W. W. Hansen Experimental Physics Laboratory
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305 - 4085
Gravity Probe B Relativity Mission

DATABASE BACKUP PROCEDURES FOR SCIENCE NETWORK DATABASES

P0996 Revision A

August 5, 2005

Approvals

NAME	**SIGNATURE**	**DATE**
Carin Kahn Database Administrator		
Paul Shestople *Data Processing IPT*		
Ron Sharbaugh *S/W Manager*		
Marcie Smith *MOC Project Manager*		
Dorrene Ross *Software Quality Assurance*		

_____

Tom Langenstein ITAR Assessment Performed, ITAR Control Req'd? □ Yes □ No

Table of Contents

1 Revision History

REV	DATE	AUTHOR	COMMENTS
-	7/29/2004	CWK	Initial version
A	8/5/2005	CWK	Reflects hardware changes and changes necessary later in mission due to larger databases.

2 Scope of Document

This document describes the procedures for backing up the Sybase Science Network Databases. The backup system supports the following goals

- Provide for long-term archival copies of databases
- Support recovery of one or several databases in case of media failures or any other event that render them unusable
- Support database recovery in a way such that both downtime and data loss is minimized
- Maintain a warm standby database server to allow for a quick switchover when the primary database system fails

The document describes the various activities that support the abovementioned goals. Backing up the science network databases is an ongoing activity. Most of the steps are automatically scheduled using cron. Step-by-step instructions are provided for activities that are manually done.

Knowledge of Unix and Sybase is required to perform the activities described in this document.

3 Reference Documents

Document	Document No.	ALIAS
Data Management Plan	S0331	
Science Data Network 24 Hour Recovery Plan During IOC	S0976	
Database Backup and Recovery Plan for Science Network Databases	S0979	
Database Backup Procedures Using the Standby Database Server	P1070	
Database Recovery Procedures for Science Network Databases	P1071	
Assessment of Science Network Database Errors and Failures	P1081	
Telemetry Data Processing (TDP) in the Non-Real-Time System	P0826	

4 Operational Personnel

The procedures described in this document may only be conducted by the following persons:

- Carin Kahn or the designated Database Administrator
- Dorrene Ross Qualified QA Rep

5 Overview Backup Plan

In the backup plan a distinction is made between user and system databases. User databases contain the data that are accessed and manipulated by the end user, system databases control the operation of the database management system itself. GPB_L0, GPB_L1, GPB_L1A, L2, orbit_determ and the MSS databases are all user databases. Master, model and sybsystemprocs, among others, are examples of system databases.

5.1 Backup Plan for User Databases

- Fix allocation errors for text and image data for affected tables immediately before starting full backups. See Contingency 1 below.
- Full backups of the user databases are made to disks weekly. These are RAID-5 disks attached to sci-base, directoris /sybackup1/carin/sybbackups & /sybackup2/carin/sybbackups.
- Transaction logs for user databases are dumped to a disk attached to sci-crunch every 30 minutes, using directory /sybackup1/carin/syblogs. In addition, a dump of the transaction log may be triggered by an sp_thresholdaction procedure when the transaction log is starting to fill up.
- The day after the full backups are made the user databases are recovered to the warm standby Sybase server.
- Once the full backups have been loaded on the standby server, the databases are checked for logical and physical consistency using the dbcc checkstorage command. If time allows more time-consuming checks are run as well.
- When the newly loaded databases on the standby server have passed the consistency check they are backed up to two sets of tapes for archival. The two sets of tapes are stored in different physical locations. Tape backups are kept for three months, both copies for one set of backups per month are kept for the life of the mission and beyond.
- Verify one set of backup tapes by restoring at least one database from tape to the standby server.
- The preloading of transaction logs to the warm standby server resumes after the archive tapes have been created and verified. Tranlogs are recovered about four times a day, lagging 24 hours behind the production databases.
- Once a day a faster but less thorough consistency checker is run on the standby databases.
- The new transaction log files are copied to a secondary directory, /home/sybase/syblogs, several times a day.
- Weeknights the transaction log files are copied to the incremental OS-backup tapes managed by the SA. These tapes are recycled weekly. Additional incremental OS-backup tapes are generated weekly and recycled monthly. It should thus be possible to recover any transaction log file from an OS-tape made within the last month. The OS-backups are not action item for the dba.
- Tranlog files are removed from the primary backup disk, /sybackup/carin/syblogs after four weeks and from the secondary repository, /home/sybase/syblogs, after one week.
- Once a day scripts for comparing disk devices and database allocations between the production and standby servers are run. This ensures that the two are set up identically, minimizing surprises in case an emergency switchover to standby server has to be initiated.
- After IOC, towards the end of the mission, a job that prunes the semi-temporary table Snaptemp was instituted, thereby avoiding additional manual table maintenance.

Towards the end of the mission, as databases have grown larger and larger, and the uptime requirements have been relaxed somewhat, the backup plan above has been modified slightly. To ensure a turnaround of the backup/preload/check/tape backup-cycle within a week, preloads are no longer delayed for a full 24

hours and dbcc's are performed concurrent with and sometimes even after the tape backups. The most crucial check, dbcc checkstorage, is run once a week for all databases but the very time-consuming and less crucial dbcc checkdb check might only be run every few weeks. Tranlog dumps are preloaded to the standby server as time allows.

Contingency 1: Dealing with tables with text or image data:

A bug in Sybase 12.5.0.3, EBF 11331, causes allocated text and image pages to not be deallocated when inserted or updated rows are rolled back. These allocation errors should be fixed before full backups are started.

Contingency 2: When full backups no longer fit on the backup disk:

Due to time constraints, a large enough disk to hold the full user backups could not be attached to sci-crunch before start of IOC. Two months after launch the full backups thus no longer fit on the backup disk. The workaround was to backup the production databases directly to the two sets of tapes that are kept for archival. The standby databases are loaded from these tapes. No additional tape backups are made and the load itself constitutes a read verification of the tapes. Other workarounds, although never used, is to backup only a few databases at a time, decreasing the need for simultaneous disk space, or, if a single database no longer fits on one backup disk, to stripe it to two disks.

