D4Science Infrastructure - Task #7426

Add a general purpose postgresql server to collect the dev and preprod databases

Mar 06, 2017 06:14 PM - Andrea Dell'Amico

Status:	Closed	**Start date:**	Mar 07, 2017
Priority:	Normal	**Due date:**	
Assignee:	_InfraScience Systems Engineer	**% Done:**	100%
Category:	System Application	**Estimated time:**	0.00 hour
Target version:	PostgreSQL cluster		
Infrastructure:	Development, Pre-Production		

Description

With the postgis extensions installed.

Subtasks:

Task # 7434: Activate backups for the postgresql dev service	**Closed**

Related issues:

Related to D4Science Infrastructure - Incident #7647: Cannot edit Service end...	**Closed**	**Mar 22, 2017**	**Mar 23, 2017**
Related to D4Science Infrastructure - Task #7708: Please move Databases of de...	**Closed**	**Mar 23, 2017**	

History

#3 - Mar 07, 2017 10:01 AM - Andrea Dell'Amico

- Status changed from New to In Progress

hostname and IP address: postgresql-srv-dev.d4science.org 146.48.123.24

#4 - Mar 07, 2017 12:50 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 0 to 100

The VM is ready. We should migrate there all the postgres 9.x db instances used in dev and preproduction but the portal ones (and maybe them too):

geonetwork and geoservers when they are able to work against postgresql 9.6+
SimulFishGrowth_pre
SimulFishGrowth_prod
survey-db-1
survey-db-2
survey-db-3
surveyTestCreation
quota_dev_db
Others that I don't know about?

The database creation will be ansible provisioned until the service of #7429 will be ready.

#5 - Mar 22, 2017 03:59 PM - Massimiliano Assante

Let's do this, dev.d4science.org:5432/gcube-releases should be migrated too.

How should we proceed? We need one performing actions on the DB side and on Infrastructure Side to update the resources pointing to the db(s)

#6 - Mar 22, 2017 04:12 PM - Andrea Dell'Amico

Massimiliano Assante wrote:

> Let's do this, dev.d4science.org:5432/gcube-releases should be migrated too.

This is a production one despite it running on dev.d4science.org, correct?

> How should we proceed? We need one performing actions on the DB side and on Infrastructure Side to update the resources pointing to the db(s)

We first need to collect the database name, the username that needs to access the database (some of them are accessed as postgres but we need an unprivileged user; better if we have a different user for every database), and the passwords.
DB names and users are easy to collect, the passwords are not.

Then we need to perform a dump and a restore, and the applications data sources need to be modified.

#7 - Mar 22, 2017 07:10 PM - Andrea Dell'Amico

> gcube-releases-dev has been imported on postgresql-srv-dev.d4science.org

unfortunately I could not test this because of Incident #7647: Cannot edit Service endpoint on /gcube

#8 - Mar 22, 2017 07:24 PM - Massimiliano Assante

- Related to Incident #7647: Cannot edit Service endpoint on /gcube added

#9 - Mar 23, 2017 04:59 PM - Massimiliano Assante

- Related to Task #7708: Please move Databases of dev2.d4science.org and preprod1.d4science.org to postgresql-srv-dev.d4science.org added

#10 - Apr 11, 2017 07:27 PM - Andrea Dell'Amico

- Status changed from Feedback to Closed