Pacemaker

Recent changes and future plans

ClusterLabs Summit 2020

Ken Gaillot <kgaillot@redhat.com>
Senior Software Engineer, Red Hat

Recent changes (2.0.1 - 2.0.3)

Coming soon (2.0.4 - 2.0.5)

Future directions

Fence history display

Synchronized across all nodes

Failures/pending shown in crm_mon by default

```
# crm_mon -1 --fence-history=3
... <snip> ...
Fencing History:
  * reboot of rhel8-3 successful: delegate=rhel8-1, client=stonith_admin.5516,
origin=rhel8-4, completed='2020-01-20 13:06:40 -06:00'
  * reboot of rhel8-3 successful: delegate=rhel8-1, client=stonith_admin.1633,
origin=rhel8-4, completed='2020-01-20 12:05:43 -06:00'
```

Can be erased manually

```
stonith_admin -H '*' --cleanup
```

Questions?

Dynamic recheck interval

cluster-recheck-interval

Applies to failure-timeout and date-based rules
(date_expression with an operation of gt, lt, or in_range but not date_spec)

Questions?

Pacemaker Remote hardening

Start-up environment variables (/etc/sysconfig, /etc/default)

Listen address and port
PCMK_remote_address="192.0.2.1"
PCMK_remote_port=3121

TLS priority preferences
(see https://gnutls.org/manual/html_node/Priority-Strings.html)
PCMK_tls_priorities="NORMAL:-VERS-SSL3.0:-VERS-TLS1.0:-MD5:-3DES-CBC"

TLS Diffie-Hellman prime length
PCMK_dh_min_bits=1024
PCMK_dh_max_bits=2048

Questions?

Bundle enhancements

- Support for Docker, podman, and rkt
- Can be used in cluster with sbd
- Per-node environment variables for Pacemaker Remote

```
<bundle id="httpd-bundle">
  <podman image="pcmk:http" replicas="3" options=”-e MY_VAR=true”/>
  <storage>
    <storage-mapping id="httpd-env" options=”r”
      source-dir="/srv/httpd-files/pacemaker-environment-vars"
      target-dir="/etc/pacemaker/pcmk-init.env" />
  </storage>
  <primitive class="ocf" id="httpd" provider="heartbeat" type="apache"/>
</bundle>
```

Questions?

Machine-friendly tool output

--output-as/--output-to

```
# stonith_admin --list-installed --output-as=xml
<pacemaker-result api-version="2.0" request="stonith_admin --list-installed --output-as=xml">
  <list name="Installed fence devices" count="3">
        <item name="device">fence_virt</item>
        <item name="device">fence_virtd</item>
        <item name="device">fence_xvm</item>
  </list>
  <status code="0" message="OK"/>
</pacemaker-result>
```

API XML schema

Output formats can take options (such as HTML stylesheet)

```
crm_mon -1 --output-as=html
--html-stylesheet="https://example.com/css/main.css"
```

File Edit View History Bookmarks Tools Help

Cluster Status

file:///home/kgaillot/cluster-status.html

Search

Cluster Summary

- **Stack:** corosync
- **Current DC:** rhel8-4 (version 2.0.3-4.el8-4b1f869f0f) - partition with quorum
- **Last updated:** Mon Jan 20 12:07:20 2020
- **Last change:** Mon Jan 20 12:05:00 2020 by root via cibadmin on rhel8-2
- 5 nodes configured
- 3 resource instances configured

Config Options

- STONITH of failed nodes enabled
- Cluster is symmetric
- No Quorum policy: Stop ALL resources
- Resource management enabled

Node List

- Node: rhel8-1 online
- Node: rhel8-2 online
- Node: rhel8-3 OFFLINE
- Node: rhel8-4 online
- Node: rhel8-5 online

Active Resources

- Fencing (stonith:fence_xvm): Started rhel8-1
- rsc1 (ocf::pacemaker:Dummy): Started rhel8-4
- rsc2 (ocf::pacemaker:Dummy): Started rhel8-5

Failed Resource Actions

- rsc1_monitor_10000 on rhel8-4 'not running' (7): call=23, status='complete', exitreason='', last-rc-change='2020-01-20 12:05:19 -06:00', queued=0ms, exec=0ms

