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

Database Recovery Procedures
Using the Standby Database Server

P1070 Revision C

August 9, 2005

Approvals

NAME	SIGNATURE	DATE
Carin Kahn Author		
Jennifer Spencer *Data Processing Lead*		
Rodney Torii *Data Processing IPT*		
Ron Sharbaugh *S/W Manager*		
Marcie Smith *MOC Project Manager*		
Kelly Burlingham *Software Quality Assurance*		

Required Signatures prior to Execution

NAME	SIGNATURE	DATE
Name: *Data Base Administrator*		
Name: *System Administrator*		
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
-	11/19/03	CWK	Initial version
A	3/8/04	CWK	Expansion of sections for maintenance and accessing science_offline. Addition of switchback after emergency switchover section. Changed order of switchover, loading dbs/tranlogs during an emergency switchover. Improved application testing sections and signoffs for testing. General enhancements based on two trial runs to increase robustness and improve communication
B	12/17/04	CWK	Production server now called science_prod, not science. Alternate procedure for L1A-maintenance. Minor corrections.
C	8/9/05	CWK	Adjust for additional hardware

2 Scope

In the event of a database failure or a system failure that may require database recovery and/or a switch to the standby server it is imperative that the database administrator or designee is contacted immediately. The dba/designee will help identify the problem and assess what recovery strategy to employ.

The purpose of this document is to provide step-by-step instructions for the various recovery scenarios that all involve switching production over to the standby database server and later switching back to the original production server. Instructions are provided for both scheduled and unscheduled switchovers. A scheduled switchover is typically executed in order to do maintenance on the original production databases. An unscheduled switchover is initiated when a database or system failure occurs and recovery on the production system cannot be done within the allotted twenty-four hour recovery window or whenever a switchover is deemed appropriate in order to get the users back online as soon as possible. The document also addresses accessing the original production server and doing database maintenance during the switchover.

This document does not address how the database and system failures are assessed and resolutions arrived at, that process is covered in document P1081, "Assessment of Science Network Database Errors and Failures". This document assumes that a decision has already been made, through the proper channels, to switch over to the standby server.

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 for Science Network Databases	P0996	
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

This procedure may only be conducted by the following persons:

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

5 QA Provisions

This procedure shall be conducted on a formal basis to its latest approved and released version. Software QA (D.Ross) shall be notified 24 hours prior to the start of this procedure. QA may monitor the execution of all or part of this procedure should they elect to do so.

In case of an emergency it will not be possible to adhere to the 24 hour notification-rule. If the Data Processing IPT gives the go-ahead to implement a procedure, QA is notified by pager and by phone.

RED-LINE AUTHORITY:

Authority to red-line (make minor changes during execution) this procedure is given solely to the DBA and shall be approved by QA.

QA NOTIFICATION: Date/time:_____
GP-B QA (D. Ross)

6 Recovery Scenario Start

Recovery Scenario Used: Scheduled/Unscheduled started on (date&time) _____
ended on (date&time) _____

Executed by _____ Signature: _____

Witnessed by _____ Signature: _____

Switchback Scenario Used: Scheduled/Unscheduled: started on (date&time) _____
ended on (date&time) _____

Executed by _____ Signature: _____

Witnessed by _____ Signature: _____

7 Recovery Scenarios Using the Standby Server

The following procedures should be used when implementing a scheduled or unscheduled (emergency) switchover from the production database server to the standby database server.

The reason for a scheduled switchover is typically that maintenance has to be done on the production system in order to assure good long-term performance and space utilization. These activities severely limit and often block user access during extended periods of time. In order to minimize user downtime, users are switched over to a standby system while the production system is unavailable and then switched back once maintenance has completed. The impact on users is minimized, two short system outages, typically lasting less than one hour each, while the actual switchovers are effected. No data is lost, user-initiated database updates are carried from one system to another, and vehicle data files continue to be processed and immediately accessible to the users while the original production databases are being maintained.

An unscheduled, emergency switchover will be initiated when the production system unexpectedly becomes unavailable due to hardware or software problems and it is furthermore deemed unlikely that the problems can be corrected and the production system made fully operational within the allotted twenty-four hour recovery window. In many situations it will make sense to start preparing for an emergency switchover while still assuming the production system can be readily fixed and the switchover won't have to take place. If the system outage makes the production databases inaccessible to the database administrator, updates made to Level 2, MSS and OD databases since the last transaction log dump will be lost. Data files that were processed after the last transaction log dump occurred must be reprocessed to prevent data loss.