5.2 Backup Plan for System Databases

- Several times a week complete consistency checking is done on the master and other system databases for the production and standby systems.
- Full backups of the master database are made to the /sybackup1/carin/sybbackups directory daily. Nine generations of these backups are kept, representing slightly more than one week.
- In addition to the master backups, the most critical tables in the master databases on both installations are bulk-copied out, printouts are written to file of important system tables, the configuration file and configuration block are backed up and scripts are generated for recreating disk devices and user and system databases. This is done daily and the files are stored in the /sybackup1/carin/sybbackups directory.
- The last three generations of master database backups and data files are copied to the /home/sybase/sybsystem directly nightly, well before the OS-backup of that file system is scheduled to start.
- Weeknights the master dump and data files are copied to the incremental OS-backup tapes managed by the SA. These tapes are recycled and overwritten every week. In addition to the full OS-backup, which is run monthly, every weekend an incremental OS-backup is generated. The weekly incremental backup is recycled monthly.
- System databases other than the master database are backed up to the /home/sybase/sybsystem directory, which is backed up to tape as part of the OS-backups as run by the SA. This is done on an as-needed basis, whenever changes have been made. It is the responsibility of the dba to make sure backups are duly made.

5.3 Backup Architecture

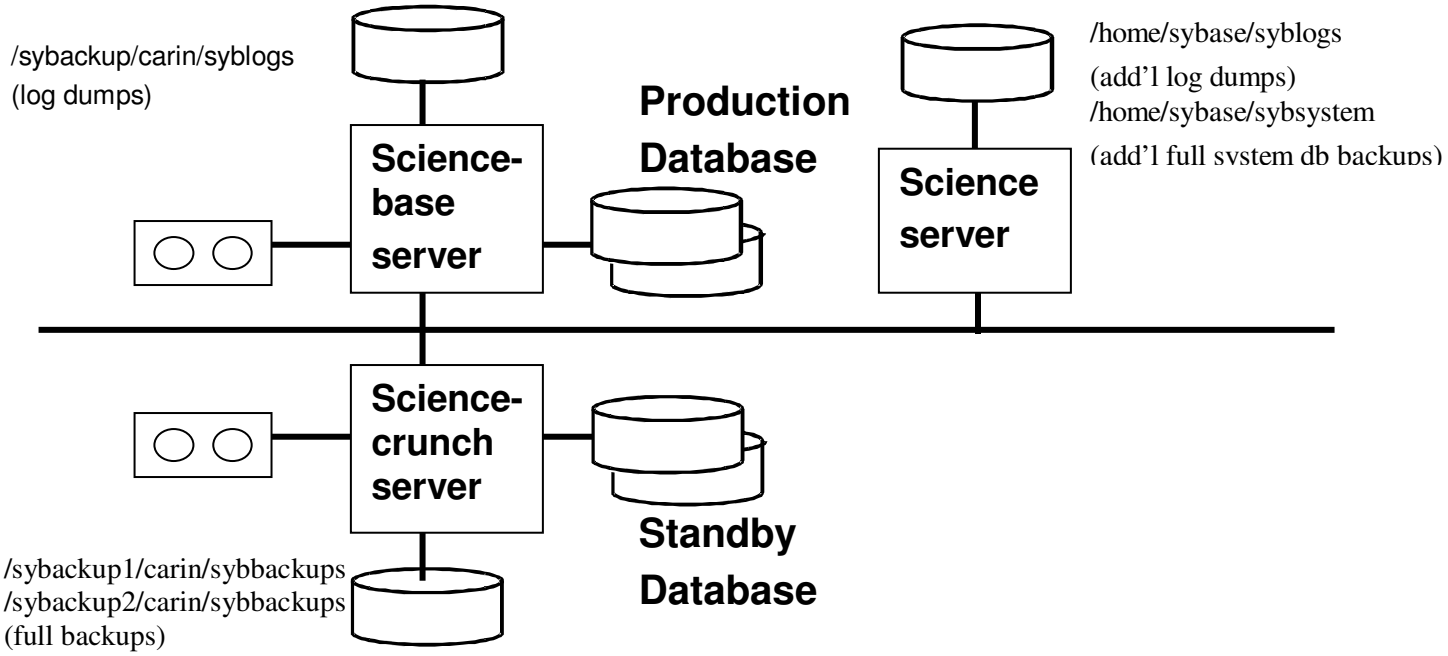

Figure 5.3

The tape drives on the science-base and science-crunch servers have been configured the exact same way. A tape written on one tape drive can thus be read on the other and vice versa. The configuration must not be altered, a change would render existing tapes unreadable. Although the tape drives are interchangeable, a remote tape dump will take five to ten times longer than a local dump.

To verify the tape drive setup issue:

- mt –f /dev/rmt/0cn status

Note that two slightly different outputs are valid:
MIRROR with two tapes loaded or AUTOLOAD with one tape loaded:

Sony AIT-x 8mm tape drive:
sense key(0x6)= Unit Attention residual= 0 retries= 0 file no= 0 block no= 0

MIRROR with one tape loaded or AUTOLOAD with two tapes loaded:

Sony AIT-x 8mm tape drive:
sense key(0x0)= No additional sensor residual= 0 retries= 0 file no= 0 block no= 0

- cat /kernel/drv/st.conf

locate a tape-config-list label with the following setup (others may co-exist):

```
tape-config-list=
"SONY     SD", "Sony AIT-x 8mm", "SONY_AIT"
SONY_AIT = 1, 0x36, 0, 0xd679, 4, 0x00, 0x00, 0x00, 0x00, 0
```

5.4 Overview of Backup Schedule

The following graph depicts the more important steps in the backup plan for the science network databases. A more complete picture, including detailed descriptions and in some cases step-by-step instructions of the different activities can be found in the sections below.

