

Infrastructure - Task #4459

Story # 3470 (Closed): CN cluster communication needs to be monitored

Monitor Hazelcast Logs with Splunk

2014-03-14 18:42 - Robert Waltz

Status:	Closed	Start date:	2014-03-14
Priority:	Normal	Due date:	
Assignee:	David Doyle	% Done:	100%
Category:	d1_monitor	Estimated time:	0.00 hour
Target version:	Maintenance Backlog	Story Points:	
Milestone:	CCI-1.2		
Product Version:	*		
Description			
Hazelcast spits a lot of output to a variety of log files. These logs files should be analysed to determine if errors or partitions have occurred in the cluster			

History

#1 - 2014-03-14 18:43 - Robert Waltz

- Milestone changed from None to CCI-1.2

#2 - 2014-03-14 18:59 - David Doyle

Will start by:

- build hazelcast indexes into Splunk test
- add /var/log/dataone/daemon/hazelcast-process.log and /var/metacat/logs/hazelcast-storage.log into Splunk test

From there I'll see how Splunk treats those logs and see if any changes need to be made at the input/index level.

Going to add bwilson as a watcher to see if he has any input.

#3 - 2014-03-17 19:51 - David Doyle

- Status changed from New to In Progress

Instead of having to try to figure out how to get every single log we want to monitor in the future routed through rsyslog into Splunk, I'm going back to looking into universal forwarders on the nodes that monitor files based on simple Splunk input config stanzas.

Starting out with using a forwarder on cn-sandbox-orc-1 to send /var/log/syslog and one or two Hazelcast logs to the Splunk test environment. Will see how that works out and build out further if this option works with a minimal amount of hiccups.

#4 - 2014-03-17 22:39 - David Doyle

Test setup running /var/log/dataone/daemon/hazelcast-process.log and /var/metacat/logs/hazelcast-storage.log into Splunk prod is working.

#5 - 2014-03-19 22:59 - David Doyle

Hazelcast logs are moving into Splunk from all sandbox CNs. Building out to other environments next.

#6 - 2014-03-20 03:31 - David Doyle

Logging now coming in from all non-prod CNs except for cn-stage-unm-2, for which I don't have sudo access.

#7 - 2014-03-21 04:39 - David Doyle

Robert reported some hazelcast messages that should indicate an error state. Leaving a brief IRC transcript below for Bruce and I to work back from.

Also, monitoring VM performance on ORC sandbox/dev/stage machines that have Splunk forwarders and hazelcast logging installed. (Mentioning this here since I'm rolling out hazelcast logging with the Splunk forwarders.) Seeing minor upticks in CPU and memory usage, all well within expected levels. Will in all likelihood build out to prod tomorrow/over the weekend.

```
[10:08pm] robert: i just saw the messages you should be looking for from cn-dev-ucsb-1
[10:08pm] robert: [ INFO] 2014-03-20 01:58:56,709 (AddOrRemoveConnection:process:58) [128.111.54.78]:5701 [DataONE] Removing Address
Address[64.106.40.9]:5701
[10:08pm] robert: [ INFO] 2014-03-20 01:58:56,710 (ConnectionManager:destroyConnection:338) [128.111.54.78]:5701 [DataONE] Connection
[Address[64.106.40.9]:5701] lost. Reason: Explicit close
[10:09pm] robert: [ INFO] 2014-03-20 01:58:56,743 (AddOrRemoveConnection:process:58) [128.111.54.78]:5701 [DataONE]
[10:09pm] robert: Members [2] {
[10:09pm] robert: Member [160.36.13.153]:5701
[10:09pm] robert: Member [128.111.54.78]:5701 this
[10:09pm] robert: }
[10:09pm] robert: [ INFO] 2014-03-20 01:58:56,792 (MemberRemover:process:40) [128.111.54.78]:5701 [DataONE] Removing Address
Address[64.106.40.9]:5701

[10:17pm] robert: Most important is the 'Members[d] {' string followed by the Member [ip-address] lines
[10:17pm] robert: also Connection Lost
[10:21pm] robert: and ...

[10:21pm] robert: [ INFO] 2014-03-20 01:59:21,028 (SocketAcceptor$1:run:111) [128.111.54.78]:5701 [DataONE] 5701 is accepting socket
connection from /64.106.40.9:58655
[10:21pm] robert: [ INFO] 2014-03-20 01:59:21,029 (SocketAcceptor$1:run:122) [128.111.54.78]:5701 [DataONE] 5701 accepted socket connection
from /64.106.40.9:58655

[10:55pm] robert: hmm, ya, so you'll see the Connect Lost, and I think that indicates an error condition until you see the accepted socket connection
```

