | 74 kB    00:00
(6/6): wget-1.13.4-2.fc17.x86_64.rpm        | 495 kB   00:01
-----
Total                                          238 kB/s | 1.1 MB    00:04
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : apr-1.4.6-1.fc17.x86_64          1/6
  Installing : apr-util-1.4.1-2.fc17.x86_64     2/6
  Installing : apr-util-ldap-1.4.1-2.fc17.x86_64 3/6
  Installing : httpd-tools-2.2.22-3.fc17.x86_64 4/6
  Installing : httpd-2.2.22-3.fc17.x86_64       5/6
  Installing : wget-1.13.4-2.fc17.x86_64        6/6
  Verifying  : apr-util-ldap-1.4.1-2.fc17.x86_64 1/6
  Verifying  : httpd-tools-2.2.22-3.fc17.x86_64 2/6
  Verifying  : apr-util-1.4.1-2.fc17.x86_64     3/6
  Verifying  : apr-1.4.6-1.fc17.x86_64          4/6
  Verifying  : httpd-2.2.22-3.fc17.x86_64       5/6
  Verifying  : wget-1.13.4-2.fc17.x86_64        6/6

Installed:
httpd.x86_64 0:2.2.22-3.fc17          wget.x86_64 0:1.13.4-2.fc17

Dependency Installed:
apr.x86_64 0:1.4.6-1.fc17          apr-util.x86_64 0:1.4.1-2.fc17
apr-util-ldap.x86_64 0:1.4.1-2.fc17 httpd-tools.x86_64 0:2.2.22-3.fc17

Complete!
```

6.3. Preparazione

Prima di tutto è necessario creare la pagina che Apache servirà. Su Fedora la docroot di default di Apache è `/var/www/html`, qui verrà creato il file `index`.

```
# cat <<-END >/var/www/html/index.html
<html>
<body>My Test Site - pcmk-1</body>
</html>
END
```

Per il momento l'architettura verrà semplificata in modo da servire unicamente un sito statico e sincronizzare manualmente i dati tra i due nodi. Quindi lo stesso comando andrà lanciato su `pcmk-2`.

```
[root@pcmk-2 ~]# cat <<-END >/var/www/html/index.html
<html>
```

```
<body>My Test Site - pcmk-2</body>
</html>
END
```

6.4. Enable the Apache status URL

In order to monitor the health of your Apache instance, and recover it if it fails, the resource agent used by Pacemaker assumes the server-status URL is available. Look for the following in `/etc/httpd/conf/httpd.conf` and make sure it is not disabled or commented out:

```
<Location /server-status>
    SetHandler server-status
    Order deny,allow
    Deny from all
    Allow from 127.0.0.1
</Location>
```

6.5. Aggiornare la configurazione

At this point, Apache is ready to go, all that needs to be done is to add it to the cluster. Lets call the resource `WebSite`. We need to use an OCF script called `apache` in the heartbeat namespace ¹, the only required parameter is the path to the main Apache configuration file and we'll tell the cluster to check once a minute that apache is still running.

```
# crm configure primitive WebSite ocf:heartbeat:apache \
    params configfile=/etc/httpd/conf/httpd.conf \
    statusurl="http://localhost/server-status" \
    op monitor interval=1min
WARNING: WebSite: default timeout 20s for start is smaller than the advised 40s
WARNING: WebSite: default timeout 20s for stop is smaller than the advised 60s
```

The easiest way resolve this, is to change the default:

```
# crm configure op_defaults timeout=240s
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPAddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
```

Dopo una breve attesa il cluster dovrbbe avviare apache

¹ Compare the key used here `ocf:heartbeat:apache` with the one we used earlier for the IP address: `ocf:heartbeat:IPAddr2`

```
# crm_mon -1
=====
Last updated: Tue Apr  3 11:54:29 2012
Last change: Tue Apr  3 11:54:26 2012 via crmd on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
2 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf:heartbeat:IPaddr2):      Started pcmk-2
WebSite (ocf:heartbeat:apache):      Started pcmk-1
```

Attenzione però, la risorsa WebSite non sta girando sullo stesso host dell'indirizzo IP!



Nota

If, in the **crm_mon** output, you see:

```
Failed actions:
  WebSite_start_0 (node=pcmk-2, call=301, rc=1, status=complete): unknown error
```

Then you've likely not enabled the status URL correctly. You can check if this is the problem by running:

```
wget http://127.0.0.1/server-status
```

If you see **Connection refused** in the output, then this is indeed the problem. Check to ensure that **Allow from 127.0.0.1** is present for the **<Location /server-status>** block.

6.6. Assicurarsi che le risorse funzionino sullo stesso host

To reduce the load on any one machine, Pacemaker will generally try to spread the configured resources across the cluster nodes. However we can tell the cluster that two resources are related and need to run on the same host (or not at all). Here we instruct the cluster that WebSite can only run on the host that ClusterIP is active on.

For the constraint, we need a name (choose something descriptive like website-with-ip), indicate that its mandatory (so that if ClusterIP is not active anywhere, WebSite will not be permitted to run anywhere either) by specifying a score of INFINITY and finally list the two resources.



Nota

If ClusterIP is not active anywhere, WebSite will not be permitted to run anywhere.

**Importante**

Colocation constraints are "directional", in that they imply certain things about the order in which the two resources will have a location chosen. In this case we're saying **WebSite** needs to be placed on the same machine as **ClusterIP**, this implies that we must know the location of **ClusterIP** before choosing a location for **WebSite**.

```
# crm configure colocation website-with-ip INFINITY: WebSite ClusterIP
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
colocation website-with-ip inf: WebSite ClusterIP
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
# crm_mon -1
=====
Last updated: Tue Apr  3 11:57:13 2012
Last change: Tue Apr  3 11:56:10 2012 via cibadmin on pcmk-1
Stack: corosync
Current DC: pcmk-2 (1719314624) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
2 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf:heartbeat:IPaddr2):   Started pcmk-2
WebSite        (ocf:heartbeat:apache):    Started pcmk-2
```

6.7. Controllare l'ordinamento di start e stop di una risorsa

When Apache starts, it binds to the available IP addresses. It doesn't know about any addresses we add afterwards, so not only do they need to run on the same node, but we need to make sure ClusterIP is already active before we start WebSite. We do this by adding an ordering constraint.

We need to give it a name (choose something descriptive like `apache-after-ip`), indicate that its mandatory (so that any recovery for ClusterIP will also trigger recovery of WebSite) and list the two resources in the order we need them to start.

```
# crm configure order apache-after-ip mandatory: ClusterIP WebSite
# crm configure show
```

```
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
colocation website-with-ip inf: WebSite ClusterIP
order apache-after-ip inf: ClusterIP WebSite
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
```

6.8. Specificare una Location preferita

Pacemaker does not rely on any sort of hardware symmetry between nodes, so it may well be that one machine is more powerful than the other. In such cases it makes sense to host the resources there if it is available. To do this we create a location constraint.

Again we give it a descriptive name (prefer-pcmk-1), specify the resource we want to run there (WebSite), how badly we'd like it to run there (we'll use 50 for now, but in a two-node situation almost any value above 0 will do) and the host's name.

```
# crm configure location prefer-pcmk-1 WebSite 50: pcmk-1
WARNING: prefer-pcmk-1: referenced node pcmk-1 does not exist
```

This warning should be ignored.

```
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
location prefer-pcmk-1 WebSite 50: pcmk-1
colocation website-with-ip inf: WebSite ClusterIP
order apache-after-ip inf: ClusterIP WebSite
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
# crm_mon -1
=====
Last updated: Tue Apr  3 12:02:14 2012
Last change: Tue Apr  3 11:59:42 2012 via cibadmin on pcmk-1
Stack: corosync
```



```

Current DC: pcmk-2 (1719314624) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
2 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf:heartbeat:IPaddr2):      Started pcmk-2
WebSite        (ocf:heartbeat:apache): Started pcmk-2

```

Attenzione però, le risorse sono comunque su pcmk-2!

Anche se al momento viene preferito pcmk-1 a pcmk-2, questa preferenza è (intenzionalmente) minore dello stickiness della risorsa (quanto è stato preferito non avere un downtime superfluo).

To see the current placement scores, you can use a tool called `crm_simulate`

```

# crm_simulate -sl
Current cluster status:
Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf:heartbeat:IPaddr2):      Started pcmk-2
WebSite        (ocf:heartbeat:apache): Started pcmk-2

Allocation scores:
native_color: ClusterIP allocation score on pcmk-1: 50
native_color: ClusterIP allocation score on pcmk-2: 200
native_color: WebSite allocation score on pcmk-1: -INFINITY
native_color: WebSite allocation score on pcmk-2: 100

Transition Summary:

```

6.9. Spostamento manuale delle risorse all'interno del cluster

There are always times when an administrator needs to override the cluster and force resources to move to a specific location. Underneath we use location constraints like the one we created above, happily you don't need to care. Just provide the name of the resource and the intended location, we'll do the rest.

