

D4Science Infrastructure - Task #2227

test replication tools for postgres

Feb 10, 2016 05:24 PM - Lucio Lelii

Status:	Closed	Start date:	Feb 10, 2016
Priority:	Immediate	Due date:	
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	High-Throughput-Computing	Estimated time:	0.00 hour
Target version:	UnSprintable		
Infrastructure:	Development		
Description install and test one of the tools for postgresSQL replication reported here https://wiki.postgresql.org/wiki/Replication,_Clustering,_and_Connection_Pooling Pgpool-II looks interesting			
Related issues: Blocked by D4Science Infrastructure - Task #4144: Two VMs are needed to test ... Closed May 31, 2016			

History

#1 - Feb 26, 2016 06:03 PM - Lucio Lelii

Do you have any update on this ?

#2 - Feb 26, 2016 06:05 PM - Pasquale Pagano

- Priority changed from Urgent to Immediate

@tommaso.piccioli@isti.cnr.it please react asap since we need to move forward on this activity.

#3 - Mar 01, 2016 02:39 PM - Tommaso Piccioli

- % Done changed from 0 to 20

New empty VM created:

postgres1-d-d4s.d4science.org
postgres2-d-d4s.d4science.org

ubuntu 14.04

#4 - Mar 01, 2016 02:45 PM - Andrea Dell'Amico

- Assignee changed from _InfraScience Systems Engineer to Andrea Dell'Amico

I'm taking it.

#5 - Mar 02, 2016 04:09 PM - Andrea Dell'Amico

- Status changed from New to In Progress

#6 - Mar 02, 2016 05:45 PM - Andrea Dell'Amico

- % Done changed from 20 to 40

pgpool II and postgres are installed on both machines. I'm studying the pgpool configuration to make it usable.

#7 - Mar 07, 2016 07:58 PM - Andrea Dell'Amico

- % Done changed from 40 to 50

pgpool is actually running on postgres1-d-d4s.d4science.org. It's configured to run on replication mode and the load balancer is active.

I don't have configured the cache memory yet, nor I did touch the connection pool parameters.

The failover capabilities are also not configured.

#9 - Mar 08, 2016 12:50 PM - Andrea Dell'Amico

- % Done changed from 50 to 60

The pgpool main configuration file is now a template, and we have a set of tasks to configure the authentication.

#10 - Mar 22, 2016 04:52 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

#11 - Apr 08, 2016 07:27 PM - Andrea Dell'Amico

- Status changed from Feedback to In Progress

The replication scripts need more testing. And we need some automated tools to remove the WAL archives regularly because they grow fast.

#12 - May 30, 2016 07:46 PM - Andrea Dell'Amico

- Assignee changed from Andrea Dell'Amico to _InfraScience Systems Engineer

I need two different VMs to test the replication without breaking the actual instances.

#13 - May 31, 2016 12:27 PM - Andrea Dell'Amico

- Blocked by Task #4144: Two VMs are needed to test the pgpool replication functionalities added

#14 - Jul 08, 2016 01:23 PM - Andrea Dell'Amico

@lucio.lelii@isti.cnr.it I need a short downtime to update the pgpool configuration on those hosts to be identical to the one seen on [#4144](#), otherwise the recovery is not going to work.

#15 - Jul 12, 2016 09:19 AM - Lucio Lelii

you can restart the node with pgpool, we are not using it yet.

#16 - Jul 12, 2016 03:28 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 60 to 100

The two servers are now up to date.

#17 - Jul 12, 2016 05:21 PM - Andrea Dell'Amico

- Status changed from Feedback to Closed

#18 - Jul 12, 2016 07:12 PM - Andrea Dell'Amico

Added ganglia monitoring, within a general "D4Science Development servers" cluster.

http://monitoring.research-infrastructures.eu/ganglia/?c=D4Science%20Development%20servers&m=load_one&r=hour&s=by%20name&hc=4&mc=2