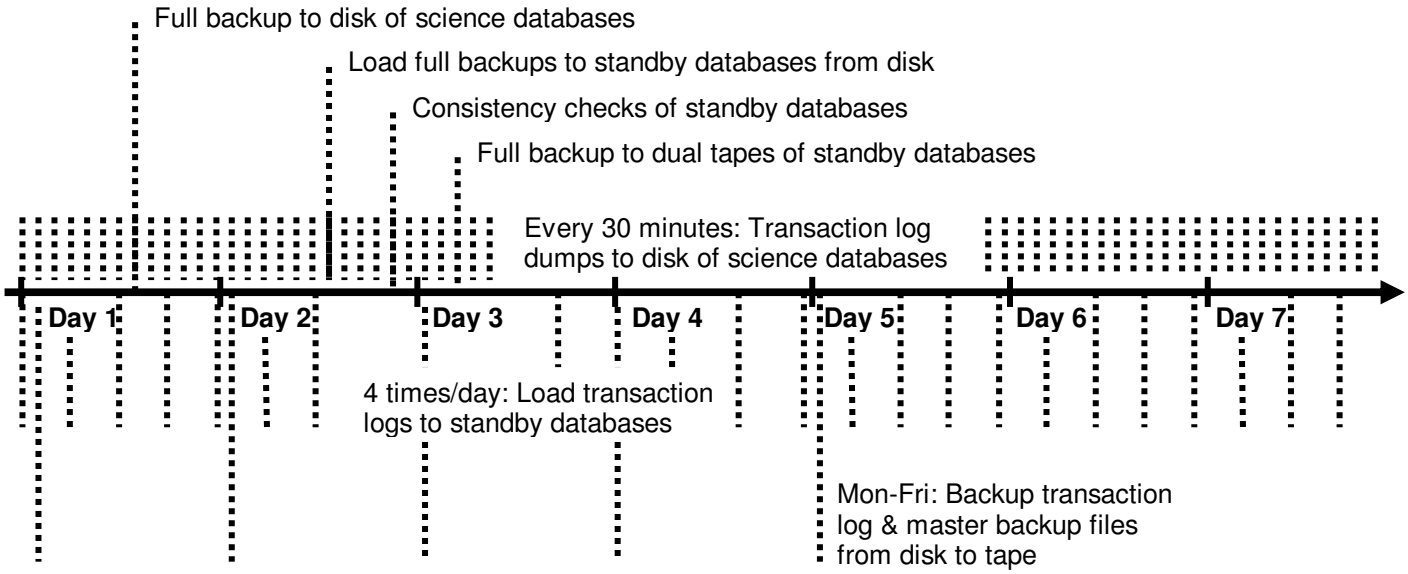

6 Backing up the Science Network Databases

6.1 Backing up User Databases

Most jobs, both those automatically scheduled as well as those initiated by the dba, mail the log to the dba and send a page to the dba-pager whenever an error occurs. This system is not fail-safe however. If a job hangs and never completes or any one of the job scheduler, cron, the mail system or mail forwarding, is malfunctioning, the dba will simply not know that an error has occurred or that a crucial job did not even run.

It is the responsibility of the dba to routinely verify that jobs **are** run and that errors do not go undetected. This is easily done by reviewing the output in /home/sybase/sybdba/log for the pertinent jobs.

6.1.1 Fix allocation errors for text and image data

A bug in the Sybase software, version 12.5.0.3 EBF 11331, causes problems with tables containing image or text data. When a row is rolled back, the rolling back of the row itself succeeds but the allocated text or image pages are not deallocated. Not only is this a waste of space but it means that an inordinate number of errors are generated when running dbcc's, possibly masking more serious and potentially fatal problems.

The indication that a table with text/image data is experiencing this problem is that errors 2540, 2546, 7940 and 7949 are generated when dbcc textalloc or checkstorage is run. The remedy is to run dbcc textalloc on the table in question with the fix-option.

There are four different tables in the science network databases that have text or image data, GPB_L0..Snapshots, GPB_L1..GPS_Packets, GPB_L1..GPS_Telemetry and L2..run_info.

To date this has only been a problem with the Snapshots table in the GPB_L0 database. When rows are added to this table the application code does not eliminate duplicates but relies on the Sybase index constraints, "ignore duplicate keys", to reject them. In effect, an attempt is made to insert all rows, duplicates are rolled back and, alas, text/image data pages erroneously are not deallocated.

The "ignore duplicate keys" index option has never been set on the indices for the GPS_Packets and GPS_Telemetry tables. The good news is that the erroneous space allocations thus never occur. The bad news is that a small portion of the telemetry data never makes it into these tables. Instead of rejecting individual rows base on index, the entire batch is rejeceted if at least one row in the batch contains a duplicate key. Since the tables are not being used, a decision was made not to correct the index to add the "ignore duplicate keys" option.

The run_info table has never been affected since duplicate data is programmatically removed by the science application before the new data is added.

<u>Jobs automatically scheduled with cron:</u>

3-4 times a week on science, once before the full_backup is initiated:

/home/sybase/sybdba/script/dbcc_fix_text.csh science_prod GPB_L0 Snapshots

runs dbcc textalloc

Dealing with errors:

When errors are encountered in the script the log file is mailed to the dba of record and a page sent to the dba-pgr. The dba corrects the error and manually restarts the job as appropriate.

6.1.2 Full backups of user databases to disk

Full backups of science network databases are made to RAID-5 disks that are attached to sci-crunch, using the /sybackup1/carin/sybbackus and /sybackup2/carin/sybbackups directories. The disks are not backed up to tape.

Full backups of the tqsm database is made to /home/sybase/sybsystem. This disk partition is backed up to tape weeknights by the SA, the tapes are recycled weekly. Monthly the current tqsm backup is copied to the "permanent" full backup of disk partitions as handled by the SA. Two generations of the backup files are kept in the /home/sybase/sybsystem directory.

Jobs automatically scheduled with cron:

Weekly on science, currently Friday evening Pacific time:

/home/sybase/sybdba/script/full_backup.csh science_prod GPB_L1.L2.orbit_determ.gpb_3_4_3.GPB_DP

/home/sybase/sybdba/script/full_backup.csh science_prod GPB_L0.GPB_L1A

/home/sybase/sybdba/script/full_backup.csh tqsm_server GPB_MOC # tqsm data

Jobs/tasks to be initiated by the dba:

None. But full backups may be generated as needed by issuing

/home/sybase/sybdba/script/full_backup.csh <server> <db1.db2.db3..> \
<db backups to keep> <tranlog cycles to keep>

Dealing with errors:

When errors are encountered in the script, the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/<server>_back.log. The dba corrects the error and manually restarts the job as appropriate.

Contingency: When full backups from science_prod no longer fit on disk:

If the backups no longer fit on the backup disk they are backed up directly to tape from the production server. See section 6.1.6 below, "Backing up user databases to tape" for detailed instructions on how to backup the databases to tape. The exact command to use is listed in step 6.1.6.6 below. Transaction log dumps cannot take place while a full backup is occurring. It's thus extremely important to switch tapes in a timely manner on an actively updated database to avoid filling up the tranlog.

An alternative approach when all the backups no longer fit on the backup disk may be used:

Day 1: Backup all databases except GPB_L1A to disk

Day 2: Preload the dumps from day 1 to the standby server and perform tape backups

Day 3: Delete the disk backups and backup GPB_L1A to disk

Day 4: Preload GPB_L1A to standby server and perform tape backups of GPB_L1A alone.