```

# crm resource move WebSite pcmk-1
# crm_mon -1
=====
Last updated: Tue Apr  3 12:03:41 2012
Last change: Tue Apr  3 12:03:37 2012 via crm_resource on pcmk-1
Stack: corosync
Current DC: pcmk-2 (1719314624) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
2 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf:heartbeat:IPaddr2):      Started pcmk-1
WebSite        (ocf:heartbeat:apache): Started pcmk-1

```

Notice how the colocation rule we created has ensured that ClusterIP was also moved to pcmk-1. For the curious, we can see the effect of this command by examining the configuration

```
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
location cli-prefer-WebSite WebSite \
    rule $id="cli-prefer-rule-WebSite" inf: #uname eq pcmk-1
location prefer-pcmk-1 WebSite 50: pcmk-1
colocation website-with-ip inf: WebSite ClusterIP
order apache-after-ip inf: ClusterIP WebSite
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
```

The automated constraint used to move the resources to **pcmk-1** is the line beginning with **location cli-prefer-WebSite**.

6.9.1. Restituire il controllo nuovamente al cluster

Una volta terminata una qualsiasi attività che richiede lo spostamento della risorsa in pcmk-1, in questo caso nessuna, è possibile consentire al cluster di riprendere le normali operazioni con il comando `unmove`. Dal momento che in precedenza è stata configurata una default stickiness, la risorsa rimarrà in pcmk-1.

```
# crm resource unmove WebSite
# crm configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
location prefer-pcmk-1 WebSite 50: pcmk-1
colocation website-with-ip inf: WebSite ClusterIP
order apache-after-ip inf: ClusterIP WebSite
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
```

Note that the constraint is now gone. If we check the cluster status, we can also see that as expected the resources are still active on pcmk-1.

```
# crm_mon
=====
Last updated: Tue Apr  3 12:05:08 2012
Last change: Tue Apr  3 12:03:37 2012 via crm_resource on pcmk-1
Stack: corosync
Current DC: pcmk-2 (1719314624) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
2 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf:heartbeat:IPaddr2):      Started pcmk-1
WebSite        (ocf:heartbeat:apache): Started pcmk-1
```


Storage replicato con DRBD

Indice

7.1. Background	49
7.2. Installare i pacchetti DRBD	49
7.3. Configurare DRBD	50
7.3.1. Creare una partizione per DRBD	50
7.3.2. Scrivere la configurazione DRBD	51
7.3.3. Inizializzare e caricare DRBD	51
7.3.4. Popolare DRBD con i dati	52
7.4. Configurare il cluster per DRBD	53
7.4.1. Testare la migrazione	56

7.1. Background

Even if you're serving up static websites, having to manually synchronize the contents of that website to all the machines in the cluster is not ideal. For dynamic websites, such as a wiki, it's not even an option. Not everyone care afford network-attached storage but somehow the data needs to be kept in sync. Enter DRBD which can be thought of as network based RAID-1. See <http://www.drbd.org/> for more details.

7.2. Installare i pacchetti DRBD

Since its inclusion in the upstream 2.6.33 kernel, everything needed to use DRBD has shipped with Fedora since version 13. All you need to do is install it:

```
# yum install -y drbd-pacemaker drbd-udev
```

```
Loaded plugins: langpacks, presto, refresh-packagekit
Resolving Dependencies
--> Running transaction check
---> Package drbd-pacemaker.x86_64 0:8.3.11-5.fc17 will be installed
--> Processing Dependency: drbd-utils = 8.3.11-5.fc17 for package: drbd-
pacemaker-8.3.11-5.fc17.x86_64
---> Package drbd-udev.x86_64 0:8.3.11-5.fc17 will be installed
--> Running transaction check
---> Package drbd-utils.x86_64 0:8.3.11-5.fc17 will be installed
--> Finished Dependency Resolution
```

Dependencies Resolved

Package	Arch	Version	Repository	Size
Installing:				
drbd-pacemaker	x86_64	8.3.11-5.fc17	updates-testing	22 k
drbd-udev	x86_64	8.3.11-5.fc17	updates-testing	6.4 k
Installing for dependencies:				
drbd-utils	x86_64	8.3.11-5.fc17	updates-testing	183 k

Transaction Summary

Install 2 Packages (+1 Dependent package)

Total download size: 212 k

```

Installed size: 473 k
Downloading Packages:
(1/3): drbd-pacemaker-8.3.11-5.fc17.x86_64.rpm | 22 kB 00:00
(2/3): drbd-udev-8.3.11-5.fc17.x86_64.rpm | 6.4 kB 00:00
(3/3): drbd-utils-8.3.11-5.fc17.x86_64.rpm | 183 kB 00:00
-----
Total 293 kB/s | 212 kB 00:00
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : drbd-utils-8.3.11-5.fc17.x86_64 1/3
  Installing : drbd-pacemaker-8.3.11-5.fc17.x86_64 2/3
  Installing : drbd-udev-8.3.11-5.fc17.x86_64 3/3
  Verifying : drbd-pacemaker-8.3.11-5.fc17.x86_64 1/3
  Verifying : drbd-udev-8.3.11-5.fc17.x86_64 2/3
  Verifying : drbd-utils-8.3.11-5.fc17.x86_64 3/3

Installed:
  drbd-pacemaker.x86_64 0:8.3.11-5.fc17 drbd-udev.x86_64 0:8.3.11-5.fc17

Dependency Installed:
  drbd-utils.x86_64 0:8.3.11-5.fc17

Complete!

```

7.3. Configurare DRBD

Prima di configurare DRBD è necessario definire a parte lo spazio che questo dovrà utilizzare.

7.3.1. Creare una partizione per DRBD

Se si dispone di più di 1Gb liberi è bene servirsene. Negli scopi di questa guida 1Gb è fin troppo spazio per un file html e sufficiente per ospitare più avanti i metadata di GFS2.

```

# vgdisplay | grep -e Name -e Free
VG Name          vg_pcmk1
Free PE / Size   31 / 992.00 MiB
# lvs
LV      VG      Attr      LSize   Pool Origin Data%  Move Log Copy%  Convert
lv_root vg_pcmk1 -wi-ao-- 8.56g
lv_swap vg_pcmk1 -wi-ao-- 960.00m
# lvcreate -n drbd-demo -L 1G vg_pcmk1
Logical volume "drbd-demo" created
# lvs
LV      VG      Attr      LSize   Pool Origin Data%  Move Log Copy%  Convert
drbd-demo vg_pcmk1 -wi-a--- 1.00G
lv_root vg_pcmk1 -wi-ao-- 8.56g
lv_swap vg_pcmk1 -wi-ao-- 960.00m

```

La medesima operazione va ripetuta sul secondo nodo, avendo cura di utilizzare la stessa grandezza per la partizione.

```

# ssh pcmk-2 -- lvs
LV      VG      Attr      LSize   Origin Snap% Move Log Copy%  Convert
lv_root vg_pcmk1 -wi-ao-- 8.56g
lv_swap vg_pcmk1 -wi-ao-- 960.00m
# ssh pcmk-2 -- lvcreate -n drbd-demo -L 1G vg_pcmk1
Logical volume "drbd-demo" created
# ssh pcmk-2 -- lvs
LV      VG      Attr      LSize   Origin Snap% Move Log Copy%  Convert
drbd-demo vg_pcmk1 -wi-a--- 1.00G
lv_root vg_pcmk1 -wi-ao-- 8.56g

```

```
lv_swap    vg_pcmk1 -wi-ao-- 960.00m
```

7.3.2. Scrivere la configurazione DRBD

There is no series of commands for building a DRBD configuration, so simply copy the configuration below to `/etc/drbd.conf`

Informazioni dettagliate sulle direttive utilizzate in questa configurazione (e le alternative disponibili) è disponibile presso <http://www.drbd.org/users-guide/ch-configure.html>



Avvertimento

Be sure to use the names and addresses of your nodes if they differ from the ones used in this guide.