Section 8 below, "Accessing science_prod while science_pc is production server", explains how to access the original production server for maintenance and database repairs while the standby server, science_pc, acts as the production server of record. Section 9 below, "Performing Maintenance", shows how maintenance can be performed on Level 0 and 1 as well as Level2, MSS and OD databases.

Once maintenance has been completed or the problems causing the emergency switchover have been corrected, the users need to be switched back to the production system. The procedure for reversing a switchover is detailed in sections 7.3, "Switching Back to the Production Server ", and 7.4 below, "Switching Back to Production Server After an Emergency Switchover".

7.1 Scheduled Switchover to Standby Server

Well before the scheduled switchover (a few days to a week):

7.1.1 Determine date of full production backups to pre-load to the standby server. If appropriate, schedule extra full backups and preloads before the switchover. Make sure there will be time to preload the backups, run full dbcc's and preload transaction log dumps well before the switchover _____

7.1.2 Estimate time required for preparing for the switchover, performing the switchover, doing database maintenance, reprocessing data files against the offlined dataserver, preparing for the switchback. Pad the time to allow for unanticipated delays _____

7.1.3 Based on what kind of maintenance activities will take place, determine when to do maintenance on Level 2, MSS, and OD databases. Add this time to outage estimates for switchover or switchback. See section 9.1 below, "Maintenance of Level2, MSS and OD Databases", to decide on strategy.

_____

7.1.4 Schedule the switchover as well as the anticipated switchback with the mission/flight director. This may be done by the Data IPT

Activity	Time required	Planned Start	Planned End	Actual **Start**	Actual **End**	Actual Time
Prepare standby for switchover						
Switchover						
Optional: L2 etc. maintenance						
Maintenance level 0/1						
Full backup						
Reprocess data files against science_offline						
Prepare switchback						
Switchback						
Optional: L2 etc. maintenance						

7.1.5 Confirm the timing with the data processing team, their schedule will be affected. Also remind the DP team that they will have to reprocess the data for the period when science_pc is the server of record and load the data to science_offline/science_prod before the switchback _____

7.1.6 Start the extra full backups/preloads as appropriate

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

_____

7.1.7 Notify QA of the coming switchover. QA & System Administrator to sign off on P1070doc.

_____

7.1.8 Notify users of the planned outage that will occur while operation is switched to the standby system. The outage will last one-two hours excluding maintenance on Level 2, MSS and OD databases. Estimate additional outage for these databases if maintenance will be performed during this switchover and not the switchback. Send notification to all@relgyro.stanford.edu with cc's to fd & arb@relgyro.stanford.edu _____

Once the desired production databases have been preloaded to the standby server:

7.1.9 Put the crontab jobs preload_standby & preload_tranlog jobs on hold. (i.e. comment out) _____

7.1.10 Check the dbcc's on the standby server. Correct any problems found: first on standby db to verify the fix works, then on production db, redo the backup/preload of the affected database(s)_____

7.1.11 Optional. Start full backups of newly loaded databases on science_pc (or you may wait til AFTER the switchover is complete)

```
/home/sybase/sybdba/script/tape_backup.csh science_pc      <YYMMDD>
```

_____

7.1.12 Prep the backup directory for storing future tranlog dumps for all the user databases post-switchover. That is, verify /sybackup/carin/syblogs/science_pc/<user-db>/gen_0 and /sybackup1/carin/sybbackups/science_pc are empty _____

7.1.13 Add a crontab job that preloads tranlogs without using a time-lag and runs more frequently. The command line should be similar to this (last param is the number 0):

```
15 7,13,19,1 * * * /home/sybase/sybdba/script/preload_tranlog.csh all 0
```

_____

Four – five hours before the switchover:

7.1.14 Stop data processing cron. Login as tdp on moc-server, issue cron stop, cron status _____

Once the last data file has finished processing or about two hours before the switchover at the latest:

7.1.15 Broadcast a warning over the voice-loop that the science Sybase server will be coming down in a couple of hours for a planned switchover to standby server _____

7.1.16 Verify with the data processing team that they are on target for processing files _____

7.1.17 Remove the currently active preload_transaction job from the crontab schedule. _____

7.1.18 Put all backup jobs that run against the production sybase server on hold by commenting out in crontab. (full_backup, tranlog_backup, etc.) _____

7.1.19 Review all other jobs that run against the production sybase server and put on hold as appropriate by commenting out in crontab _____