If GPB_L1A no longer fits on one backup disk, the dump file must be striped to two disks. The backup command will then look something like

dump database GPB_L1A to /sybackup1/carin/sybbackups/science_prod/gen_0/GPB_L1A.full.<date>

stripe on /sybackup2/carin/sybbackups/science_prod/gen_0/GPB_L1A.full.<date>

6.1.3 Backup transaction logs for user databases to disk

Backups of the transaction logs are written to the /sybackup/carin/syblogs directory on the backup disk attached to sci-base. The log dumps stay in this directory for four weeks. Tranlog dumps for the last week can also be found in the /home/sybase/syblogs directory.

Jobs automatically scheduled with cron:

Every 30 minutes on science:

/home/sybase/sybdba/script/tranlog_backup.csh science_prod # active user tranlogs on science_prod

/home/sybase/sybdba/script/ tranlog_backup.csh tqsm_server GPB_MOC # tqsm data

Jobs/tasks to be initiated by the dba:

None. But additional tranlog dumps may be run as needed by issuing

/home/sybase/sybdba/script/tranlog_backup.csh <server> <db1.db2.db3…>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/<server>_back.log. The dba corrects the error and manually restarts the job as appropriate.

6.1.4 Preload full backups of user databases to standby server

In order to minimize downtime if a switchover to the standby server is necessary, the full backups are preloaded to the standby server, science_pc. The preloading also offloads maintenance activities from the production server. The standby server typically runs 24 hours behind the production server. This is to allow for the detection of data errors before they are reproduced on science_pc.

Only the user databases that are actively being updated are preloaded to the standby server. Older versions of the MSS databases are set to read-only and are not preloaded.

Jobs automatically scheduled with cron:

Weekly, on sci-crunch, typically at least 24 hours after the full backup started:

/home/sybase/sybdba/script/preload_standby.csh L2.GPB_L1.orbit_determ.gpb_3_4_3.GPB_DP

/home/sybase/sybdba/script/preload_standby.csh GPB_L1A.GPB_L0

preloads user dbs to science_pc

Jobs/tasks to be initiated by the dba:

None. But additional preloads may be run as needed by issuing

/home/sybase/sybdba/script/preload_standby.csh <db1.db2.db3…> <gen_to_load> \

<dbcc-switch> <db to load TO>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/science_pc_back.log. The dba corrects the error and manually restarts the job as appropriate.

Occasionally it will not be possible to put a preloaded database online for standby access. This is caused by uncommitted transactions being present in the log. When this happens, simply preload the next transaction log in the sequence and try again. Do not online the database without the standby access option, you will not be able to preload subsequent tranlogs once this has been done.

Contingency: When full backups from science_prod are made directly to tape:

If the standby server is preloaded from backup on tape, write-protect the tapes, set the ASP-mode to Autoload Idle and run the tape_load.csh script. Detailed Instructions for how to recover databases from tape can be found in Document P1071, "Database Recovery Procedures for Science Network Databases". The load should be run as a local tape load on sci-crunch. Run the command:

/home/sybase/sybdba/script/tape_load.csh science_pc <YYMMDD> science_pc_back <list of dbs> same unload

6.1.5 Perform consistency checking of user databases

The various consistency checker programs supplied by Sybase, dbcc's, verify page and row level linkage and pointers as well as the page allocations. The most complete checking is performed by running a combination of dbcc checkstorage, checkcatalog and checktable commands.

Dbcc checktable is very time-consuming, to check 100 GB of data takes 8-9 hours. This check is therefore only run once a week during IOC and even less often later in the mission. Dbcc checkcatalog takes less than a minute and is run daily, together with dbcc checkstorage during IOC, once a week later in the mission. About 100 gb of allocated database space can be checked per hour using dbcc checkstorage if set up for parallel processing.

Even though text/image allocations are routinely fixed in the production system, those changes are not recorded in the tranlog. This means that the fixes do not replicate to the standby server when the tranlogs are preloaded. Dbcc checkstorage or textalloc cannot thus be run daily on GPB_L0, numerous errors would be generated by erroneously allocated Snapshots-pages. In its place dbcc checkalloc & checkdb are run.

Jobs automatically scheduled with cron:

Run after preloads to standby server. Daily during IOC, weekly late in mission:

/home/sybase/sybdba/script/dbcc.csh science_pc <MSS>.orbit_determ.L2.GPB_L1.GPB_L1A.GPB_L0 \
checkcatalog.checkstorage

Run after preloads to standby server. Weekly during IOC, less often later in mission:

/home/sybase/sybdba/script/dbcc.csh science_pc <MSS>.orbit_determ.L2.GPB_L1.GPB_L1A.GPB_L0
checkcatalog.checkdb

Jobs/tasks to be initiated by the dba:

None. But additional preloads may be run as needed by issuing

/home/sybase/sybdba/script/dbcc.csh <server> <db1.db2.db3> <dbcc_option1.dbcc_option2>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record.

Errors detected by dbcc checkcatalog, checkdb and checkalloc are listed with the appropriate error number. Look up the error in the Sybase Troubleshooting guide and perform the actions necessary to correct it. If appropriate, test the correction on science_pc first, then correct on the source server, science_prod, generate a new backup and overwrite the corrupt database on sci-crunch.

Dbcc checkstorage only lists the number of errors detected. After dbcc checkstorage completes, the script automatically runs dbcc checkverify in the targeted database, on the standby server to see if the errors detected are transient or "real". If the error is a true database error it must be corrected. To get the error numbers, log in to the affected database on the standby server and run the stored procedure sp_dbcc_faultreport "short". Look up the listed errors in the Sybase Troubleshooting guide and perform the actions necessary to correct it. If appropriate, test the correction on science_pc first, then correct on the source server, science_prod, generate a new backup and overwrite the corrupt database on sci-crunch.

6.1.6 Backing up user databases to tape

Tape backups should always be run locally, using the backup server and tape drive that is associated with the targeted dataserver. Because of tape handling sequences in the backup script the job itself must be started on the host where the tape drive is attached.

As the GPB_L0 database and the Snapshots table grow during the mission, it is becoming less likely that existing allocation errors can be fixed and the database fully backed up before new errors are introduced. (It would be necessary to stop data processing to avoid this problem). Since erroneously allocated datapages for image data in the Snapshots table may still be present in GPB_L0 as it is being backed up to tape, the space accounting does not work properly and end-of-tape may be reached prematurely. The workaround is to back up GPB_L0 to tape separately, setting the tape capacity much lower than for the regular tape jobs.

