



RHCS and Pacemaker in Drupal Hosting

Marek Salwerowicz - IT Department, Web Frameworks Section
`marek.salwerowicz@cern.ch`



Agenda

Drupal Service Presentation

Red Hat Cluster Suite (HA for RHEL 6)

Corosync and Pacemaker (HA for RHEL 7)

Conclusion, Q&A

Agenda

Drupal Service Presentation

Red Hat Cluster Suite (HA for RHEL 6)

Corosync and Pacemaker (HA for RHEL 7)

Conclusion, Q&A

Drupal Service

Purpose

Provide infrastructure for building sites and web applications based on Drupal Content Management System

Statistics

- 1020+ sites
- 300 HTTP requests/second

Drupal Service

Main sites

- CERN Home site
- LHC Experiments' (ATLAS, ALICE, CMS, LHCb) sites
- Press Office site
- HR sites (internal, recruitment)

Drupal Service

Hardware

- workers: 8 servers (2 quads)
 - 2x Intel Xeon CPU
 - 64 GB RAM
 - 2x 10 Gbit network interfaces
- Storage:
 - NetAPP FAS 3210
 - 2 controllers in 7-mode
 - 20 TB storage space (35 % used)
 - 4x 10Gbit network interfaces
 - NFS and iSCSI

Drupal Service

Network

- 2x 10Gbit switches
- 2x independent 10 Gbit uplinks to CERN Network
- IPv4 + IPv6

Software

- Scientific Linux CERN 6
- Red Hat High Availability Add-On
- Puppet Configuration Management System

Agenda

Drupal Service Presentation

Red Hat Cluster Suite (HA for RHEL 6)

Corosync and Pacemaker (HA for RHEL 7)

Conclusion, Q&A

RHCS – cluster.conf file

Logical description of the whole infrastructure:

- Cluster nodes, controlled by IPMI
- iSCSI quorum disk configuration
- Services (memcache, apache, rsyslog...) startup control
- Failover domains configured
- Floating IP addresses allocation

RHCS – cluster overview (1/2)

Cluster Status for drupal @ Mon Jan 16 14:16:16 2017
Member Status: Quorate

<u>Member Name</u>	<u>ID</u>	<u>Status</u>
drupal11-storage.cern.ch	1	Online , Local , rgmanager
drupal12-storage.cern.ch	2	Online , rgmanager
drupal13-storage.cern.ch	3	Online , rgmanager
drupal14-storage.cern.ch	4	Online , rgmanager
drupal21-storage.cern.ch	5	Online , rgmanager
drupal22-storage.cern.ch	6	Online , rgmanager
drupal23-storage.cern.ch	7	Online , rgmanager
drupal24-storage.cern.ch	8	Online , rgmanager
/dev/sdc	0	Online , Quorum Disk

RHCS – cluster overview (2/2)

Service Name	Owner (Last)	State
service : drupal1b01	drupal11-storage.cern.ch	started
service : drupal1b02	drupal12-storage.cern.ch	started
service : drupal1b03	drupal11-storage.cern.ch	started
service : drupal11	drupal11-storage.cern.ch	started
service : drupal12	drupal12-storage.cern.ch	started
service : drupal13	drupal13-storage.cern.ch	started
service : drupal14	drupal14-storage.cern.ch	started
service : drupal21	drupal21-storage.cern.ch	started
service : drupal22	drupal22-storage.cern.ch	started
service : drupal23	drupal23-storage.cern.ch	started
service : drupal24	drupal24-storage.cern.ch	started
service : drupal5b	drupal24-storage.cern.ch	started
service : drupalws	drupal24-storage.cern.ch	started
service : dsmcad	drupal24-storage.cern.ch	started
service : filebeat	drupal24-storage.cern.ch	started
service : rsyslog2	drupal24-storage.cern.ch	started

RHCS – advantages

- Single Point of Service Control (`clustat` command)
- Allows transparent maintenance of hardware / software installed
- Reliable Quorum disk with $n-1$ votes (n = number of nodes) – split-brain scenario mitigation, ability to provide service with one node up.
- `cluster.conf` easily maintained and versioned (using Puppet and Git behind)

RHCS – limitations

- Corosync Totem protocol sensitive to network disruption
- Floating IP address re-allocation takes around 10 seconds

Agenda

Drupal Service Presentation

Red Hat Cluster Suite (HA for RHEL 6)

Corosync and Pacemaker (HA for RHEL 7)

Conclusion, Q&A

Corosync and Pacemaker

- Drupal 8 hosting : future release of the service (currently under development)
- CERN CentOS 7
- RHCS-like use case

Corosync

corosync.conf

- Nodes definition
- Totem protocol setup
(network / cryptography)
- Quorum setup: external quorum service since RHEL7.3, but **missing** quorum disk

Corosync

quorum device

- In SLC6 common iSCSI drive used successfully as quorum device
- HA for RHEL 7 offers quorum TCP service – how to implement it correctly ?
- Quorum disk model to be ported soon ?^a

^aCorosync developers claim it is "going to happen"

Pacemaker

- Services description (cluster.conf-like)
- Fencing (STONITH) service
- Resources group status – detailed pcs status output
(next slide)

```
# pcs status
```

```
[..]
```

```
Resource Group: drupal8dev11-worker
```

```
drupal8dev11-rhcs.cern.ch : Started drupal8dev11-storage
```

```
drupal8dev11_drupal2etc   : Started drupal8dev11-storage
```

```
drupal8dev11_drupal2tmpfs : Started drupal8dev11-storage
```

```
drupal8dev11_memcached    : Started drupal8dev11-storage
```

```
drupal8dev11_apache       : Started drupal8dev11-storage
```

```
[..]
```

Pacemaker

Initial configuration done using `pcs` tool.

How to control configuration file using Puppet ?

clusterlabs.org Pacemaker documentation

Q: How Should the Configuration be Updated?

*A: Rule 1 - **Never** edit the `cib.xml` file manually. Ever. I'm not making this up.*

Rule 2 - Read Rule 1 again.

Rule 3 - The cluster will notice if you ignored rules 1 & 2 and refuse to use the configuration.

Pacemaker

- cibadmin tool – export / import xml (parts of)config file
- config file verification does not work
https://bugzilla.redhat.com/show_bug.cgi?id=1213946
- Syntax: XML config files **sensitive** to whitespaces, newlines etc.

HA for RHEL7 wishlist

- Quorum iSCSI disk
- Easily maintained configuration
- Resources failure count – no restart is possible without `cleanup` operation
- `cleanup` operation results in occasional cluster failures – timeouts tuning ?

Agenda

Drupal Service Presentation

Red Hat Cluster Suite (HA for RHEL 6)

Corosync and Pacemaker (HA for RHEL 7)

Conclusion, Q&A

Conclusion, Q&A

- HA for RHEL 6 makes clustered Web Service administration easier.
- HA for RHEL 7 has big potential, but misses essential functionalities (quorum disk, config file verification and versioning)

Thank you!



www.cern.ch