```
global {
    usage-count yes;
}
common {
    protocol C;
}
resource wwwdata {
    meta-disk internal;
    device /dev/drbd1;
    syncer {
        verify-alg sha1;
    }
    net {
        allow-two-primaries;
    }
    on pcmk-1 {
        disk /dev/vg_pcmk1/drbd-demo;
        address 192.168.122.101:7789;
    }
    on pcmk-2 {
        disk /dev/vg_pcmk1/drbd-demo;
        address 192.168.122.102:7789;
    }
}
```



Nota

TODO: Spiegare la ragione dell'opzione `allow-two-primaries`

7.3.3. Inizializzare e caricare DRBD

Sistemata la configurazione, è possibile inizializzare DRBD

```
# drbdadm create-md wwwdata
Writing meta data...
initializing activity log
NOT initialized bitmap
New drbd meta data block successfully created.
```

```
success
```

Caricando il modulo DRBD del kernel e verificando che tutto sia a posto

```
# modprobe drbd
# drbdadm up wwwdata
# cat /proc/drbd
version: 8.3.11 (api:88/proto:86-96)
srcversion: 0D2B62DEDB020A425130935

1: cs:Connected ro:Secondary/Secondary ds:Inconsistent/Inconsistent C r-----
   ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:1015740
```

Repeat on the second node

```
# ssh pcmk-2 -- drbdadm --force create-md wwwdata
Writing meta data...
initializing activity log
NOT initialized bitmap
New drbd meta data block successfully created.
success
# ssh pcmk-2 -- modprobe drbd
WARNING: Deprecated config file /etc/modprobe.conf, all config files belong into /etc/
modprobe.d/.
# ssh pcmk-2 -- drbdadm up wwwdata
# ssh pcmk-2 -- cat /proc/drbd
version: 8.3.11 (api:88/proto:86-96)
srcversion: 0D2B62DEDB020A425130935

1: cs:Connected ro:Secondary/Secondary ds:Inconsistent/Inconsistent C r-----
   ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:1015740
```

Ora va indicato a DRBD quale set di dati utilizzare. Dal momento che entrambe le parti contengono dati non necessari è possibile lanciare il seguente comando su pcmk-1:

```
# drbdadm -- --overwrite-data-of-peer primary wwwdata
# cat /proc/drbd
version: 8.3.11 (api:88/proto:86-96)
srcversion: 0D2B62DEDB020A425130935

1: cs:SyncSource ro:Primary/Secondary ds:UpToDate/Inconsistent C r-----
   ns:8064 nr:0 dw:0 dr:8728 al:0 bm:0 lo:0 pe:1 ua:0 ap:0 ep:1 wo:f oos:1007804
   [>.....] sync'ed: 0.9% (1007804/1015740)K
   finish: 0:12:35 speed: 1,320 (1,320) K/sec
```

After a while, the sync should finish and you'll see:

```
# cat /proc/drbd
version: 8.3.11 (api:88/proto:86-96)
srcversion: 0D2B62DEDB020A425130935

1: cs:Connected ro:Primary/Secondary ds:UpToDate/UpToDate C r-----
   ns:1015740 nr:0 dw:0 dr:1016404 al:0 bm:62 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:0
```

pcmk-1 is now in the Primary state which allows it to be written to. Which means it's a good point at which to create a filesystem and populate it with some data to serve up via our WebSite resource.

7.3.4. Popolare DRBD con i dati

```
# mkfs.ext4 /dev/drbd1
```



```

mke2fs 1.42 (29-Nov-2011)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
63488 inodes, 253935 blocks
12696 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=260046848
8 block groups
32768 blocks per group, 32768 fragments per group
7936 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376

Allocating group tables: done
Writing inode tables: done
Creating journal (4096 blocks): done
Writing superblocks and filesystem accounting information: done

```

Now mount the newly created filesystem so we can create our index file

```

# mount /dev/drbd1 /mnt/
# cat <<-END >/mnt/index.html
<html>
<body>My Test Site - drbd</body>
</html>
END
# umount /dev/drbd1

```

7.4. Configurare il cluster per DRBD

One handy feature of the crm shell is that you can use it in interactive mode to make several changes atomically.

First we launch the shell. The prompt will change to indicate you're in interactive mode.

```

# crm
crm(live) #

```

Next we must create a working copy of the current configuration. This is where all our changes will go. The cluster will not see any of them until we say it's ok. Notice again how the prompt changes, this time to indicate that we're no longer looking at the live cluster.

```

cib crm(live) # cib new drbd
INFO: drbd shadow CIB created
crm(drbd) #

```

Now we can create our DRBD clone and display the revised configuration.

```

crm(drbd) # configure primitive WebData ocf:linbit:drbd params drbd_resource=wwwdata \
op monitor interval=60s
crm(drbd) # configure ms WebDataClone WebData meta master-max=1 master-node-max=1 \
clone-max=2 clone-node-max=1 notify=true
crm(drbd) # configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
params ip="192.168.122.120" cidr_netmask="32" \
op monitor interval="30s"

```

```
primitive WebData ocf:linbit:drbd \  
    params drbd_resource="wwwdata" \  
    op monitor interval="60s" \  
primitive WebSite ocf:heartbeat:apache \  
    params configfile="/etc/httpd/conf/httpd.conf" \  
    op monitor interval="1min" \  
ms WebDataClone WebData \  
    meta master-max="1" master-node-max="1" clone-max="2" clone-node-max="1" \  
    notify="true" \  
location prefer-pcmk-1 WebSite 50: pcmk-1 \  
colocation website-with-ip inf: WebSite ClusterIP \  
order apache-after-ip inf: ClusterIP WebSite \  
property $id="cib-bootstrap-options" \  
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \  
    cluster-infrastructure="corosync" \  
    stonith-enabled="false" \  
    no-quorum-policy="ignore" \  
    last-lrm-refresh="1333446866" \  
rsc_defaults $id="rsc-options" \  
    resource-stickiness="100" \  
op_defaults $id="op-options" \  
    timeout="240s"
```

Once we're happy with the changes, we can tell the cluster to start using them and use `crm_mon` to check everything is functioning.

```
crm(drbd) # cib commit drbd  
INFO: committed 'drbd' shadow CIB to the cluster  
crm(drbd) # quit  
bye  
# crm_mon -1  
=====
```

Last updated: Tue Apr 3 13:50:01 2012
Last change: Tue Apr 3 13:49:46 2012 via crm_shadow on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
4 Resources configured.
=====

Online: [pcmk-1 pcmk-2]

ClusterIP	(ocf::heartbeat:IPaddr2):	Started pcmk-1
WebSite	(ocf::heartbeat:apache):	Started pcmk-1
Master/Slave Set: WebDataClone [WebData]		
Masters: [pcmk-1]		
Slaves: [pcmk-2]		



Nota

TODO: Include details on adding a second DRBD resource

Ora che DRBD sta funzionando è possibile configurare una risorsa filesystem per utilizzarlo. In aggiunta alla definizione del filesystem si necessita inoltre di indicare al cluster dove questa sia localizzata (cioè solo sul nodo Primary di DRBD) e quando ad essa è consentito avviarsi (solo dopo che il nodo Primary è stato attivato).

Once again we'll use the shell's interactive mode

```
# crm
crm(live) # cib new fs
INFO: fs shadow CIB created
crm(fs) # configure primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="ext4"
crm(fs) # configure colocation fs_on_drbd inf: WebFS WebDataClone:Master
crm(fs) # configure order WebFS-after-WebData inf: WebDataClone:promote WebFS:start
```

Vi è inoltre la necessità di indicare al cluster che Apache necessita di funzionare sullo stessa macchina del filesystem e che questo debba essere attivato prima che Apache venga avviato.