Jobs automatically scheduled with cron:

Weekly, after preload and dbcc's have completed (don't forget to manually change YYMMDD first):

/home/sybase/sybdba/script/tape_backup.csh science_pc <YYMMDD> science_pc_back
L2.GPB_L1.orbit_determ.gpb_3_4_3.GPB_DP.GBP_L1A init 210
backs up science_prod user dbs after preload on science_pc

Jobs/tasks to be initiated by the dba - See step-by-step instructions below:

Weekly on sci-crunch, well before tape backups are due to fire off:

Manually go in and change YYMMDD to reflect the date that the full backup was created on science_prod.

Weekly on sci-crunch, after full backups have been preloaded and dbcc's come out clean:

/home/sybase/sybdba/script/tape_backup.csh science_pc <YYMMDD> science_pc_back GPB_L0 init 120
backs up science_prod user dbs after preload on science_pc

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/<server>_back.log. The dba corrects the error and manually restarts the job as appropriate.

Step-by-step instructions for backing up databases to tape:

6.1.6.1 Always do a local tape dump, i.e. use the tape drive that is associated with the data server that you are backing up from. So, if you are backing up from science_prod, use the tape drive attached to sci-base, managed by the backup server science_prod_back. If you are backing up from science_pc, use the tape drive attached to sci-crunch, managed by the science_pc_back backup server. A remote tape dump takes up to 10 times longer than corresponding local dump.

6.1.6.2 Two sets of backup tapes will be created, Make sure the tape drive is set up for mirroring, i.e. the ASP-mode is "Mirror" as shown in the LCD-displays on the drive itself. If necessary, change the ASP-mode using the three push-buttons marked Menu, Select and Enter on the right-hand side of the tape drive. While making changes, current settings and messages are displayed in the left-hand LCD-display window. The settings cannot be changed while the tape drive is in use

```
Push Menu to initiate a change
=> Configuration menu
Push Enter to change a configuration
Push Select until "ASP mode" is displayed
Push Enter to change the ASP mode
Push Select until the desired ASP mode, Mirror, is displayed.
Push Enter when the correct ASP mode is displayed
Push Menu to exit
Push Enter to save the changes before exiting; Menu to cancel
Verify the settings are correct
```

6.1.6.3 If you wish to review what's on a tape that you are recycling, mount one tape at a time in the left bay and list it's contents with the commands

```
isql -Usa -S<server> -w222
load transaction tempdb from "/dev/rmt/0cn" with listonly=full
```

Dumpvolume is stored as "Volume id:", highlighted below. Note that the leading 0 in the year 04 is dropped since the info is stored as an integer. The file name is always the source database name.

```
Backup Server session id is:  43.  Use this value when executing the
'sp_volchanged' system stored procedure after fulfilling any volume change
request from the Backup Server.
Backup Server: 4.35.1.1: Device '/dev/rmt/0cn':
Label name:     'VOL1'
Volume id:      '40712 '
Access code:    ' '
Reserved:       '                          '
Owner id:       '              '
Reserved:       '                            '
Labeling
version:        7

Backup Server: 4.37.1.1: Device '/dev/rmt/0cn':
Label id:       'HDR1'
File name:      'GPB_L1A          '
Stripe count:   1
Device typecount:       1
Archive volume number:  1
Stripe position:        0
Generation number:      0001
Generation version:
00

Backup Server: 4.148.1.1:
Create date & time:     Wednesday, Jul 14, 2004, 14:35:26
Expiration date & time: Wednesday, Jul 14, 2004, 00:00:00
Access code:    ' '
File block count:       0
Sybase id string:       'Sybase  '
Reserved:       '   '
```

```
Etc. . . . .
```

6.1.6.4 Label the tapes using the smaller label that goes on the edge of the cassette. It should look like this:

Science db mm/dd/yy <db1.db2.db3…> YYMMDD #/## <AorB>	Science db 7/15/04 L2.L1.L1A 040715 1/2A

The date is the date as of which the databases are current, not the date the tape backup was made if dumping the standby databases. This date is also used to electronically label the tape with a dumpvolume, but in the format YYMMDD. List this value on the third line, followed by the sequence number of this tape, out of total number of tapes and add an A for the left bay, a B for the right bay. Last but not least, on the second line, list the databases present on each particular tape. (It speeds up recovery considerably if the load-program doesn't have to scan several tapes to find the database it is trying to recover).

6.1.6.5 Mount the first two tapes in the local tape drive. Make sure the tapes are not write-protected

6.1.6.6 Start the dump to tape. Because of some tape-handling sequences in the script, the load must be started from the host where the tape drive is attached. If starting the backup manually, ssh to the appropriate server. If starting the backup automatically using cron, add the job to the crontab on the host where the backup is run. See step 6.1.6.4 above for determining the dumpvolume value. In general the command looks like this

```
ssh -l sybase <host where tape drive is attached>
cd /home/sybase/sybdba/script
tape_backup.csh <source sybase server> <dumpvolume on tape>
      <backup server running on host where the tape drive is attached>
      <database(s)_to_dump> <dbcc/nodbcc switch> <unload/nounload switch>
```

- Weekly tape backup of preloaded user databases on science_pc:

The db's in this example were preloaded as of 6/15/04. A local tape-drive is used and the tape will be initialized, i.e. this is a new, blank tape or we wish to overwrite what's on an existing tape. The tapes will be unloaded when the backup completes:

```
Mount the tapes on the tape drive on sci-crunch
ssh -l sybase sci-crunch
tape_backup.csh science_pc 040615 L2.GPB_L1.GPB_L1A.orbit_determ.gpb_3_4_3
                        # defaults to local backup, all dbs, init, unload
```

- Contingency backup of user databases on science_prod directly to tape:

A local tape-drive is used and the tape will be initialized, i.e. this is a new, blank tape or we wish to overwrite what's on an existing tape. The tapes will be unloaded when the backup completes:

```
Mount the tapes on the tape drive on sci-base
ssh -l sybase sci-base
tape_backup.csh science_prod 041231
                        # defaults to local backup, all dbs, init, unload
```

- Append one database backup to an existing tape, do a remote backup:

In this example the tape drive on sci-base is not available when we want to add orbit_determ to an existing tape. This is a small db, it is ok to do a remote backup. Do not unload the tapes at the end.:

```
Mount the tapes on the tape drive on sci-crunch
ssh -l sybase sci-crunch      #this is where the tape drive is mounted
tape_backup.csh science_prod 041231 science_pc_back orbit_determ noinit
nounload
```