7.1.20 Dump the tranlogs and preload to the standby server to make sure Level 0 and 1 databases will be fully preloaded before the switchover is initiated and thus not delay it. Use the commands

```
/home/sybase/sybdba/script/tranlog_backup.csh science_prod
/home/sybase/sybdba/script/preload_tranlog.csh all 0  ##number 0
```

_____

7.1.21 Keep running tranlog dumps against science_prod and preloads on science_pc as time allows.

```
/home/sybase/sybdba/script/tranlog_backup.csh  science_prod
/home/sybase/sybdba/script/preload_tranlog.csh all 0  ##number 0
```

_____

ACTUAL SWITCHOVER:

7.1.22 Prevent users from logging in to the science_prod sybase server by locking all logins except "sa", "carin" and "probe". Log in to the science_prod server as a system administrator and issue:

```
isql –U<user> –Sscience_prod
select 'exec sp_locklogin ' + name +', '"lock"'
      from master..syslogins where name not in ("sa", "carin", "probe")
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 != 2
```

The only logins returned should be carin, sa and probe _____

7.1.23 Kill any active Sybase processes in at least Level 2, MSS and OD databases.

```
/home/sybase/sybdba/script/kill_spids science_prod L2
```

```
/home/sybase/sybdba/script/kill_spids science_prod gpb_3_4_3 ##current MSS
/home/sybase/sybdba/script/kill_spids science_prod orbit_determ
```

7.1.24 Manually run the last transaction log dump on the production server and preload on standby server

```
/home/sybase/sybdba/script/tranlog_backup.csh  science_prod
/home/sybase/sybdba/script/preload_tranlog.csh all 0  ##number 0
```

7.1.25 Fix the name of the SYB_BACKUP server to facilitate starting science_prod as science_offline later

```
isql -U<user> -Sscience_prod
shutdown SYB_BACKUP [with nowait]   -- shut down first to be safe
sp_helpserver
sp_dropserver  SYB_BACKUP
sp_addserver   SYB_BACKUP, null, science_offline_back
sp_helpserver
```

7.1.26 Shut down both production and standby Sybase servers to force all connections to be severed.

```
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_stop
ssh -l sybase sci-crunch
/apps/licensed/sybase_local/etc/sybase_stop
```

7.1.27 Switch interfaces files on science to redirect users to sci-crunch and set up alias science_offline for the original science_prod server.

```
ssh -l sybase science
cd $SYBASE  ##/apps/licensed/sybase-12.5
cp interfaces.standby interfaces
```

7.1.28 Start the standby Sybase server. It will now be known as the science & science_prod server as well.

```
ssh -l sybase sci-crunch
/apps/licensed/sybase_local/etc/sybase_start
```

7.1.29 Verify you'll really be connected to science_pc when using science or science_prod

```
isql -U<user> -Sscience_prod
select @@servername           -- this should show science_pc
isql -U<user> -Sscience
select @@servername           -- this should show science_pc
```

7.1.30 Bring all user databases on science_pc online in order to remove the standby option. Substitute the current MSS database for database gpb_3_4_3 below:

```
isql -U<user> -Sscience_prod
select @@servername           -- this should show science_pc
online database GPB_L0
online database GPB_L1
online database GPB_L1A
online database L2
online database gpb_3_4_3     -- change to current MSS database
online database orbit_determ
online database GPB_DP
```

7.1.31 Unlock all the previously locked users on the standby server.

```
isql -U<user> -Sscience_pc
select 'exec sp_locklogin ' + name +', '"unlock"'
      from master..syslogins where status & 2 = 2
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 = 2
```

No logins should be returned. (unless there are logins locked for other reasons)

7.1.32 Switch interfaces files on the moc-server to redirect users to sci-crunch - Moc-authorization required

```
log on to the moc server as user database
```

```
cd /apps/licensed/sybase-12.5.0.3
cp interfaces.standby interfaces
cd /apps/licensed/sybase-11.9.2
cp interfaces.standby interfaces
```

7.1.33 Test connectivity for mission planning against science_pc (gpb_n_n_n, orbit_determ)

```
Log on to a moc-client with your personal user-id
MP; userid cmdops; pick MSS db; Param Gen; Application; GPS Ephemeris; Report
A valid report should be displayed
```

7.1.34 Test tcad connectivity on a science-client, verifying it is connecting to science_pc

```
Log on to a science-client with your personal user-id
tcad; select DISPLAY & TMdata; a few telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

7.1.35 Test tcad connectivity on a moc-client, verifying it is connecting to science_pc

