

VegBank - Bug #2395

Fix Backup Strategy -- delete extra backups from time to time

04/05/2006 10:37 AM - Michael Lee

Status:	Resolved	Start date:	04/05/2006
Priority:	Immediate	Due date:	
Assignee:	Chad Berkley	% Done:	0%
Category:	misc	Estimated time:	0.00 hour
Target version:	UnMilestoned	Spent time:	0.00 hour
Bugzilla-Id:	2395		
Description			
<p>Our current backup strategy is to backup the database every night (which is good). However, we don't delete any old backups, so we will inevitably run out of space. This needs to be fixed. We certainly don't need backups for EACH day the database has existed, but having some older copies might not be horrible either. I propose:</p> <p>Keep all the backups for the last month. Keep one weekly backup (or every 10 days, whatever is easier) from one month ago to 6 months ago. Keep one monthly backup for anything older than 6 months.</p> <p>In addition, we need to make sure that this backup is not just to VegBank- a copy needs to be distributed, either by email or network to an alternate machine. I don't know the current status of that "other-machine/off-site" backup, but it would be nice to know. PMark?</p>			
Related issues:			
Blocked by VegBank - Bug #2481: Fix Backup Strategy: where backups go		Resolved	06/28/2006

History

#1 - 04/05/2006 11:25 AM - Michael Lee

I just deleted backups from September 2005 and previous, keeping one per month in a new dir: /usr/vegbank/backup/oldarchive

This created 9GB free space.

#2 - 04/06/2006 12:41 PM - P. Anderson

Just off the top of my head, the simplest (but not best) approach would be to remove files after X days. This crontab line sets X=180.

```
10 2 * * * find /usr/vegbank/backup -name "*.bz2" -mtime +180 -maxdepth 1 -exec rm \{} \;
```

#3 - 04/06/2006 01:52 PM - Michael Lee

(In reply to comment [#2](#))

Just off the top of my head, the simplest (but not best) approach would be to remove files after X days. This crontab line sets X=180.

```
10 2 * * * find /usr/vegbank/backup -name "*.bz2" -mtime +180 -maxdepth 1 -exec rm \{} \;
```

This doesn't solve the issue of keeping old backups (one per month). Also, if our system time gets off for some reason, this script could kill all our backups. Not sure what to do about that latter issue, as any approach likely would cause problems if the system time got off. Might be nice to somehow check that not all the backups are being deleted before deleting some? Just throwing ideas out...

#4 - 04/06/2006 06:03 PM - Michael Lee

Suggestion from Matt on conference call 06-APR-2006:

We need to copy the backups to a place on our hard drive, as we now do, which will be backups if we need them before tape backups are performed. Also, we need to copy the files to a machine that is tape-backedup. Older backups can then be deleted, because they are on tape.

IMPORTANT We also have items NOT in CVS that are large binary files that should be backed up when they are changed. These items are located in:

/usr/www/vegbank/htdocs/noncvs
on vegbank

#5 - 05/08/2006 10:23 AM - Chad Berkley

Colby needs to be the one to do this.

#6 - 06/28/2006 10:13 AM - Michael Lee

This bug has been split. This portion of the bug now deals only with removing extra old backups in the /usr/vegbank/backup dir of VegBank. They never get deleted, thus eventually we'll run out of space, see other comments for details.

#7 - 09/04/2006 11:09 AM - Chad Berkley

the backup script now keeps 5 backups (from the last 5 days) around but deletes older ones so that the disk does not get filled up. We should also have this directory backed up to tape at least once every 5 days to make sure we have an archive.

#8 - 03/27/2013 02:20 PM - Redmine Admin

Original Bugzilla ID was 2395