Fencing History

- reboot of rhel8-3 successful: delegate=rhel8-1, client=stonith_admin.1633, origin=rhel8-4, last-successful='2020-01-20 12:05:43 -06:00'

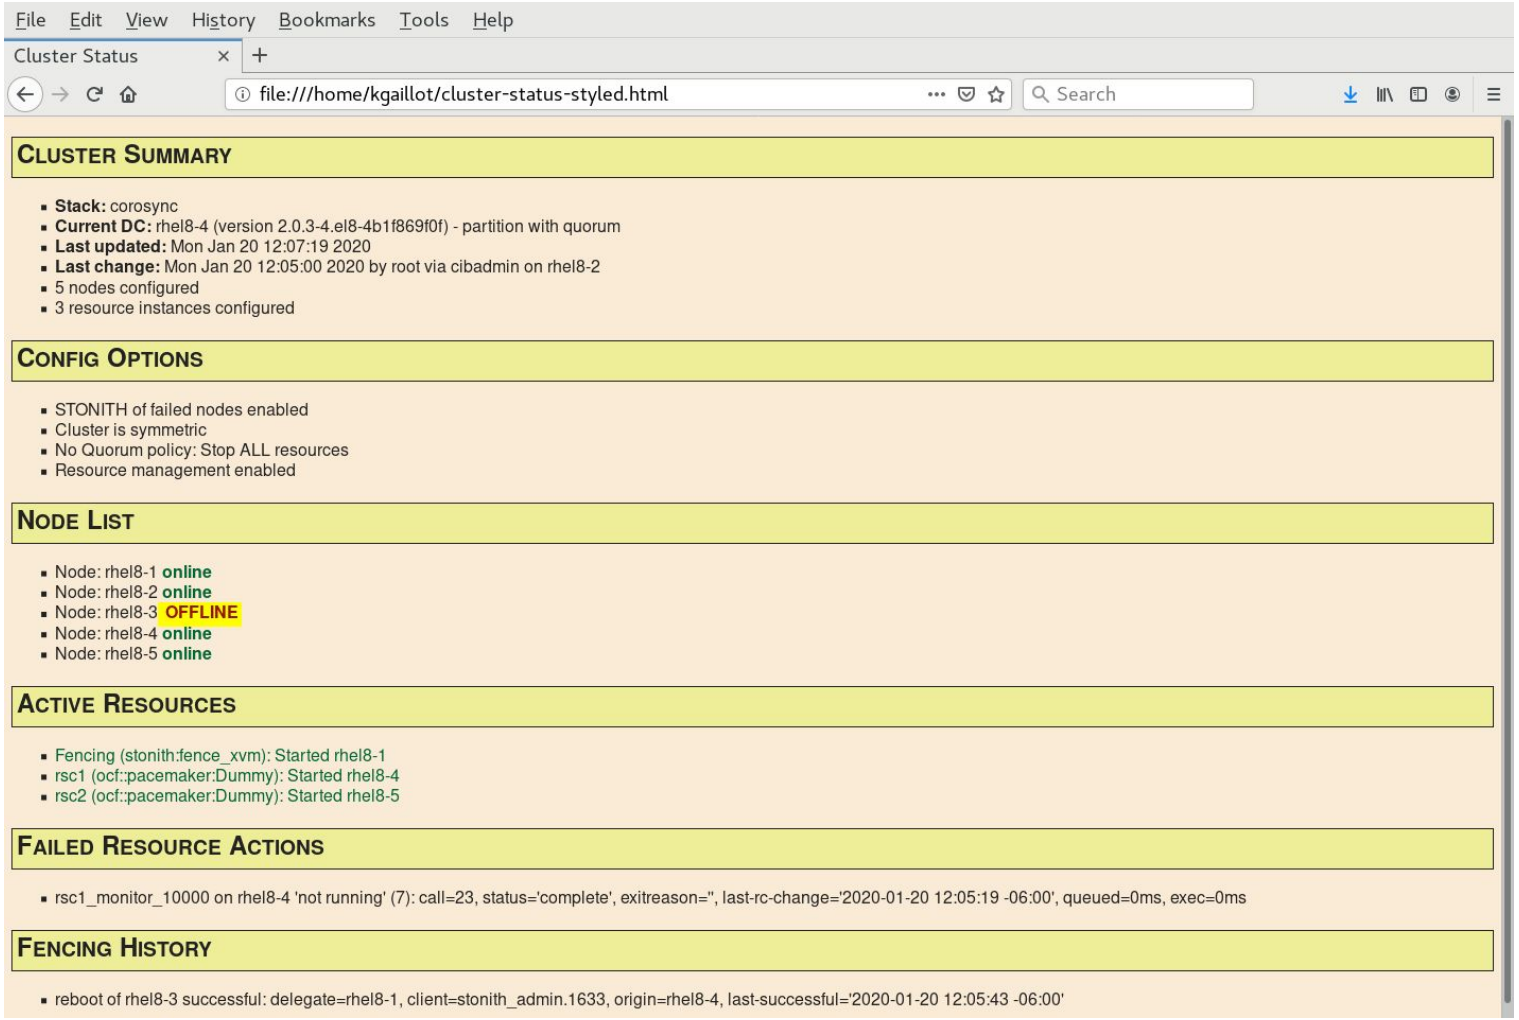
File Edit View History Bookmarks Tools Help
Cluster Status
file:///home/kgaillot/cluster-status-styled.html
Search
Cluster Summary
Stack: corosync
Current DC: rhel8-4 (version 2.0.3-4.el8-4b1f869f0f) - partition with quorum
Last updated: Mon Jan 20 12:07:19 2020
Last change: Mon Jan 20 12:05:00 2020 by root via cibadmin on rhel8-2
5 nodes configured
3 resource instances configured
Config Options
STONITH of failed nodes enabled
Cluster is symmetric
No Quorum policy: Stop ALL resources
Resource management enabled
Node List
Node: rhel8-1 online
Node: rhel8-2 online
Node: rhel8-3 OFFLINE
Node: rhel8-4 online
Node: rhel8-5 online
Active Resources
Fencing (stonith:fence_xvm): Started rhel8-1
rsc1 (ocf::pacemaker:Dummy): Started rhel8-4
rsc2 (ocf::pacemaker:Dummy): Started rhel8-5
Failed Resource Actions
rsc1_monitor_10000 on rhel8-4 'not running' (7): call=23, status='complete', exitreason='', last-rc-change='2020-01-20 12:05:19 -06:00', queued=0ms, exec=0ms
Fencing History
reboot of rhel8-3 successful: delegate=rhel8-1, client=stonith_admin.1633, origin=rhel8-4, last-successful='2020-01-20 12:05:43 -06:00'