6.1.6.7 Be prepared to change tapes if the dump is larger than 210 GB, the amount of data that may be written to one tape, compressed. It takes 4-5 hours to write one tape volume, more on a heavily loaded system. When end-of-tape is reached the dump is suspended and an entry made in the backup-log indicating that the tape volumes need to be changed. This log-entry initiates a page to the dba with the same msg. If the database being backed up is actively being updated, it is

important to not let the dump be suspended indefinitely. Tranlogs are not dumped while a full backup is taking place and you may run the risk of the tranlog filling up. Sample backup-log entry:

```
Jul  6 22:13:17 2004: Backup Server: 6.54.1.1: OPERATOR: Mount the next volume to write
Jul  6 22:13:17 2004: Backup Server: 6.78.1.1: EXECUTE sp_volchanged
        @session_id = 130,
        @devname = '/dev/rmt/0cn',
        @action = { 'PROCEED' | 'RETRY' | 'ABORT' },
        @vname = <new_volume_name>
```

6.1.6.8 When additional tapes are requested, insert two new tapes into the tape drive. When the tapes are properly queued, notify the backup server that the tapes have been changed. The simplest is to run a script with the parameters dataserver and session id. The script may be run from any client. You may have to reissue the command several times, especially if overwriting previously used tapes:

```
/home/sybase/sybdba/script/vol_changed.csh science_pc 130
```

(If no session id-parameter is used, the script will cycle through all values from 1 to 499, trying to make an sp_volchanged take effect)

Another way to communicate with the backup session is to log in to the sybase server to issue sp_volchanged as noted in the backuplog.

```
isql -Usa -S<server> -w222
Copy and paste the sp_volchanged command from the backup-log.
For example: exec sp_volchanged
             @session_id = <id from log>,
             @dev_name = = "/dev/rmt/0cn",
             @action = "PROCEED"
```

6.1.6.9 Verify that the backup really was restarted. Tail the backup log. Are there more messages indicating you will need to reissue sp_volchanged to verify it is ok to proceed or is data now being dumped to the tape?

```
tail -f $SYBASE/ASE/ASE/install/<server>_back.log
Jul  6 22:16:17 2004: Backup Server: 6.54.1.1: OPERATOR: Volume on device '/dev/rmt/0cn' is
expired and will be over written.
Jul  6 22:16:17 2004: Backup Server: 6.78.1.1: EXECUTE sp_volchanged
        @session_id = 130,
        @devname = '/dev/rmt/0cn',
        @action = { 'PROCEED' | 'RETRY' | 'ABORT' },
        @vname = <new_volume_name>
Jul  6 22:18:55 2004: Backup Server: 6.28.1.1: Dumpfile name 'GPB_L1A          ' section
number 2 mounted on tape drive '/dev/rmt/0cn'
Jul  6 22:19:19 2004: Backup Server: 4.58.1.1: Database GPB_L1A: 108764564 kilobytes DUMPed.
Jul  6 22:20:37 2004: Backup Server: 4.58.1.1: Database GPB_L1A: 110719054 kilobytes DUMPed.
Jul  6 22:22:08 2004: Backup Server: 4.58.1.1: Database GPB_L1A: 112739674 kilobytes DUMPed.
Jul  6 22:23:32 2004: Backup Server: 4.58.1.1: Database GPB_L1A: 114694164 kilobytes DUMPed.
```

6.1.6.10 Keep feeding tapes and issue c until the tape backup completes

6.1.6.11 Write-protect all tapes by flipping the tabs on them to "safe" position.

6.1.6.12 Store the A-set of tapes in the DBA's office, currently in one of the rooms off the MOC, and the B-set of tapes in the box underneath the sci-crunch server in HEPL 127.

6.1.7 Verify the backup tape

Verify the tapes can be read by restoring at least one of the databases from the tape to the standby server. Later in the mission, as the databases grow larger, this is no longer possible to do for all tape backups. At that point, the databases are only restored from the tapes that will become part of the permanent archive and not recycled after three months. For the remainder of the tapes, perform a minimal verification by reading the labels on the tape, verifying that all databases are present.

Jobs automatically scheduled with cron:

None

Jobs/tasks to be initiated by the dba - See step-by-step instructions below:

Weekly on sci-crunch, after tape backups have been run:

Write-protect the tapes, move the tape to the bay where it was not created in and issue

/home/sybase/sybdba/script/tape_load.csh science_pc <YYMMDD> science_pc_back <db> same unload nodbcc

For tapes that will not be part of the permanent archive:

Isql –Usa –Sscience_pc

load tran tempdb from "/dev/rmt/0cn" with listonly=full, unload

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/science_pc_back.log. The dba corrects the error, which might entail recreating the tape dump. Redo the tape verification as appropriate.

6.1.8 Preload tranlog dumps from user databases to standby server

During IOC, transaction log dumps are continuously loaded to the standby server except for 24-48 hours when the full backups are loaded to the standby server, dbcc's are run and tape backups created. The tranlogs are loaded with at least a 24 hour delay to allow for the detection of data errors before they are reproduced on science_pc. Later in the mission, as the databases grow and preloads and tape backups take a lot longer, transaction logs will be preloaded less frequently.

Jobs automatically scheduled with cron:

2-4 times a day, on sci-crunch, except for when full backups are loaded, dbcc's and tape dumps run:

/home/sybase/sybdba/script/preload_tranlog.csh

Jobs/tasks to be initiated by the dba:

None. But additional preloads may be run as needed by issuing

/home/sybase/sybdba/script/preload_tranlog.csh <db1.db2.db3> <delay in hours>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/science_pc_back.log. The dba corrects the error and manually restarts the job as needed.

6.1.9 Copy tranlog dumps to backup disk

In case the backup disk holding the tranlog dump fails, these files are copied to /home/sybase/syblogs. This file system is backed up to tape, as part of the nightly and weekly OS-backups run by the SA. Because of space limitations the log files may only stay in this directory for 8 days before being removed. To speed up copying, files for the last three days only are copied.

Jobs automatically scheduled with cron:

Several times a day:

/home/sybase/sybdba/script/backup_tranlog_files.csh 3 #copies new log dumps to /home/sybase/syblogs

Jobs/tasks to be initiated by the dba:

None.