```
crm(fs) # configure colocation WebSite-with-WebFS inf: WebSite WebFS
crm(fs) # configure order WebSite-after-WebFS inf: WebFS WebSite
```

Time to review the updated configuration:

```
crm(fs) # configure show
node $id="1702537408" pcmk-1
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebData ocf:linbit:drbd \
    params drbd_resource="wwwdata" \
    op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="ext4"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
ms WebDataClone WebData \
    meta master-max="1" master-node-max="1" clone-max="2" clone-node-max="1"
    notify="true"
location prefer-pcmk-1 WebSite 50: pcmk-1
colocation WebSite-with-WebFS inf: WebSite WebFS
colocation fs_on_drbd inf: WebFS WebDataClone:Master
colocation website-with-ip inf: WebSite ClusterIP
order WebFS-after-WebData inf: WebDataClone:promote WebFS:start
order WebSite-after-WebFS inf: WebFS WebSite
order apache-after-ip inf: ClusterIP WebSite
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
```

Dopo la revisione della nuova configurazione questa va caricata e si potrà osservarla in azione all'interno del cluster.

```
crm(fs) # cib commit fs
INFO: committed 'fs' shadow CIB to the cluster
crm(fs) # quit
bye
# crm_mon -1
=====
Last updated: Tue Apr  3 13:52:21 2012
Last change: Tue Apr  3 13:52:06 2012 via crm_shadow on pcmk-1
```

```
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
5 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-1
WebSite        (ocf::heartbeat:apache):      Started pcmk-1
Master/Slave Set: WebDataClone [WebData]
  Masters: [ pcmk-1 ]
  Slaves: [ pcmk-2 ]
WebFS (ocf::heartbeat:Filesystem):      Started pcmk-1
```

7.4.1. Testare la migrazione

We could shut down the active node again, but another way to safely simulate recovery is to put the node into what is called "standby mode". Nodes in this state tell the cluster that they are not allowed to run resources. Any resources found active there will be moved elsewhere. This feature can be particularly useful when updating the resources' packages.

Mettendo il nodo locale in standby consentirà di osservare il cluster spostare tutte le risorse sull'altro nodo. Da notare inoltre come lo stato del nodo cambierà, indicando che non può più ospitare risorse.

```
# crm node standby
# crm_mon -1
=====
Last updated: Tue Apr  3 13:59:14 2012
Last change: Tue Apr  3 13:52:36 2012 via crm_attribute on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
5 Resources configured.
=====

Node pcmk-1 (1702537408): standby
Online: [ pcmk-2 ]

ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-2
WebSite        (ocf::heartbeat:apache):      Started pcmk-2
Master/Slave Set: WebDataClone [WebData]
  Masters: [ pcmk-2 ]
  Stopped: [ WebData:1 ]
WebFS (ocf::heartbeat:Filesystem):      Started pcmk-2
```

Fatto quanto necessario su pcmk-1 (in questo caso nulla, si è solo osservato le risorse spostarsi), è possibile consentire il nodo ad essere nuovamente parte integrante del cluster.

```
# crm node online
# crm_mon -1
=====
Last updated: Tue Apr  3 14:00:06 2012
Last change: Tue Apr  3 14:00:00 2012 via crm_attribute on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
5 Resources configured.
=====
```

```
Online: [ pcmk-1 pcmk-2 ]
```

```
ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-2  
WebSite        (ocf::heartbeat:apache):        Started pcmk-2  
Master/Slave Set: WebDataClone [WebData]  
  Masters: [ pcmk-2 ]  
  Slaves: [ pcmk-1 ]  
WebFS (ocf::heartbeat:Filesystem):      Started pcmk-2
```

Notare come ora l'impostazione di stickiness delle risorse prevenga la migrazione dei servizi nuovamente verso pcmk-1.

Conversione in Active/Active

Indice

8.1. Requisiti	59
8.1.1. Installing the required Software	59
8.2. Creazione un Filesystem GFS2	60
8.2.1. Preparazione	60
8.2.2. Creazione e popolamento della partizione GFS2	60
8.3. Riconfigurare il cluster per GFS2	63
8.4. Riconfigurare Pacemaker per l'Active/Active	64
8.4.1. Testare il recovery	67

8.1. Requisiti

The primary requirement for an Active/Active cluster is that the data required for your services is available, simultaneously, on both machines. Pacemaker makes no requirement on how this is achieved, you could use a SAN if you had one available, however since DRBD supports multiple Primaries, we can also use that.

The only hitch is that we need to use a cluster-aware filesystem. The one we used earlier with DRBD, ext4, is not one of those. Both OCFS2 and GFS2 are supported, however here we will use GFS2 which comes with Fedora 17.

8.1.1. Installing the required Software

```
# yum install -y gfs2-utils dlm kernel-modules-extra
```

```
Loaded plugins: langpacks, presto, refresh-packagekit
Resolving Dependencies
--> Running transaction check
---> Package dlm.x86_64 0:3.99.4-1.fc17 will be installed
---> Package gfs2-utils.x86_64 0:3.1.4-3.fc17 will be installed
---> Package kernel-modules-extra.x86_64 0:3.4.4-3.fc17 will be installed
--> Finished Dependency Resolution
```

Dependencies Resolved

Package	Arch	Version	Repository	Size
Installing:				
dlm	x86_64	3.99.4-1.fc17	updates	83 k
gfs2-utils	x86_64	3.1.4-3.fc17	fedora	214 k
kernel-modules-extra	x86_64	3.4.4-3.fc17	updates	1.7 M

Transaction Summary

Install 3 Packages

Total download size: 1.9 M

Installed size: 7.7 M

Downloading Packages:

(1/3): dlm-3.99.4-1.fc17.x86_64.rpm	83 kB	00:00
(2/3): gfs2-utils-3.1.4-3.fc17.x86_64.rpm	214 kB	00:00

```
(3/3): kernel-modules-extra-3.4.4-3.fc17.x86_64.rpm | 1.7 MB 00:01
-----
Total 615 kB/s | 1.9 MB 00:03
Running Transaction Check
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : kernel-modules-extra-3.4.4-3.fc17.x86_64 1/3
  Installing : gfs2-utils-3.1.4-3.fc17.x86_64 2/3
  Installing : dlm-3.99.4-1.fc17.x86_64 3/3
  Verifying : dlm-3.99.4-1.fc17.x86_64 1/3
  Verifying : gfs2-utils-3.1.4-3.fc17.x86_64 2/3
  Verifying : kernel-modules-extra-3.4.4-3.fc17.x86_64 3/3

Installed:
  dlm.x86_64 0:3.99.4-1.fc17
  gfs2-utils.x86_64 0:3.1.4-3.fc17
  kernel-modules-extra.x86_64 0:3.4.4-3.fc17

Complete!
```

8.2. Creazione un Filesystem GFS2

8.2.1. Preparazione

Before we do anything to the existing partition, we need to make sure it is unmounted. We do this by telling the cluster to stop the WebFS resource. This will ensure that other resources (in our case, Apache) using WebFS are not only stopped, but stopped in the correct order.

```
# crm resource stop WebFS
# crm_mon -1
=====
Last updated: Tue Apr  3 14:07:36 2012
Last change: Tue Apr  3 14:07:15 2012 via cibadmin on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
5 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-2
Master/Slave Set: WebDataClone [WebData]
Masters: [ pcmk-2 ]
Slaves: [ pcmk-1 ]
```



Nota

Notare come sia Apache che WebFS sono stati fermati.

8.2.2. Creazione e popolamento della partizione GFS2

Ora che il cluster e le sue parti stanno funzionando senza problemi, è possibile creare la partizione GFS2.



Avvertimento

Questa operazione cancellerà qualsiasi contenuto sul device DRBD. Assicurarsi quindi di avere una copia di qualsiasi dato cruciale.

Nella creazione di una partizione GFS2 è necessario specificare diversi parametri aggiuntivi.

Per cominciare, va utilizzata l'opzione `-p` per specificare la volontà di utilizzare il DLM del Kernel. Proseguendo, l'opzione `-j` indica che la risorsa dovrà riservare abbastanza spazio per due journal (uno per ogni nodo che accede al filesystem).

Lastly, we use `-t` to specify the lock table name. The format for this field is **clustername:fsname**. For the **fsname**, we need to use the same value as specified in *corosync.conf* for **cluster_name**. Just pick something unique and descriptive and add somewhere inside the **totem** block. For example:

```
totem {
    version: 2

    # crypto_cipher and crypto_hash: Used for mutual node authentication.
    # If you choose to enable this, then do remember to create a shared
    # secret with "corosync-keygen".
    crypto_cipher: none
    crypto_hash: none
    cluster_name: mycluster
    ...
}
```



Importante

Do this on each node in the cluster and be sure to restart them before continuing.



Importante

We must run the next command on whichever node last had `/dev/drbd` mounted. Otherwise you will receive the message:

```
/dev/drbd1: Read-only file system
```

```
# ssh pcmk-2 -- mkfs.gfs2 -p lock_dlm -j 2 -t mycluster:web /dev/drbd1
This will destroy any data on /dev/drbd1.
It appears to contain: Linux rev 1.0 ext4 filesystem data, UUID=dc45fff3-c47a-4db2-96f7-
a8049a323fe4 (extents) (large files) (huge files)
Are you sure you want to proceed? [y/n]y
Device: /dev/drbd1
Blocksize: 4096
Device Size 0.97 GB (253935 blocks)
```

```
Filesystem Size:          0.97 GB (253932 blocks)
Journals:                2
Resource Groups:         4
Locking Protocol:        "lock_dlm"
Lock Table:              "mycluster"
UUID:                    ed293a02-9eee-3fa3-ed1c-435ef1fd0116
```

```
# crm
crm(live)# cib new dlm
INFO: dlm shadow CIB created
crm(dlm)# configure primitive dlm ocf:pacemaker:controld \
    op monitor interval=60s
crm(dlm)# configure clone dlm_clone dlm meta clone-max=2 clone-node-max=1
crm(dlm)# configure show
node $id="1702537408" pcmk-1 \
    attributes standby="off"
node $id="1719314624" pcmk-2
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.120" cidr_netmask="32" \
    op monitor interval="30s"
primitive WebData ocf:linbit:drbd \
    params drbd_resource="wwwdata" \
    op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="ext4" \
    meta target-role="Stopped"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
primitive dlm ocf:pacemaker:controld \
    op monitor interval="60s"
ms WebDataClone WebData \
    meta master-max="1" master-node-max="1" clone-max="2" clone-node-max="1"
    notify="true"
clone dlm_clone dlm \
    meta clone-max="2" clone-node-max="1"
location prefer-pcmk-1 WebSite 50: pcmk-1
colocation WebSite-with-WebFS inf: WebSite WebFS
colocation fs_on_drbd inf: WebFS WebDataClone:Master
colocation website-with-ip inf: WebSite ClusterIP
order WebFS-after-WebData inf: WebDataClone:promote WebFS:start
order WebSite-after-WebFS inf: WebFS WebSite
order apache-after-ip inf: ClusterIP WebSite
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff" \
    cluster-infrastructure="corosync" \
    stonith-enabled="false" \
    no-quorum-policy="ignore" \
    last-lrm-refresh="1333446866"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
op_defaults $id="op-options" \
    timeout="240s"
crm(dlm)# cib commit dlm
INFO: committed 'dlm' shadow CIB to the cluster
crm(dlm)# quit
bye
# crm_mon -1
=====
Last updated: Wed Apr  4 01:15:11 2012
Last change: Wed Apr  4 00:50:11 2012 via crmd on pcmk-1
Stack: corosync
Current DC: pcmk-1 (1702537408) - partition with quorum
Version: 1.1.7-2.fc17-ee0730e13d124c3d58f00016c3376a1de5323cff
2 Nodes configured, unknown expected votes
7 Resources configured.
```

```

=====
Online: [ pcmk-1 pcmk-2 ]

ClusterIP      (ocf::heartbeat:IPaddr2):      Started pcmk-1
Master/Slave Set: WebDataClone [WebData]
  Masters: [ pcmk-1 ]
  Slaves: [ pcmk-2 ]
Clone Set: dlm_clone [dlm]
  Started: [ pcmk-1 pcmk-2 ]

```

A questo punto è possibile ripopolare il nuovo filesystem con i dati (le pagine web). Per ora verrà creata una versione alternativa dell'home page.

```

# mount /dev/drbd1 /mnt/
# cat <<-END >/mnt/index.html
<html>
<body>My Test Site - GFS2</body>
</html>
END
# umount /dev/drbd1
# drbdadm verify wwwdata

```

8.3. Riconfigurare il cluster per GFS2

```

# crm
crm(live) # cib new GFS2
INFO: GFS2 shadow CIB created
crm(GFS2) # configure delete WebFS
crm(GFS2) # configure primitive WebFS ocf:heartbeat:Filesystem params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"

```

Now that we've recreated the resource, we also need to recreate all the constraints that used it. This is because the shell will automatically remove any constraints that referenced WebFS.

```

crm(GFS2) # configure colocation WebSite-with-WebFS inf: WebSite WebFS
crm(GFS2) # configure colocation fs_on_drbd inf: WebFS WebDataClone:Master
crm(GFS2) # configure order WebFS-after-WebData inf: WebDataClone:promote WebFS:start
crm(GFS2) # configure order WebSite-after-WebFS inf: WebFS WebSite
crm(GFS2) # configure show
node pcmk-1
node pcmk-2
primitive WebData ocf:linbit:drbd \
  params drbd_resource="wwwdata" \
  op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
  params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"
primitive WebSite ocf:heartbeat:apache \
  params configfile="/etc/httpd/conf/httpd.conf" \
  op monitor interval="1min"
primitive ClusterIP ocf:heartbeat:IPaddr2 \
  params ip="192.168.122.101" cidr_netmask="32" \
  op monitor interval="30s"
ms WebDataClone WebData \
  meta master-max="1" master-node-max="1" clone-max="2" clone-node-max="1" notify="true"
colocation WebSite-with-WebFS inf: WebSite WebFS
colocation fs_on_drbd inf: WebFS WebDataClone:Master
colocation website-with-ip inf: WebSite ClusterIP
order WebFS-after-WebData inf: WebDataClone:promote WebFS:start
order WebSite-after-WebFS inf: WebFS WebSite
order apache-after-ip inf: ClusterIP WebSite

```

```
property $id="cib-bootstrap-options" \  
  dc-version="1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f" \  
  cluster-infrastructure="openais" \  
  expected-quorum-votes="2" \  
  stonith-enabled="false" \  
  no-quorum-policy="ignore" \  
rsc_defaults $id="rsc-options" \  
  resource-stickiness="100"
```

Prima di caricare la configurazione nel cluster questa va revisionata, la shell andrà terminata e si dovrà osservare il responso del cluster

```
crm(GFS2) # cib commit GFS2  
INFO: committed 'GFS2' shadow CIB to the cluster  
crm(GFS2) # quit  
bye  
# crm_mon  
=====  
Last updated: Thu Sep 3 20:49:54 2009  
Stack: openais  
Current DC: pcmk-2 - partition with quorum  
Version: 1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f  
2 Nodes configured, 2 expected votes  
6 Resources configured.  
=====  
  
Online: [ pcmk-1 pcmk-2 ]  
  
WebSite (ocf::heartbeat:apache): Started pcmk-2  
Master/Slave Set: WebDataClone  
  Masters: [ pcmk-1 ]  
  Slaves: [ pcmk-2 ]  
ClusterIP (ocf::heartbeat:IPaddr): Started pcmk-2  
WebFS (ocf::heartbeat:Filesystem):  
  Started pcmk-1
```

8.4. Riconfigurare Pacemaker per l'Active/Active

Quasi tutto è a posto. Le versioni recenti di DRBD supportano l'operare in modalità Primary/Primary ed inoltre il filesystem utilizzato è cluster aware. Tutto ciò che rimane da fare è configurare nuovamente il cluster per sfruttare queste peculiarità.

This will involve a number of changes, so we'll again use interactive mode.

```
# crm  
# cib new active
```

There's no point making the services active on both locations if we can't reach them, so let's first clone the IP address. Cloned IPaddr2 resources use an iptables rule to ensure that each request only gets processed by one of the two clone instances. The additional meta options tell the cluster how many instances of the clone we want (one "request bucket" for each node) and that if all other nodes fail, then the remaining node should hold all of them. Otherwise the requests would be simply discarded.

