

OpenNebula - Feature #4809

Simplify HA management in OpenNebula

09/21/2016 08:49 AM - Ruben S. Montero

Status:	Closed	Start date:	09/21/2016
Priority:	Normal	Due date:	
Assignee:	Vlastimil Holer	% Done:	0%
Category:	Core & System	Estimated time:	0.00 hour
Target version:	Release 5.4	Pull request:	
Resolution:	fixed		
Description			
<p>This issue is to simplify the process of setting OpenNebula services in HA. Currently the system is designed to be compatible with classical Linux cluster suites. In practice, this impose to many requirements (e.g. need of a reliable fencing) and leads to deployments that are not easy to maintain.</p>			
Related issues:			
Related to Bug # 4796: When failing over HA controllers, hypervisor collectd ...		Closed	09/19/2016

Associated revisions

Revision e1cb2c92 - 04/10/2017 05:25 PM - Ruben S. Montero

F #4809: Add a base class for extended template attributes

Revision d03ed158 - 04/10/2017 11:32 PM - Ruben S. Montero

F #4809: Add zone server list to zone data

Revision ebc810f4 - 04/11/2017 07:34 PM - Ruben S. Montero

F #4809: Add Server to Zone API call

Revision c360b015 - 04/11/2017 07:41 PM - Ruben S. Montero

F #4809: AddServer zone API in federated setups

Revision cd580714 - 04/13/2017 04:31 PM - Ruben S. Montero

F #4809: Delete servers from zone

Revision a6d4ab3c - 04/18/2017 03:12 PM - Ruben S. Montero

F #4809: Update Sql database interface to include read/write/bootstrap operations

Revision e4848c55 - 04/19/2017 06:35 PM - Ruben S. Montero

F #4809: Overwrite SqlDB db_exec_* methods in DB implementations

Revision bae57600 - 04/19/2017 06:44 PM - Ruben S. Montero

F #4809: Some notes on the implementation of the log replication for zone servers