Dealing with errors:

For some errors that occur, the output from the job is mailed to the dba of record and a page sent to the dba-pgr. To ensure errors do not go undetected, a status job runs daily that shows what tranlog dump files reside in the /home/sybase/syblogs directory. It is up to the dba to verify this list is reasonable. The dba corrects the error and manually restarts the job as appropriate.

6.1.10 Cleanup tranlog dumps

Because of space limitations, the tranlog dump files may only stay in the /home/sybase/syblogs directory for 8 days. They are retained long enough to be backed up to several nightly and one weekly OS-backup tape. This facilitates recovery from any of the incremental or full backups if the primary backup disk for logdumps fails. Tranlog dump files are retained for 32 days in the primary repository, /sybackup/carin/syblogs, before they are removed.

Jobs automatically scheduled with cron:

Once a day:

/home/sybase/sybdba/script/cleanup_tranlog_files.csh

Jobs/tasks to be initiated by the dba:

None.

Dealing with errors:

For some errors that occur, the output from the job is mailed to the dba of record and a page sent to the dba-pgr. To ensure errors do not go undetected, a status job runs daily that shows what tranlog dump files reside in the /home/sybase/syblogs directory. It is up to the dba to verify this list is reasonable. The dba corrects the error and manually restarts the job as appropriate.

6.1.11 Comparing the production and standby database server setups

The purpose of the warm standby database server, science_pc, is to function as the production server during an emergency or planned maintenance, when the original production system, science_prod, is unavailable. In order for a switchover to the standby server to be seamless, the two database systems must be set up identically. Jobs are run daily that compare disk devices, database setups, logins and configurations between science_prod and science_pc. The dba is notified via email if differences between the two systems are found. The dba will then make the corrections necessary to realign the two systems. Whenever it is necessary to maintain slight variations in setups the comparison programs are adjusted to ignore these differences.

Jobs automatically scheduled with cron:

Once a day:

/home/sybase/sybdba/script/cmp_disk_init.csh science_prod science_pc
/home/sybase/sybdba/script/cmp_create_db.csh science_prod science_pc
/home/sybase/sybdba/script/cmp_logins.csh science_prod science_pc
/home/sybase/sybdba/script/cmp_cfg_file.csh science_prod science_pc

Jobs/tasks to be initiated by the dba:

None.

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record.The dba corrects the error and manually restarts the job as appropriate.

Dealing with differences found between the two systems:

When differences between the two systems are found, the diverging rows are mailed to the dba. The dba corrects the differences as appropriate and reruns the comparison job to make sure the systems now look the same.

- When disk devices are added or renamed, run sp_dropdevice and disk init's to adjust the servers.
- When databases have been added, add them in the exact same way, making sure that the dbid's and database allocations are identical.
- When databases have been extended, run alter database in the exact same way on the other server, using the lstart order in the sysusages table. If contiguous segments are allocated to the same device, online the database in between each allocation to possibly prevent the segments from being combined.
- When new logins have been added or passwords changed, bcp out master..syslogins, remove all rows except the ones to add or change, run sp_configure "allow updates, 1 on the target server. If the password has changed delete the old row from the master..syslogins table. bcp in the logins with batch size 1, run sp_configure "allow updates", 0 to turn off this option.
- Changes other than passwords in logins may be replicated by running sp_modifylogin on the other server
- Changes in configurations are duplicated by issuing the appropriate commands on the differing server, for instance sp_configure, sp_cacheconfig, sp_poolconfig

In some cases it is not possible to have the two systems look exactly the same. For instance, contiguous database allocations on the same disk device may have been combined by the dbms on one server and registered separately on the other. In cases like these, adjust the sql code that pulls the data to ignore the offending rows.

6.1.12 Prune the Snaptemp table

The GPB_L0..Snaptemp table is a holding table used for generating Snapshot data out of components that may come down at different times. Once the Snapshots have been built, the original Snaptemp records are no longer needed. After IOC, the Snaptemp table is pruned on a regular basis, to prevent it from getting too large.

Jobs automatically scheduled with cron:

Once a week:

/home/sybase/sybdba/script/prune_Snaptemp.csh science_prod GPB_L0 14 12 # keeps 14 days, cycle 12

Jobs/tasks to be initiated by the dba:

None.

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. The dba corrects the error and manually restarts the job as appropriate.

6.2 Backing up System Databases

Most jobs, both those automatically scheduled as well as those initiated by the dba, mail the log to the dba and send a page to the dba-pager whenever an error occurs. This system is not fail-safe however. If a job hangs and never completes or the job scheduler, cron, the mail system or mail forwarding is malfunctioning, the dba will simply not know that an error has occurred or that a crucial job did not even run.

It is the responsibility of the dba to routinely verify that jobs **are** run and that errors do not go undetected. This is easily done by reviewing the output in /home/sybase/sybdba/log for the pertinent jobs.

6.2.1 Perform consistency checking of system databases

The system databases, except for dbccdb, are all very small. Performing a complete set of dbccs, checkdb, checkalloc, textalloc, checkcatalog, takes 1-2 minutes. Dbcc checkstorage is not performed on these databases.

Jobs automatically scheduled with cron:

Several times a week, on science

/home/sybase/sybdba/script/dbcc.csh science_prod master.model.sybsystemprocs.sybsystemdb.sybdba

Several times a week, on sci-crunch

/home/sybase/sybdba/script/dbcc.csh science_pc master.model.sybsystemprocs.sybsystemdb

Jobs/tasks to be initiated by the dba:

None. But additional dbcc's may be run as needed by issuing

/home/sybase/sybdba/script/dbcc.csh <server> <db1.db2.db3> <dbcc_option1.dbcc_option2>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr.

Errors detected by dbcc checkcatalog, checkdb and checkalloc are listed with the appropriate error number. Look up the error in the Sybase Troubleshooting guide and perform the actions necessary to correct it. If appropriate, test the correction on science_pc first, then correct on the source server, science, generate a new backup and overwrite the corrupt database on sci-crunch.

6.2.2 Full backups of master databases to disk

Backups of the master database are made to the /sybackup1/carin/sybbackups/<server>/master directory.

In addition to the full database backups, operating system-level backups of the entire master device are made using dd. The first two blocks are copied out separately and then a complete dd-dump is made. The most critical tables in the master database are bulk-copied out and printouts generated to a file with the same information. Scripts for regenerating the disk devices and database allocations are created. Lastly, a copy is made of the configuration file. All this is written to abovementioned directory.

Ten generations of the backup files, typically representing one week, are retained in the /sybackup1/carin/sybbackups directory. The backup files are also copied to the Sybase home directory, which is backed up to tape as part of the OS-backups nightly and weekly, see below.

