

## Operations - Support #66

### Get backups running on gdp-0[1234]

10/07/2016 09:01 AM - Eric Allman

<b>Status:</b>	Closed	
<b>Priority:</b>	Normal	
<b>Assignee:</b>	Nitesh Mor	
<b>Category:</b>		
<b>Target version:</b>		
<b>Description</b>		
For Terraswarm review. Back up 01→02, 02→03, 03→04, 04→01. Related: <a href="#">#65</a> .		
<b>Related issues:</b>		
Related to Operations - Support #65: Upgrade gdp-01 to Ubuntu 16.04		<b>Closed</b>

#### History

##### #1 - 10/12/2016 11:59 PM - Nitesh Mor

- Status changed from New to In Progress

- Split the 4TB disk on gdp-0[1234] into two logical volumes each, one for own data and one for the backup of a remote log-server.
- Created a different user gdp-backup, with SSH keys setup for password-less login.
- Changed the umask setting for the log-server to be a little more relaxed, such that gdp-backup can read the data

Data transfer in progress. Looks like a few hours for gdp-01's 100+GB data files.

##### #2 - 10/13/2016 12:00 AM - Nitesh Mor

- Related to Support #65: Upgrade gdp-01 to Ubuntu 16.04 added

##### #3 - 10/14/2016 01:28 AM - Nitesh Mor

It turns out that naively doing an rsync is not a very good strategy. Default behavior of rsync is

- make a decision whether a file needs to be synced or not based on modification time and filesize.
- if so, then do a full sync based on a rolling checksum.

With such an approach, if all the logs get a single new record appended, this kind of rsync will require reading the entire contents of the disk. However, with the knowledge that the .gdplog files are only appended to, we can use the --append flag safely. For .gdpndx and .gdptidx files, this assumption does not hold true necessarily. But we can always regenerate them from the data file, so we don't even need to do any backups for them.

##### #4 - 02/08/2017 11:55 AM - Eric Allman

- Status changed from In Progress to Resolved

I believe backups are working now. Nitesh, can you confirm?

##### #5 - 02/13/2017 02:07 PM - Nitesh Mor

*- Status changed from Resolved to Closed*

Confirmed that backups are working. Closing the issue.