Revision c8981e82 - 04/20/2017 02:13 PM - Ruben S. Montero

F #4809: Some implementation files. Fix compilation issues

Revision 116425fc - 04/20/2017 03:11 PM - Ruben S. Montero

F #4809: Template for the LogDBManager

Revision dd0598aa - 04/21/2017 02:52 PM - Ruben S. Montero

F #4809: Work on log management and replication

Revision bca17f4e - 04/21/2017 05:16 PM - Ruben S. Montero

F #4809: Update SqlDB method name. DO NOT replicate monitoring data

Revision ed0d64a2 - 04/21/2017 08:32 PM - Ruben S. Montero

F #4809: Updated log structure

Revision b26e5a71 - 04/22/2017 11:43 PM - Ruben S. Montero

F #4809: Replication logic

Revision f2039e02 - 04/23/2017 11:09 AM - Ruben S. Montero

F #4809: Added LogDBManger to Zone server

Revision 034e2316 - 04/23/2017 11:15 AM - Ruben S. Montero

F #4809: Start replication on new exec_wr calls

Revision 0db5478b - 04/23/2017 11:20 AM - Ruben S. Montero

F #4809: Notify clients on majority replication

Revision 0c8299f1 - 04/25/2017 09:49 AM - Ruben S. Montero

feature #4809: Bootstrap LogDB tables in DB, initialize next_index

Revision 4c577126 - 04/25/2017 03:15 PM - Ruben S. Montero

F #4809: Improve replication logic, new RaftManager to control server states

Revision 7c479b8e - 04/26/2017 11:03 PM - Ruben S. Montero

F #4809: Restructure Raft implementation

Revision 900c37fd - 04/27/2017 12:16 AM - Ruben S. Montero

F #4809: Solves some bugs

Revision cc08fd98 - 04/27/2017 09:12 AM - Ruben S. Montero

F #4809: replicate logic for leader

Revision 88513088 - 04/27/2017 10:42 AM - Ruben S. Montero

F #4809: ReplicateLog client logic

Revision 613ec634 - 04/28/2017 05:35 PM - Ruben S. Montero

F #4809: Fix replication logic

Revision 6f976c1e - 04/28/2017 08:23 PM - Ruben S. Montero

F #4809: Make some Raft events synchronous with replication threads

Revision dd51b583 - 04/28/2017 08:46 PM - Ruben S. Montero

F #4809: DO NOT monitor marketplaces, nor datastores in followers

Revision 5441a2f8 - 04/29/2017 09:25 PM - Ruben S. Montero

F #4809: Log retention system

Revision 3be9f38d - 05/01/2017 08:30 PM - Ruben S. Montero

F #4809: ActionManager can set timers in ns. Adjust Raft manager to ms timers. Include Hearbeat control from leader and follower. XML-RPC calls are now async. RaftManager db calls is now decoupled with pool classes.

Revision 071d0d84 - 05/01/2017 11:38 PM - Ruben S. Montero

F #4809: start/stop replica threads when adding/deleting servers

Revision 3f9638be - 05/02/2017 03:57 PM - Ruben S. Montero

F #4809: Condifugre raft timers through oned.conf

Revision 2a695bc8 - 05/02/2017 10:31 PM - Ruben S. Montero

F #4809: Persist raft state. Implement Vote API call

Revision c6a7500d - 05/03/2017 06:04 PM - Ruben S. Montero

F #4809: Adds vote requests from candidate, fix some bugs

Revision 03c5698a - 05/04/2017 08:56 PM - Ruben S. Montero

F #4809: Fix several issues accessing the log records

Revision 9e1fd142 - 05/05/2017 01:46 PM - Ruben S. Montero

F #4809: Debug election process

Revision c14d648e - 05/05/2017 10:17 PM - Ruben S. Montero

F #4809: Method to query raft status one servers. Updated CLI to make use of it

Revision c8f77ab8 - 05/05/2017 10:20 PM - Ruben S. Montero

F #4809: Remove cache from core pools

Revision 9d397346 - 05/06/2017 09:36 PM - Ruben S. Montero

F #4809: Keep last_index counters internally. Commit log records on heartbeats

Revision 503b2835 - 05/07/2017 09:43 PM - Ruben S. Montero

F #4809: Update API internal name to match public xml-rpc names. Do not log heartbeat/replicate log entries

Revision c1317ed6 - 05/08/2017 08:12 AM - Ruben S. Montero

F #4809: Store the leader_id for the term

Revision 59cf651d - 05/08/2017 08:47 AM - Ruben S. Montero

F #4809: API methods leader_only attribute. List, Info and Raft methods are not leader only.