```
Log on to a moc-client with your personal user-id
tcad; select DISPLAY & TMdata; pick telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

7.1.36 Test connectivity for the science application, L2. (The app dynamically reads the interfaces file)

```
Log on to a science-client (sybase user-id is fine)
matlab; ScienceDataLook('SIM5-part3'); click Separate Figures
A plot should appear in separate window (SQ_SciLPasFilt1)
```

7.1.37 Connectivity Signoff. Get signoffs by the dba/sa/engineer that tested each app for connectivity

Application	Tested on	Tested by dba/sysadmin/eng	Date/Time	Initial
Mission Planning	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.1.38 Optionally, if performing maintenance of Level 2, MSS, OD and/or GPB databases:

- Make the databases you will do maintenance on unavailable to the end users

  For databases owned by sa:

  ```
  exec sp_dboption <dbname>, "dbo use", true
  use <dbname>
  checkpoint
  ```

  For databases owned by the actual user:

  ```
  use <dbname>
  sp_changedbowner sa
  ```

- Perform the maintenance as described in section 9.1, "Maintenance of Level2, MSS and OD Databases", below. Start maintenance with MSS and OD databases, then proceed to Level 2 database. As each database is done, make available to users by undoing the dbo-use only switch or reverting the ownership to the original owner

7.1.39 Restart Data processing cron. Login as tdp on moc-server, issue cron start, cron status

7.1.40 Application Signoff. Contact the leads for each application and ask them to verify that the databases are up-to-date and fully functional. Ask them to sign off in the table below as an acknowledgement and attach their check-lists, if used

Application	Verified on	Verified by team lead/designee	Date/Time	Initial
Mission Planning	Moc-client			
Orbit Determ	Moc-client			
Data Processing	Moc-client			
TCAD	Science-client			

TCAD	Moc-client			
Science	Science-client			

7.1.41 Notify users that the science Sybase server is now available, running off the standby server. Send email to all@relgyro.stanford.edu and/or announce over the voice-loop as appropriate. If access to a database is delayed due to maintenance, indicate this in the announcements and give an estimated time of availability for each database

_____

7.1.42 Start full backups of science_pc dbs unless it's been done already
In some cases you may want to do full backups of science_prod as well _____

7.1.43 Schedule transaction dumps for science_pc in crontab on sci-crunch

```
05,35 * * * * /home/sybase/sybdba/script/tranlog_backup.csh science_pc
```

_____

7.1.44 Start science_prod as science_offline. See section 8.1 below, "Starting the data and backup servers on sci-base"

_____

You are now ready to start doing the bulk of the maintenance on the science Sybase server. See section 9.2 below, "Maintenance of Level 0 and 1 Databases on the Offline Server – Database space is at a premium", for instructions on how to perform these maintenance tasks.

7.2 Emergency Switchover to Standby Server

Before initiating an emergency switchover to the standby server it is important to determine whether the databases currently preloaded are the correct starting point. For instance, if the emergency switchover is necessitated by corrupted production databases, it is important to determine whether the standby databases have been corrupted as well. Document, "Assessment of Database Errors and Failures", guides the dba through the assessment process. If necessary, refer to document P1071, "Database Recovery Procedures – Single Database Recovery", for instructions on how to recover a specific database backup. In some cases we may know the approximate time the corruption occurred. If that is the case we may chose to preload the transaction logs for the database(s) in question up to before the time of corruption. This is all described in detail in the aforementioned document as well as cursory below.

A decision has been made by ARB/mission/flight director to switch over to the standby

7.2.1 Notify QA, System Administrator and Data IPT Lead that an emergency switchover will be initiated. If possible, get QA and SA signatures of approval on the P1070 printout used (this doc) _____

7.2.2 The dba or Data IPT Lead makes sure that the pending emergency switchover is announced to the user community. The announcement may be made by ARB/MD/FD or the dba/Data IPT Lead. An estimated time of availability for the different applications should be given as well as the time/date as of which the data will be restored. (May use one-page db briefing tool)

_____

7.2.3 Put all automatically scheduled jobs that run against the production or standby database server on hold, i.e. comment out in crontab on sci-base and sci-crunch. This includes backup jobs, preload jobs etc.