```
# configure clone WebIP ClusterIP \  
  meta globally-unique="true" clone-max="2" clone-node-  
  max="2"
```

Ora bisogna indicare a ClusterIP come decidere quali richieste sono processate e da quali host. Per fare ciò è necessario definire il parametro `clusterip_hash`.

Open the ClusterIP resource

```
# configure edit ClusterIP
```

And add the following to the params line

```
clusterip_hash="sourceip"
```

So that the complete definition looks like:

```
primitive ClusterIP ocf:heartbeat:IPaddr2 \
  params ip="192.168.122.101" cidr_netmask="32" clusterip_hash="sourceip" \
  op monitor interval="30s"
```

Here is the full transcript

```
# crm crm(live)
# cib new active
INFO: active shadow CIB created
crm(active) # configure clone WebIP ClusterIP \
  meta globally-unique="true" clone-max="2" clone-node-max="2"
crm(active) # configure shownode pcmk-1
node pcmk-2
primitive WebData ocf:linbit:drbd \
  params drbd_resource="wwwdata" \
  op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
  params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"
primitive WebSite ocf:heartbeat:apache \
  params configfile="/etc/httpd/conf/httpd.conf" \
  op monitor interval="1min"
primitive ClusterIP ocf:heartbeat:IPaddr2 \
  params ip="192.168.122.101" cidr_netmask="32" clusterip_hash="sourceip" \
  op monitor interval="30s"
ms WebDataClone WebData \
  meta master-max="1" master-node-max="1" clone-max="2" clone-node-max="1" notify="true"
clone WebIP ClusterIP \
  meta globally-unique="true" clone-max="2" clone-node-max="2"
colocation WebSite-with-WebFS inf: WebSite WebFS
colocation fs_on_drbd inf: WebFS WebDataClone:Master
colocation website-with-ip inf: WebSite WebIPorder WebFS-after-WebData inf:
  WebDataClone:promote WebFS:start
order WebSite-after-WebFS inf: WebFS WebSiteorder apache-after-ip inf: WebIP WebSite
property $id="cib-bootstrap-options" \
  dc-version="1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f" \
  cluster-infrastructure="openais" \
  expected-quorum-votes="2" \
  stonith-enabled="false" \
  no-quorum-policy="ignore"
rsc_defaults $id="rsc-options" \
  resource-stickiness="100"
```

Notice how any constraints that referenced ClusterIP have been updated to use WebIP instead. This is an additional benefit of using the crm shell.

Next we need to convert the filesystem and Apache resources into clones.

Again, the shell will automatically update any relevant constraints.

```
crm(active) # configure clone WebFSClone WebFS
crm(active) # configure clone WebSiteClone WebSite
```

Capitolo 8. Conversione in Active/Active

L'ultimo passo è quello di comunicare al cluster che ora è consentito promuovere entrambe le istanze a PPrimary (o Master).

```
crm(active) # configure edit WebDataClone
```

Change master-max to 2

```
crm(active) # configure show
node pcmk-1
node pcmk-2
primitive WebData ocf:linbit:drbd \
    params drbd_resource="wwwdata" \
    op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.101" cidr_netmask="32" clusterip_hash="sourceip" \
    op monitor interval="30s"
ms WebDataClone WebData \
    meta master-max="2" master-node-max="1" clone-max="2" clone-node-max="1" notify="true"
clone WebFSClone WebFSClone WebIP ClusterIP \
    meta globally-unique="true" clone-max="2" clone-node-max="2"
clone WebSiteClone WebSitecolocation WebSite-with-WebFS inf: WebSiteClone WebFSClone
colocation fs_on_drbd inf: WebFSClone WebDataClone:Master
colocation website-with-ip inf: WebSiteClone WebIP
order WebFS-after-WebData inf: WebDataClone:promote WebFSClone:start
order WebSite-after-WebFS inf: WebFSClone WebSiteClone
order apache-after-ip inf: WebIP WebSiteClone
property $id="cib-bootstrap-options" \
    dc-version="1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f" \
    cluster-infrastructure="openais" \
    expected-quorum-votes="2" \
    stonith-enabled="false" \
    no-quorum-policy="ignore"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
```

Prima di caricare la configurazione nel cluster questa va revisionata, la shell andrà terminata e si dovrà osservare il responso del cluster

```
crm(active) # cib commit active
INFO: committed 'active' shadow CIB to the cluster
crm(active) # quit
bye
# crm_mon
=====
Last updated: Thu Sep 3 21:37:27 2009
Stack: openais
Current DC: pcmk-2 - partition with quorum
Version: 1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f
2 Nodes configured, 2 expected votes
6 Resources configured.
=====

Online: [ pcmk-1 pcmk-2 ]

Master/Slave Set: WebDataClone
    Masters: [ pcmk-1 pcmk-2 ]
Clone Set: WebIP Started: [ pcmk-1 pcmk-2 ]
Clone Set: WebFSClone Started: [ pcmk-1 pcmk-2 ]
Clone Set: WebSiteClone Started: [ pcmk-1 pcmk-2 ]
```

```
Clone Set: dlm_clone Started: [ pcmk-1 pcmk-2 ]
```

8.4.1. Testare il recovery



Nota

TODO: mettere un nodo in standby per dimostrare il failover

Configurare STONITH

Indice

9.1. What Is STONITH	69
9.2. Quale device STONITH andrebbe utilizzato	69
9.3. Configurare STONITH	69
9.4. Esempio	70

9.1. What Is STONITH

STONITH is an acronym for Shoot-The-Other-Node-In-The-Head and it protects your data from being corrupted by rogue nodes or concurrent access.

Solo perché un nodo non è responsivo non significa che non stia accedendo ai dati. L'unica via per essere sicuri al 100% che i dati sono al sicuro è quella di utilizzare STONITH in modo da essere certi che il nodo sia effettivamente offline, prima di consentire ad altri nodi di accedere ai dati.

STONITH gioca un ruolo importante anche nel caso in cui un servizio clusterizzato non può essere stoppato. In questo caso il cluster utilizza STONITH per forzare l'intero nodo offline, rendendo di fatto sicuro l'avvio del servizio altrove.

9.2. Quale device STONITH andrebbe utilizzato

E' cruciale che il device STONITH sia in grado di differenziare un guasto del nodo da un guasto di rete.

The biggest mistake people make in choosing a STONITH device is to use remote power switch (such as many on-board IMPI controllers) that shares power with the node it controls. In such cases, the cluster cannot be sure if the node is really offline, or active and suffering from a network fault.

Likewise, any device that relies on the machine being active (such as SSH-based "devices" used during testing) are inappropriate.

9.3. Configurare STONITH

1. Find the correct driver: **stonith_admin --list-installed**
2. Since every device is different, the parameters needed to configure it will vary. To find out the parameters associated with the device, run: **stonith_admin --metadata --agent type**

The output should be XML formatted text containing additional parameter descriptions. We will endeavor to make the output more friendly in a later version.

3. Enter the shell crm Create an editable copy of the existing configuration **cib new stonith**
Create a fencing resource containing a primitive resource with a class of stonith, a type of type and a parameter for each of the values returned in step 2: **configure primitive ...**
4. If the device does not know how to fence nodes based on their uname, you may also need to set the special **pcmk_host_map** parameter. See **man stonithd** for details.

5. If the device does not support the list command, you may also need to set the special **pcmk_host_list** and/or **pcmk_host_check** parameters. See **man stonithd** for details.
6. If the device does not expect the victim to be specified with the port parameter, you may also need to set the special **pcmk_host_argument** parameter. See **man stonithd** for details.
7. Upload it into the CIB from the shell: **cib commit stonith**
8. Once the stonith resource is running, you can test it by executing: **stonith_admin --reboot nodename**. Although you might want to stop the cluster on that machine first.

9.4. Esempio

Assuming we have an chassis containing four nodes and an IPMI device active on 10.0.0.1, then we would chose the fence_ipmilan driver in step 2 and obtain the following list of parameters

Obtaining a list of STONITH Parameters

```
# stonith_admin --metadata -a fence_ipmilan
```

```
<?xml version="1.0" ?>
<resource-agent name="fence_ipmilan" shortdesc="Fence agent for IPMI over LAN">
<longdesc>
fence_ipmilan is an I/O Fencing agent which can be used with machines controlled by IPMI.
This agent calls support software using ipmitool (http://ipmitool.sf.net/).