Revision 590b3548 - 05/08/2017 05:48 PM - Ruben S. Montero

F #4809: Forward request to leader

Revision ac76ab45 - 05/08/2017 09:57 PM - Ruben S. Montero

F #4809: Reload ACL rule cache when a server becomes zone leader

Revision 5f6a627b - 05/09/2017 09:53 AM - Ruben S. Montero

F #4809: Do not schedule VMs on non-leader servers

Revision 1bcfc2bd - 05/09/2017 10:29 AM - Ruben S. Montero

F #4809: Less verbose output

Revision 741a55d1 - 05/14/2017 09:48 PM - Ruben S. Montero

F #4809: Generic replication serves. Solve minor bugs in election process

Revision b38874a0 - 05/16/2017 10:02 AM - Ruben S. Montero

F #4809: Moved zone server list to ZonePool. Added FedLogDB class

Revision faa46d9c - 05/16/2017 10:21 AM - Ruben S. Montero

F #4809: fix compilation

Revision cfe3b415 - 05/16/2017 10:56 AM - Jaime Melis

F #4809: Handle properly if only one server in the zone

Revision 5e92ec0a - 05/16/2017 03:09 PM - Ruben S. Montero

F #4809: Use int for object IDs. Added base federation replica manager

Revision 85b6ed25 - 05/18/2017 02:36 PM - Jaime Melis

F #4809: Implement Hooks on RAFT events

Revision 94052286 - 05/18/2017 02:40 PM - Jaime Melis

F #4809: Move the FT hooks to its own folder

Revision 680d49a2 - 05/18/2017 02:41 PM - Jaime Melis

F #4809: Add RAFT hooks

Revision 2676d8b7 - 05/18/2017 07:25 PM - Ruben S. Montero

feature #4809: Start the federation replica manager with OpenNebula

Revision 4291e4a8 - 05/18/2017 10:52 PM - Ruben S. Montero

F #4809: Do not start replica threads for own zone

Revision 46d4e232 - 05/18/2017 10:57 PM - Ruben S. Montero

F #4809: Fix timers in federation replica manager

Revision c5a82aba - 05/18/2017 11:25 PM - Ruben S. Montero

F #4809: Update server list on election

Revision c5f54f81 - 05/19/2017 06:08 PM - Ruben S. Montero

F #4809: First version of replication thread for federated zones

Revision 4a780941 - 05/19/2017 06:37 PM - Ruben S. Montero

F #4809: Refresh zone list on server add/remove to keep an updated list of servers

Revision 31ba4d4d - 05/22/2017 04:04 PM - Ruben S. Montero

F #4809: Fix some bugs when replicating log entries to slave zones

Revision afca14d7 - 05/22/2017 08:51 PM - Ruben S. Montero

F #4809: Initialize last_index of federated log when starting replica threads

Revision a216bb9e - 05/23/2017 11:31 AM - Ruben S. Montero

F #4809: Raft servers / zone slave lists updated on server-add in the right place (leaders/masters)

Revision 98d26e9e - 05/24/2017 06:38 PM - Ruben S. Montero

F #4809: Do not accept empty records. Increase default xml-rpc timeouts

Revision c8d88de0 - 05/24/2017 09:57 PM - Ruben S. Montero

F #4809: Do not double-unlock raft mutex in timer. Fix update of heartbeat counter

Revision 8488320e - 05/25/2017 05:33 PM - Javi Fontan

F #4809: add migrator for logdb and fed_logdb tables

Revision 02b229d7 - 05/26/2017 03:49 PM - Ruben S. Montero

F #4809: Fix cache issues when updating PCI devices on VM

Revision 3d9686e3 - 05/27/2017 09:55 AM - Ruben S. Montero

F #4809: Fix history update for non-cache vm objects

Revision 9d763ef6 - 05/29/2017 04:24 PM - Jaime Melis

F #4809: Fully working vip hooks

Revision 54c1fbaa - 05/29/2017 04:24 PM - Jaime Melis

F #4809: Pass arguments to the raft hooks

Revision 75b8889e - 05/30/2017 10:07 AM - Ruben S. Montero

F #4809: Fix cache issue when the VM outdated set is modified

Revision 234734a3 - 05/30/2017 02:16 PM - Ruben S. Montero

F #4809: Better query to purge log records

Revision f12fdfae - 05/31/2017 10:52 PM - Ruben S. Montero

F #4809: Compress DB commands.

Revision 50880bb2 - 06/01/2017 09:13 AM - Ruben S. Montero

F #4809: Do not write log in solo mode, update instead of insert for timestamps

Revision 50d0a6b4 - 06/01/2017 10:21 AM - Jaime Melis

F #4809: sql is a reserved keyword in MySQL

Revision ab63767e - 06/01/2017 10:25 AM - Jaime Melis

F #4809: Tool to backup/restore the database (with federated option)

Revision d3c7a07a - 06/01/2017 02:14 PM - Ruben S. Montero

F #4809: Check log records are properly loaded from DB

Revision 35da0ee1 - 06/01/2017 02:21 PM - Ruben S. Montero

F #4809: Fix check

Revision 654b384b - 06/02/2017 10:01 AM - Ruben S. Montero

F #4809: Do not use log after unlocking objects

Revision e80386ef - 06/03/2017 04:22 PM - Ruben S. Montero

F #4809: Better timeouts for client xmlrpc-c

Revision b768f896 - 06/03/2017 04:22 PM - Ruben S. Montero

F #4809: Execute follower hook on startup

Revision 0a2f18a7 - 06/03/2017 05:24 PM - Ruben S. Montero

F #4809: Better heartbeat management. A separated thread is created for each follower