#8 - 2014-03-23 00:46 - David Doyle

Hazelcast logging rolled out to prod CNs.

TO-DO: Find out if there are any hazelcast-related logs to be monitored on the D1 MN boxes.

#9 - 2014-03-26 18:18 - Chris Jones

David,

As Robert said, we can monitor cluster membership via log entries such as:

```
[ INFO] 2014-03-21 15:59:00,488 (?:?) [160.36.13.150]:5701 [DataONE]
```

```
Members [1] {  
Member [160.36.13.150]:5701 this  
}
```

In normal, non-split-brain operations, we would have 3 CNs in the cluster, and they would be listed by IP address in this log entry. When we upgrade CN software, we usually take them out of the cluster one at a time, so, a cluster membership of 2 can happen because of maintenance or because of a split-brain event. At the moment, we have purposefully isolated the three production CNs until we get the system metadata store consistent again, which is why you see only one member in the entry above.

We have three clusters we'll want to monitor: hzStorage, hzProcess, and hzSession.

The hzStorage cluster is run via Metacat, under Tomcat. It manages the shared system metadata map across CNs. As Robert mentioned, it logs to `/var/metacat/logs/hazelcast-storage.log`.

The hzProcess cluster is used to manage queues during `d1_synchronization`, `d1_replication`, etc., and it logs to `/var/log/dataone/daemon/hazelcast-process.log`.

The hzSession cluster is used in `d1_portal` to manage shared http session information across the CNs (basically a cookie-to-certificate mapping). From what I see in `/var/lib/tomcat6/webapps/portal/WEB-INF/log4j.properties`, it should be logging to `/var/log/tomcat6/portal.log`.

Monitoring all three clusters is a priority, but I'd say the hzStorage cluster is first priority because it involves a persistence layer (storing system metadata).

#10 - 2014-04-08 00:37 - David Doyle

I now have a (crude) scheduled search and alert in place to search hazelcast logs every minute for explicit close events and email an alert with relevant information (log info, source, etc) when it encounters more than 0 events over the last minute. This will in all likelihood need to be tweaked, especially once all the fixes on prod are in place and things are back to normal, but this is at least a starting point until things are more stable and we can hook alerts into abnormalities in the log data with more confidence. I'm meeting with Bruce Wednesday, and hopefully we can take a look at this in a little more detail then.

Just realized that I don't have hzSession cluster monitoring in place. Will put that in ASAP.

#11 - 2014-04-09 17:02 - David Doyle

Per cjones, added Chris, Matt, Dave, Robert, Skye, and Ben to alerted parties for this alert.

#12 - 2014-04-30 23:59 - David Doyle

Broke this alert up into four, one for each environment, with separate subject lines in alerts.

#13 - 2014-05-05 21:48 - Skye Roseboom

- *Target version changed from 2014.16-Block.2.4 to 2014.18-Block.3.1*

#14 - 2014-05-24 03:57 - David Doyle

- *translation missing: en.field_remaining_hours set to 0.0*

- *% Done changed from 0 to 100*

- *Status changed from In Progress to Closed*

Added hazelcast session logs to Splunk inputs on all CNs. Included archived logs. Complete pending alert changes and additions.

#15 - 2014-09-24 18:15 - Robert Waltz

- *Target version changed from 2014.18-Block.3.1 to Maintenance Backlog*