To use fence_ipmilan with HP iLO 3 you have to enable lanplus option (lanplus / -P) and
increase wait after operation to 4 seconds (power_wait=4 / -T 4)</longdesc>
<parameters>
  <parameter name="auth" unique="1">
    <getopt mixed="-A" />
    <content type="string" />
    <shortdesc>IPMI Lan Auth type (md5, password, or none)</shortdesc>
  </parameter>
  <parameter name="ipaddr" unique="1">
    <getopt mixed="-a" />
    <content type="string" />
    <shortdesc>IPMI Lan IP to talk to</shortdesc>
  </parameter>
  <parameter name="passwd" unique="1">
    <getopt mixed="-p" />
    <content type="string" />
    <shortdesc>Password (if required) to control power on IPMI device</shortdesc>
  </parameter>
  <parameter name="passwd_script" unique="1">
    <getopt mixed="-S" />
    <content type="string" />
    <shortdesc>Script to retrieve password (if required)</shortdesc>
  </parameter>
  <parameter name="lanplus" unique="1">
    <getopt mixed="-P" />
    <content type="boolean" />
    <shortdesc>Use Lanplus</shortdesc>
  </parameter>
  <parameter name="login" unique="1">
    <getopt mixed="-l" />
    <content type="string" />
    <shortdesc>Username/Login (if required) to control power on IPMI device</
shortdesc>
  </parameter>
  <parameter name="action" unique="1">
```

```

        <getopt mixed="-o" />
        <content type="string" default="reboot"/>
        <shortdesc>Operation to perform. Valid operations: on, off, reboot, status,
list, diag, monitor or metadata</shortdesc>
    </parameter>
    <parameter name="timeout" unique="1">
        <getopt mixed="-t" />
        <content type="string" />
        <shortdesc>Timeout (sec) for IPMI operation</shortdesc>
    </parameter>
    <parameter name="cipher" unique="1">
        <getopt mixed="-C" />
        <content type="string" />
        <shortdesc>Ciphersuite to use (same as ipmitool -C parameter)</shortdesc>
    </parameter>
    <parameter name="method" unique="1">
        <getopt mixed="-M" />
        <content type="string" default="onoff"/>
        <shortdesc>Method to fence (onoff or cycle)</shortdesc>
    </parameter>
    <parameter name="power_wait" unique="1">
        <getopt mixed="-T" />
        <content type="string" default="2"/>
        <shortdesc>Wait X seconds after on/off operation</shortdesc>
    </parameter>
    <parameter name="delay" unique="1">
        <getopt mixed="-f" />
        <content type="string" />
        <shortdesc>Wait X seconds before fencing is started</shortdesc>
    </parameter>
    <parameter name="verbose" unique="1">
        <getopt mixed="-v" />
        <content type="boolean" />
        <shortdesc>Verbose mode</shortdesc>
    </parameter>
</parameters>
<actions>
    <action name="on" />
    <action name="off" />
    <action name="reboot" />
    <action name="status" />
    <action name="diag" />
    <action name="list" />
    <action name="monitor" />
    <action name="metadata" />
</actions>
</resource-agent>

```

from which we would create a STONITH resource fragment that might look like this

Sample STONITH Resource

```

# crm crm(live)# cib new stonith
INFO: stonith shadow CIB created
crm(stonith)# configure primitive impi-fencing stonith::fence_ipmilan \
  params pcmk_host_list="pcmk-1 pcmk-2" ipaddr=10.0.0.1 login=testuser passwd=abc123 \
  op monitor interval="60s"

```

And finally, since we disabled it earlier, we need to re-enable STONITH. At this point we should have the following configuration.

Now push the configuration into the cluster.

```

crm(stonith)# configure property stonith-enabled="true"

```

```
crm(stonith)# configure shownode pcmk-1
node pcmk-2
primitive WebData ocf:linbit:drbd \
    params drbd_resource="wwwdata" \
    op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.101" cidr_netmask="32" clusterip_hash="sourceip" \
    op monitor interval="30s" primitive ipmi-fencing stonith::fence_ipmilan \ params
pcmk_host_list="pcmk-1 pcmk-2" ipaddr="10.0.0.1 login=testuser passwd=abc123 \ op monitor
interval="60s"ms WebDataClone WebData \
    meta master-max="2" master-node-max="1" clone-max="2" clone-node-max="1" notify="true"
clone WebFSClone WebFS
clone WebIP ClusterIP \
    meta globally-unique="true" clone-max="2" clone-node-max="2"
clone WebSiteClone WebSite
colocation WebSite-with-WebFS inf: WebSiteClone WebFSClone
colocation fs_on_drbd inf: WebFSClone WebDataClone:Master
colocation website-with-ip inf: WebSiteClone WebIP
order WebFS-after-WebData inf: WebDataClone:promote WebFSClone:start
order WebSite-after-WebFS inf: WebFSClone WebSiteClone
order apache-after-ip inf: WebIP WebSiteClone
property $id="cib-bootstrap-options" \
    dc-version="1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f" \
    cluster-infrastructure="openais" \
    expected-quorum-votes="2" \
    stonith-enabled="true" \
    no-quorum-policy="ignore"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
crm(stonith)# cib commit stonithINFO: committed 'stonith' shadow CIB to the cluster
crm(stonith)# quit
bye
```

Appendice A. Riepilogo delle configurazioni

Indice

A.1. Configurazione finale del cluster	73
A.2. Lista nodi	74
A.3. Opzioni del cluster	74
A.4. Risorse	74
A.4.1. Opzioni di default	74
A.4.2. Fencing	74
A.4.3. Servizio Address	74
A.4.4. DRBD - Storage condiviso	75
A.4.5. Cluster Filesystem	75
A.4.6. Apache	75

A.1. Configurazione finale del cluster

```
# crm configure show
node pcmk-1
node pcmk-2
primitive WebData ocf:linbit:drbd \
    params drbd_resource="wwwdata" \
    op monitor interval="60s"
primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
    op monitor interval="1min"
primitive ClusterIP ocf:heartbeat:IPAddr2 \
    params ip="192.168.122.101" cidr_netmask="32" clusterip_hash="sourceip" \
    op monitor interval="30s"
primitive ipmi-fencing stonith::fence_ipmilan \
    params pcmk_host_list="pcmk-1 pcmk-2" ipaddr=10.0.0.1 login=testuser passwd=abc123 \
    op monitor interval="60s"
ms WebDataClone WebData \
    meta master-max="2" master-node-max="1" clone-max="2" clone-node-max="1" notify="true"
clone WebFSClone WebFS
clone WebIP ClusterIP \
    meta globally-unique="true" clone-max="2" clone-node-max="2"
clone WebSiteClone WebSite
colocation WebSite-with-WebFS inf: WebSiteClone WebFSClone
colocation fs_on_drbd inf: WebFSClone WebDataClone:Master
colocation website-with-ip inf: WebSiteClone WebIP
order WebFS-after-WebData inf: WebDataClone:promote WebFSClone:start
order WebSite-after-WebFS inf: WebFSClone WebSiteClone
order apache-after-ip inf: WebIP WebSiteClone
property $id="cib-bootstrap-options" \
    dc-version="1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f" \
    cluster-infrastructure="openais" \
    expected-quorum-votes="2" \
    stonith-enabled="true" \
    no-quorum-policy="ignore"
rsc_defaults $id="rsc-options" \
    resource-stickiness="100"
```

A.2. Lista nodi

La lista dei nodi è popolata automaticamente dal cluster.

```
node pcmk-1
node pcmk-2
```

A.3. Opzioni del cluster

Qui è dove il cluster registra automaticamente alcune informazioni in merito al cluster

- dc-version - la versione (incluso l'hash del codice sorgente originale) di Pacemaker usata nel DC
- cluster-infrastructure - l'infrastruttura cluster utilizzata (heartbeat or openais)
- expected-quorum-votes - il numero massimo di nodi che ci si aspetta facciano parte del cluster

e dove l'amministratore può assegnare valori alle opzioni che controllano il modo in cui il cluster opera

- stonith-enabled=true - Fai uso di STONITH
- no-quorum-policy=ignore - Ignora la perdita di quorum e continua ad ospitare le risorse.