_____

7.2.4 Stop Data processing cron. Login as tdp on moc-server, issue cron stop, cron status _____

7.2.5

If the production server is still running and somewhat operational:

7.2.6 Make the database system unavailable by locking out the users.

```
select 'exec sp_locklogin ' + name +', '"lock"'
      from master..syslogins where name not in ("sa",  "carin", "probe")
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 != 2
```

The only logins returned should be carin, sa and probe _____

7.2.7 Kill any active Sybase processes in at least Level 2, MSS and OD databases.

```
/home/sybase/sybdba/script/kill_spids science_prod L2
/home/sybase/sybdba/script/kill_spids science_prod gpb_3_4_3 ##current MSS
/home/sybase/sybdba/script/kill_spids science_prod orbit_determ
```

_____

7.2.8 Try dumping the transaction logs

```
/home/sybase/sybdba/script/backup_tranlog.csh science
```

_____

If that doesn't work, for every database that is accessible

```
isql -Usa -Sscience
dump tran <db> to
"/sybackup/carin/syblogs/<server>/<db>/gen_0/<db>.log.CCYYMMDDHHMMSS"
with no_truncate
```

_____

7.2.9 Fix the name of the SYB_BACKUP server to facilitate starting science_prod as science_offline later

```
isql -U<user> -Sscience_prod
shutdown SYB_BACKUP [with nowait]   -- shut down first to be safe
sp_helpserver
```

```
sp_dropserver SYB_BACKUP
sp_addserver   SYB_BACKUP, null, science_offline_back
sp_helpserver
```

Optional: The standby server is currently undergoing a preload of full backups

7.2.10 If a full preload just started and it's a large db; kill the job and restart after the switchover is done

```
ps -fu sybase|grep -E "load|recover|sybmultbuf"
kill spids in unix
As user sybase, on any unix host: source ~/sybdba/tools/aliases
who.p
if necessary, `login to sybase and kill <pid> obtained
```

7.2.11 If the standby server is currently being preloaded with the latest full backups, as part of the weekly routine, allow the current load to complete if it's not going to delay the switchover too much. Make sure that no more full preloads are started. The latter is done by simply changing the names of the backup files that still remain to be preloaded. The easiest is to change the text "full" in the name to "xxxx". Do not change the date-portion of the name, it is insignificant in the preload script. The output from the job shows progress and the tag-line "The dblist contains" shows which order the databases are loaded in.

```
To determine order of databases to load:
head /home/sybase/sybdba/log/preload_standby.science_pc.<dbs/all>.log.<date>
To review progress of the ongoing load:
tail -f /home/sybase/sybdba/log/preload_standby.science_pc.<dbs/all>.log.<date>
To prevent preloading of a database:
cd /sybackup1/carin/sybbackups
# for each database not yet preloaded with new full backup:
Cd <db_not_refreshed>/gen_0
mv <db_not_refreshed>.full.<date> <db_not_refreshed>.xxxx.<date>
```

7.2.12 If you are allowing a full load to proceed you may go ahead and start preloading tranlogs for the most time-critical databases. Don't overdo it, it'll slow down the full backup somewhat

```
/home/sybase/sybdba/script/preload_tranlog.csh gpb_3_4_3.orbit_determ
```

ACTUAL SWITCHOVER:

7.2.13 Shut down the standby Sybase servers to sever all connections to it. (science_prod is down already)

```
ssh -l sybase sci-crunch
/apps/licensed/sybase_local/etc/sybase_stop
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_stop
```

7.2.14 Switch interfaces files on science to redirect users to sci-crunch and set up alias science_offline for the original science_prod server.

```
ssh -l sybase science
cd $SYBASE  ##/apps/licensed/sybase-12.5
cp interfaces.standby interfaces
```

7.2.15 Start the standby Sybase server. It will now be known as the science & science_prod server as well.

```
ssh -l sybase sci-crunch
/apps/licensed/sybase_local/etc/sybase_start
```

7.2.16 Verify you'll really be connected to science_pc when using science or science_prod

```
isql -U<user> -Sscience_prod
select @@servername           -- this should show science_pc
isql -U<user> -Sscience
select @@servername           -- this should show science_pc
```

7.2.17 If you interrupted a database preload in step 7.2.11 above, restart that preload now _____

7.2.18 Do NOT bring the user databases on science_pc online at this point!!! You need to preload the tranlogs first, to bring them up to date _____

7.2.19 Unlock all the previously locked users on the standby server.

```
isql -U<user> -Sscience_pc
select 'exec sp_locklogin ' + name +', '"unlock"'
      from master..syslogins where status & 2 = 2
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 = 2
```

No logins should be returned. (unless there are logins locked for other reasons) _____