Revision 953a7f16 - 06/05/2017 04:14 PM - Ruben S. Montero

f #4809: Do not include log info on raft.status when operating in solo mode

Revision bd4040d3 - 06/05/2017 04:40 PM - Ruben S. Montero

F #4809: Do not bootstrap DB before upgrading

Revision 486a4856 - 06/06/2017 03:43 PM - Anton Todorov

F #4809: rename sql to sqlcmd in database_schema.rb

'sql' is reserved word in MariaDB

Revision 169d153e - 06/07/2017 01:23 PM - Jaime Melis

Merge pull request #324 from atodorov-storpool/patch-16

F #4809: rename sql to sqlcmd in database_schema.rb

Revision 978d1538 - 06/08/2017 06:36 AM - Ruben S. Montero

F #4809: Load Hook manager earlier to be able to execute raft hooks on start

Revision 585cc130 - 06/08/2017 08:19 AM - Jaime Melis

F #4809: Typo in the name of the hook

Revision 6d61d510 - 06/12/2017 09:54 AM - Ruben S. Montero

F #4809: Fix deadlock when stopping replica threads

Revision f4ec2c07 - 06/12/2017 02:24 PM - Ruben S. Montero

F #4809: Fix error when deleting server from zone in standalone mode

Revision 2f5a3185 - 06/13/2017 03:48 PM - Ruben S. Montero

F #4809: Start/Stop zone replication threads when adding/deleting zones

Revision 97474282 - 06/15/2017 10:17 AM - Ruben S. Montero

F #4809: sync zone update in server add when adding a HA follower in a federation (for 2 server corner case). Check for empty records in federated log replication

Revision dff203fe - 06/15/2017 10:33 AM - Ruben S. Montero

F #4809: Adding missing unlock()

Revision d00e3798 - 06/15/2017 05:34 PM - Ruben S. Montero

F #4809: Add missing migrators

Revision 0888ae48 - 06/17/2017 08:01 PM - Ruben S. Montero

F #4809: Fix log rewrite. Disable curl timeouts.

Revision c98457ae - 06/18/2017 11:10 AM - Ruben S. Montero

F #4809: Better cancellatin of replica threads

Revision 468a104b - 06/19/2017 12:36 PM - Ruben S. Montero

F #4809: Enable curl timeouts

Revision 83d158d6 - 06/19/2017 12:36 PM - Ruben S. Montero

F #4809: Fix memory leak after disabling the pool cache

Revision a2c5a4cb - 06/21/2017 01:22 AM - Ruben S. Montero

F #4809. Fix index in PoolSQL. Update interface in all classes

Revision e4280206 - 06/21/2017 07:43 PM - Ruben S. Montero

F #4809: Fix replication on failed zones, adjust timeout for replication

Revision a09c6d48 - 06/22/2017 04:31 PM - Jaime Melis

F #4809: Update VM in DB after generating context

This is an issue now because the cache has been removed.

Revision 47050433 - 06/22/2017 04:48 PM - Ruben S. Montero

F #4809: Do not start Raft timer in solo mode

Revision 15001cb1 - 06/22/2017 07:38 PM - Ruben S. Montero

F #4809: Fix ns retries

Revision ca2a1a42 - 06/22/2017 09:37 PM - Ruben S. Montero

F #4809: Fix non-persistent state of Hosts

Revision 3224a504 - 06/22/2017 09:47 PM - Jaime Melis

F #4809: Cleanup VIP if oned dies

Revision 8e19b566 - 06/22/2017 09:53 PM - Jaime Melis

F #4809: Cleanup script

Revision cc6cc460 - 06/22/2017 10:13 PM - Ruben S. Montero

F #4809: Add replicated log index information on server zones

Revision 509d98df - 06/22/2017 10:32 PM - Jaime Melis

F #4809: Send gratuitous ARP to update VIP

Revision d5d6cb96 - 06/26/2017 05:55 PM - Ruben S. Montero

F #4809: Get fed index from the DB (needed by followers in HA). Use Zone ENDPOINT to replicate log instead of server list. Fix bug when replicate fails in a zone.