```
property $id="cib-bootstrap-options" \  
  dc-version="1.1.5-bdd89e69ba545404d02445be1f3d72e6a203ba2f" \  
  cluster-infrastructure="openais" \  
  expected-quorum-votes="2" \  
  stonith-enabled="true" \  
  no-quorum-policy="ignore"
```

A.4. Risorse

A.4.1. Opzioni di default

Qui vengono configurate le opzioni del cluster che vanno applicati a tutte le risorse

- resource-stickiness - Specifica l'impossibilità o meno di muovere risorse ad altre macchine

```
rsc_defaults $id="rsc-options" \  
  resource-stickiness="100"
```

A.4.2. Fencing

```
primitive ipmi-fencing stonith::fence_ipmilan \  
  params pcmk_host_list="pcmk-1 pcmk-2" ipaddr=10.0.0.1 login=testuser passwd=abc123 \  
  op monitor interval="60s" \  
clone Fencing rsa-fencing
```

A.4.3. Servizio Address

Users of the services provided by the cluster require an unchanging address with which to access it. Additionally, we cloned the address so it will be active on both nodes. An iptables rule (created as part of the resource agent) is used to ensure that each request only gets processed by one of the two

clone instances. The additional meta options tell the cluster that we want two instances of the clone (one "request bucket" for each node) and that if one node fails, then the remaining node should hold both.

```
primitive ClusterIP ocf:heartbeat:IPaddr2 \
    params ip="192.168.122.101" cidr_netmask="32" clusterip_hash="sourceip" \
    op monitor interval="30s"
clone WebIP ClusterIP
    meta globally-unique="true" clone-max="2" clone-node-max="2"
```



Nota

TODO: Il RA quando clonato dovrebbe controllare l'opzione globally-unique=true

A.4.4. DRBD - Storage condiviso

Qui viene definito il servizio DRBD e specificata quale risorsa DRBD (da drbd.conf) questi debba controllare. La risorsa viene definita come master/slave e, per rispettare il setup active/active, entrambe le istanze vengono abilitate ad essere promosse specificando master-max=2. Viene valorizzata inoltre l'opzione notify, così che il cluster comunicherà all'agent DRBD quando il suo nodo cambierà stato.

```
primitive WebData ocf:linbit:drbd \
    params drbd_resource="wwwdata" \
    op monitor interval="60s"
ms WebDataClone WebData \
    meta master-max="2" master-node-max="1" clone-max="2" clone-node-max="1" notify="true"
```

A.4.5. Cluster Filesystem

Il Cluster Filesystem si assicura che i file vengano letti e scritti nella maniera corretta. E' necessario specificare il block device (fornito da DRBD), dove si vuole effettuare l'operazione di mount e che viene utilizzato GFS2. Di nuovo questo è un clone, perché è inteso essere attivo su entrambi i nodi. La constraint aggiuntiva assicura che la risorsa possa essere avviata su nodi con gfs-control attivo e istanze drbd.

```
primitive WebFS ocf:heartbeat:Filesystem \
    params device="/dev/drbd/by-res/wwwdata" directory="/var/www/html" fstype="gfs2"
clone WebFSClone WebFS
colocation WebFS-with-gfs-control inf: WebFSClone gfs-clone
colocation fs_on_drbd inf: WebFSClone WebDataClone:Master
order WebFS-after-WebData inf: WebDataClone:promote WebFSClone:start
order start-WebFS-after-gfs-control inf: gfs-clone WebFSClone
```

A.4.6. Apache

Infine viene definito il servizio Apache. E' necessario solamente specificare al cluster dove trovare il file di configurazione principale e costringere questo ad essere eseguito su nodi con il filesystem richiesto montato e l'indirizzo IP attivo.

```
primitive WebSite ocf:heartbeat:apache \
    params configfile="/etc/httpd/conf/httpd.conf" \
```

Appendice A. Riepilogo delle configurazioni

```
    op monitor interval="1min"
clone WebSiteClone WebSite
colocation WebSite-with-WebFS inf: WebSiteClone WebFSClone
colocation website-with-ip inf: WebSiteClone WebIP
order apache-after-ip inf: WebIP WebSiteClone
order WebSite-after-WebFS inf: WebFSClone WebSiteClone
```

Appendice B. Sample Corosync Configuration

Sample Corosync.conf for a two-node cluster using multicast.

```
# Please read the corosync.conf.5 manual page
totem {
    version: 2

    # crypto_cipher and crypto_hash: Used for mutual node authentication.
    # If you choose to enable this, then do remember to create a shared
    # secret with "corosync-keygen".
    crypto_cipher: none
    crypto_hash: none

    # interface: define at least one interface to communicate
    # over. If you define more than one interface stanza, you must
    # also set rrp_mode.
    interface {
        # Rings must be consecutively numbered, starting at 0.
        ringnumber: 0
        # This is normally the *network* address of the
        # interface to bind to. This ensures that you can use
        # identical instances of this configuration file
        # across all your cluster nodes, without having to
        # modify this option.
bindnetaddr: 192.168.122.0
        # However, if you have multiple physical network
        # interfaces configured for the same subnet, then the
        # network address alone is not sufficient to identify
        # the interface Corosync should bind to. In that case,
        # configure the *host* address of the interface
        # instead:
        # bindnetaddr: 192.168.1.1
        # When selecting a multicast address, consider RFC
        # 2365 (which, among other things, specifies that
        # 239.255.x.x addresses are left to the discretion of
        # the network administrator). Do not reuse multicast
        # addresses across multiple Corosync clusters sharing
        # the same network.
mcastaddr: 239.255.1.1
        # Corosync uses the port you specify here for UDP
        # messaging, and also the immediately preceding
        # port. Thus if you set this to 5405, Corosync sends
        # messages over UDP ports 5405 and 5404.
mcastport: 4000
        # Time-to-live for cluster communication packets. The
        # number of hops (routers) that this ring will allow
        # itself to pass. Note that multicast routing must be
        # specifically enabled on most network routers.
        ttl: 1
    }
}

logging {
    # Log the source file and line where messages are being
    # generated. When in doubt, leave off. Potentially useful for
    # debugging.
    fileline: off
    # Log to standard error. When in doubt, set to no. Useful when
    # running in the foreground (when invoking "corosync -f")
    to_stderr: no
}
```

```
# Log to a log file. When set to "no", the "logfile" option
# must not be set.
to_logfile: yes
logfile: /var/log/cluster/corosync.log
# Log to the system log daemon. When in doubt, set to yes.
to_syslog: yes
# Log debug messages (very verbose). When in doubt, leave off.
debug: off
# Log messages with time stamps. When in doubt, set to on
# (unless you are only logging to syslog, where double
# timestamps can be annoying).
timestamp: on
logger_subsys {
    subsys: QUORUM
    debug: off
}
}

quorum {
    provider: corosync_votequorum
    expected_votes: 2
}
```

Appendice C. Approfondimenti

- Project Website <http://www.clusterlabs.org>
- Cluster Commands A comprehensive guide to cluster commands has been written by SuSE and can be found at: http://www.suse.com/documentation/sle_ha/book_sleha/?page=/documentation/sle_ha/book_sleha/data/book_sleha.html
- Corosync <http://www.corosync.org>

Appendice D. Revision History

Revisione 1-0 Mon May 17 2010

Import from Pages.app

Andrew Beekhof andrew@beekhof.net

Revisione 2-0 Wed Sep 22 2010

Italian translation

Raoul Scarazzini

rasca@miamammausainux.org

Revisione 3-0 Wed Feb 9 2011

Updated for Fedora 13

Andrew Beekhof andrew@beekhof.net

Revisione 4-0 Wed Oct 5 2011

Update the GFS2 section to use CMAN

Andrew Beekhof andrew@beekhof.net

Revisione 5-0 Fri Feb 10 2012

Generate docbook content from asciidoc sources

Andrew Beekhof andrew@beekhof.net

Revisione 6-0 Tues July 3 2012

Updated for Fedora 17

Andrew Beekhof andrew@beekhof.net

Revisione 7-0 Fri Sept 14 2012

Updated for pcs

David Vossel dvossel@redhat.com

Indice analitico

C

Creating and Activating a new SSH Key, 14

D

Domain name (Query), 11

Domain name (Remove from host name), 12

F

feedback

contact information for this manual, xii

N

Nodes

Domain name (Query), 11

Domain name (Remove from host name), 12

short name, 11

S

short name, 11

SSH, 13