Jobs automatically scheduled with cron:

Daily on sci-base:

/home/sybase/sybdba/script/save_master_info.csh science_prod # master db and misc info

Daily on sci-crunch:

/home/sybase/sybdba/script/save_master_info csh science_pc # master db and misc info

Jobs/tasks to be initiated by the dba:

None. But master backup info **must** be saved off whenever changes are made to the server setups. This is especially important when changes are made to devices or databases. Failure to do so may mean that the entire system has to be rebuilt if the master database or device is lost.

ssh –l sybase <host where dataserver runs> # necessary for “dd"

/home/sybase/sybdba/script/save_master_info.csh <server>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error may be found in the backupserver log, $SYBASE/ASE/<server>_back.log. The dba corrects the error and manually restarts the job as appropriate.

6.2.3 Copy master database backups to backup disk

As a safeguard, in case the sci-crunch backup disk holding the master backups fails, the backup files are copied to the Sybase home directory nightly. Three generations of the files are kept there, typically representing three days worth of backups. The copy script completely replaces the existing master backups with the last three versions currently in the /sybackup1/carin/sybbackups directory.

Jobs automatically scheduled with cron:

Daily:

/home/sybase/sybdba/script/backup_master_dump_files.csh #loops thru all server directories

Jobs/tasks to be initiated by the dba:

None.

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. The dba corrects the error and manually restarts the job as appropriate.

6.2.4 Full backups of system databases other than master to disk

In general these databases are very small and there is very little update activity against them. With the exception of the sybdba database, which contains database space statistics, backups must be initiated by the dba whenever changes are made.

Backups are made to the /home/sybase/sybsystem directory. This directory is backed by the SA weeknights to tapes that are recycled weekly. Full backups for long-term archival are made to tape monthly by the SA.

Dbccdb is used for running dbcc checkstorage, mainly on sci-crunch. Due to its larger size, the database is backed up to the sci-crunch backup disk. This disk is not backed up to tape but dbccdb can easily be recreated according to the scripts in /home/sybase/sybdba/ddl/dbccdb.create.

Database	Updated	Backed up to	Cron/Manual	Comments
model	Never	/home/sybase/sybsystem	Manual	Necessary for sybase to function
sybsystemprocs	Seldom	/home/sybase/sybsystem	Manual	Necessary for complete functionality
sybsystemdb	Never	/home/sybase/sybsystem	Manual	Necessary for sybase to function
sybsyntax	Never	/home/sybase/sybsystem	Manual	Optional, helps w sql syntax
sybdba	Daily	/home/sybase/sybsystem	Cron	DB space stats, OK to lose data for a week or two
dbccdb	Daily	sci-crunch backup disk	Cron/Manual	Used for dbcc checkstorage, most on sci-crunch Easy to recreate, OK to lose data

Jobs automatically scheduled with cron:

Weekly on science:

/home/sybase/sybdba/script/full_backup.csh science_prod sybdba

Jobs/tasks to be initiated by the dba:

Whenever changes to the databases are made.

/home/sybase/sybdba/script/full_backup.csh <server> <db1.db2.db3..> <db backups to keep>

Dealing with errors:

When errors are encountered in the script the output from the job is mailed to the dba of record and a page sent to the dba-pgr. Additional information about the error can usually be found in the backupserver log, $SYBASE/ASE/<server>_back.log. The dba corrects the error and manually restarts the job as appropriate.

Summary Backup Activities

6.2.5 Automatically scheduled with Cron

Unless noted differently scripts are located in /home/sybase/sybdba/scripts

Task	Runs on host	Frequency	Script
Fix Snapshots allocation errors	sci-base	Weekly	dbcc_fix_text.csh
Full backups user db's to disk	sci-base	Weekly	full_backup.csh science_prod <dbs>
Full backups tqsm database to disk	sci-base	Weekly	full_backup.csh tqsm_server GPB_MOC
Full backups sybdba db to disk	sci-base	Weekly	full_backup.csh science_prod sybdba
Tranlog dumps to disk of user db's	sci-base	30 min	tranlog_backup.csh science_prod
Tranlog dumps to disk of tqsm database	sci-base	30 min	tranlog_backup.csh tqsm_server GPB_MOC
Preload full backups to standby (user db's)	sci-base	Weekly	preload_standby.csh
Dbcc checkdb/catalog of preloaded user db's	sci-base	Weekly	dbcc.csh science_pc <dbs> <options>
Dbcc checkstorage of preloaded user db's	sci-base	Daily	dbcc.csh science_pc <dbs> <options>
Tape backups	sci-crunch	Weekly	tape_backup.csh
Preload tranlogs to standby db's	sci-base	4/day IOC	preload_tranlog.csh
Cleanup tranlogs from /home/sybase/syblogs	sci-base	Daily	cleanup_tranlog.science.csh
Compare setups science_prod vs science_pc	sci-base	Daily	cmp_disk_init.csh science_prod science_pc
			cmp_create_db.csh science_prod science_pc
			cmp_cfg_file.csh science_prod science_pc
			cmp_logins.csh science_prod science_pc
Prune snaptemp	sci-base	Weekly	prune_Snaptemp.csh science_prod GPB_L0 14 12
Consistency checking of system db's	sci-base	3/week	dbcc.csh science_prod <dbs>
Consistency checking of system db's	sci-crunch	3/week	dbcc.csh science_pc <dbs>
Full backup of master db + important data	sci-base	Daily	save_master_info.csh science_prod
Full backup of master db + important data	sci-crunch	Daily	save_master_info.csh science_pc
Copy master dumps to Sybase home dir	sci-base	Daily	backup_master_dump_files.csh

6.2.6 Manually done by DBA

Unless noted differently script are located in /home/sybase/sybdba/scripts

Task	Runs on host	Frequency	Script
Adjust archive-date label on tape backups	sci-crunch	Weekly	Change YYMMDD in cron
Backup L0 to tape (custom tape capacity)	sci-crunch	Weekly	tape_backup.csh science_pc …. **120**
Verify backup tapes can be read	sci-crunch	Weekly	tape_load.csh science_pc yymmdd \ science_pc_back <db>
Preload tranlogs to standby db's	sci-base	Weekly post-IOC	preload_tranlog.csh
Backup misc. system db's	sci-base	As needed	full_backup.csh science_prod <db>
Backup misc. system db's	sci-crunch	As needed	full_backup.csh science_pc <db>