Revision dbc47c98 - 06/26/2017 05:55 PM - Ruben S. Montero

F #4809: Remove unneeded update_zone calls when adding a server to a

zone

Revision 192ce9db - 06/27/2017 09:04 AM - Ruben S. Montero

F #4809: Fix replica log for federation

Revision 4f8195d2 - 06/27/2017 10:38 AM - Ruben S. Montero

F #4809: Better management of last index of federated log

Revision 0c9272b6 - 06/27/2017 11:14 PM - Ruben S. Montero

F #4809: Fix divergence problems when replicated the fed log in a HA zone

Revision 361409e7 - 06/28/2017 05:43 PM - Ruben S. Montero

F #4809: Compress federated log

Revision c6286c03 - 06/28/2017 05:50 PM - Ruben S. Montero

F #4809: Safer purge function for RAFT log

Revision 80d08166 - 06/29/2017 05:48 PM - Ruben S. Montero

F #4809: Re-design replicated log structure

Revision 899b460a - 06/30/2017 09:48 AM - Ruben S. Montero

F #4809: Enable federated of solo Zones

Revision fab2a07f - 06/30/2017 09:49 AM - Ruben S. Montero

F #4809: Log information to debug federated zones with HA clusters. THIS COMMIT IS MEANT TO BE REVERTED

Revision 50ab087c - 06/30/2017 12:57 PM - Ruben S. Montero

F #4809: Update onedb backup federated backup utility

Revision 6f4a45c8 - 06/30/2017 02:08 PM - Ruben S. Montero

Revert "F #4809: Log information to debug federated zones with HA clusters. THIS"

This reverts commit fab2a07f74f55528631fa5b6159e80c1fa884637.

Revision 6d24617e - 06/30/2017 03:02 PM - Ruben S. Montero

F #4809: Pre-allocate lastoid to prevent stale id's in the pool in case of leader failure

Revision fcf08d42 - 07/04/2017 10:42 AM - Ruben S. Montero

F #4809: Update migrator. There is no longer need to add servers to a zone to configure a federation if not using HA

Revision 87b5e5cb - 07/04/2017 04:34 PM - Ruben S. Montero

F #4809: Re-design replicated log structure

Revision 215bc0df - 07/04/2017 04:34 PM - Ruben S. Montero

F #4809: Enable federated of solo Zones

Revision 7b22d875 - 07/04/2017 04:34 PM - Ruben S. Montero

F #4809: Log information to debug federated zones with HA clusters. THIS COMMIT IS MEANT TO BE REVERTED

Revision 25b48b1e - 07/04/2017 04:34 PM - Ruben S. Montero

F #4809: Update onedb backup federated backup utility

Revision 3378c9a2 - 07/04/2017 04:34 PM - Ruben S. Montero

Revert "F #4809: Log information to debug federated zones with HA clusters. THIS"

This reverts commit fab2a07f74f55528631fa5b6159e80c1fa884637.

Revision cfd29830 - 07/04/2017 04:34 PM - Ruben S. Montero

F #4809: Pre-allocate lastoid to prevent stale id's in the pool in case of leader failure

Revision 031da2d3 - 07/04/2017 04:34 PM - Ruben S. Montero

F #4809: Update migrator. There is no longer need to add servers to a zone to configure a federation if not using HA

Revision b315b710 - 07/12/2017 02:04 PM - Ruben S. Montero

F #4809: Fix race condition updating next index in federated log.

History

#1 - 09/21/2016 08:50 AM - Ruben S. Montero

- Related to Bug #4796: When failing over HA controllers, hypervisor collectd probes do not switch to new controller added

#2 - 05/25/2017 02:17 PM - Ruben S. Montero

- Assignee set to Ruben S. Montero

#3 - 07/04/2017 03:32 PM - Ruben S. Montero

- Assignee changed from Ruben S. Montero to Vlastimil Holer

#4 - 07/13/2017 02:51 PM - Ruben S. Montero

- Status changed from New to Closed

- Resolution set to fixed