Machine-friendly tool output

--output-as/--output-to

```
# stonith_admin --list-installed --output-as=xml
<pacemaker-result api-version="2.0" request="stonith_admin --list-installed --output-as=xml">
  <list name="Installed fence devices" count="3">
        <item name="device">fence_virt</item>
        <item name="device">fence_virtd</item>
        <item name="device">fence_xvm</item>
  </list>
  <status code="0" message="OK"/>
</pacemaker-result>
```

API XML schema

Output formats can take options such as HTML stylesheet

crm_mon -1 --output-as=html
--html-stylesheet="https://example.com/css/main.css"

Questions?

Expected this year

- Section selection in crm_mon
- Machine-friendly output: high-level C API
- Fencing reasons
- Shutdown locks
- Enhancements to access control lists (ACLs)
  - Colorized display
  - Groups
  - PAM integration

Questions?

Big ideas

- More intelligent disaster recovery support
  - Coordinated status across related clusters
  - Coordinated configuration changes
  - Easier failover testing (live or dry run)
- Event-driven resources
  - “Push” rather than poll for systemd resource status
  - Persistent daemonized resources
- More configurable failure response
  - Colocation constraint option for “noncritical resources”
  - failure-restart + failure-escalation

Questions?

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.


linkedin.com/company/red-hat


youtube.com/user/RedHatVideos


facebook.com/redhatinc


twitter.com/RedHat

Red Hat