7.2.20 Switch interfaces files on the moc-server to redirect users to sci-crunch - Moc-authorization required

```
log on to the moc server as user database
cd /apps/licensed/sybase-12.5.0.3
cp interfaces.standby interfaces
cd /apps/licensed/sybase-11.9.2
cp interfaces.standby interfaces
```

_____

7.2.21 Test connectivity for mission planning against science_pc (gpb_n_n_n, orbit_determ)

```
Log on to a moc-client with your personal user-id
MP; userid cmdops; pick MSS db; Param Gen; Application; GPS Ephemeris; Report
A valid report should be displayed
```

_____

7.2.22 Test tcad connectivity on a science-client, verifying it is connecting to science_pc

```
Log on to a science-client with your personal user-id
tcad; select DISPLAY & TMdata; a few telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

_____

7.2.23 Test tcad connectivity on a moc-client, verifying it is connecting to science_pc

```
Log on to a moc-client with your personal user-id
tcad; select DISPLAY & TMdata; pick telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

_____

7.2.24 Test connectivity for the science application, L2. (The app dynamically reads the interfaces file)

```
Log on to a science-client (sybase user-id is fine)
matlab; ScienceDataLook('SIM5-part3'); click Separate Figures
A plot should appear in separate window (SQ_SciLPasFilt1)
```

_____

7.2.25 Connectivity Signoff. Get signoffs by the dba/sa/engineer that tested each app for connectivity

Application	Tested on	Tested by dba/sysadmin/eng	Date/Time	Initial
Mission Planning	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.2.26 Optionally, preload logs for a particular database only up to a particular point in time, for instance

```
home/sybase/sybdba/script/preload_tranlog.csh GPB_L1A      0 "4/17/04 21:50"
```

_____

7.2.27 Preload the tranlogs. Start with the most time-critical databases as the situation necessitates.

```
/home/sybase/sybdba/script/preload_tranlog.csh gpb_3_4_3.orbit_determ
OR if there's no criticality; run maybe two load streams in parallel
/home/sybase/sybdba/script/preload_tranlog.csh GPB_L0.GPB_L1.L2.GPB_DP.orbit_determ
/home/sybase/sybdba/script/preload_tranlog.csh gpb_3_4_3.GPB_L1A
```

_____

7.2.28 Optionally, if one or several standby databases are corrupted, load them from backup. Document P1071, "Database Recovery Procedures for Science Network Databases", details how to recover in various situations from diverse media. If you simply need to recover the most recent backup, issue

```
/home/sybase/sybdba/script/preload_standby.csh <db1>.<db2> etc
```

_____

7.2.29 As each set of databases, a set defined by which dbs are needed for a particular application, is fully recovered, perform any checks that this particular failure call for and bring the database online when completely satisfied. Do not online the database until you are sure it is the right thing to do, any remaining tranlogs cannot be loaded thereafter. Quickly retest connectivity per app.

```
isql –Usa –Sscience_pc
online <database>
sp_helpdb <database>
```

Application	**Databases**	**online <database>**	**Retest Connectivity**
Mission Planning	gpb_n_n_n orbit_determ		(Test from science-client)
Science	L2		
TDP/TCAD	GPB_L0 GPB_L1 GPB_L1A		(Test from science-client
	GPB_DP		(Just verify it's online)

7.2.30 As each application comes online, contact the lead for that application and ask them to verify that the database is up-to-date and fully functional. Make sure to inform the lead as of which date each database was recovered. The application leads are responsible for recapturing lost data _____

7.2.31 Restart Data processing cron. Login as tdp on moc-server, issue cron start, cron status _____

7.2.32 Application Signoff. As each application comes online, the lead for that application is contacted and asked to verify that the databases are up-to-date and fully functional. The leads will report back the status of the application/data to the dba, noting possible data loss and plans for recapture. The team leads will sign off the sheet below and supply check-lists to attach to this document as appropriate.

Application	Verified on	Verified by team lead/designee	Date/Time	Initial
Mission Planning	Moc-client			
Orbit Determ	Moc-client			
Data Processing	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.2.33 Report back to ARB/MD/FD when system has been fully recovered, stating when each application became available and any data loss. Use the briefing tool in document, "Assessment of Database Errors and Failures" _____

7.2.34 Notify the user community that the Sybase databases are now available. Indicate as of what time/date the data has been recovered in the different applications and what, if any, efforts are made to recapture lost data. Send email to all@relgyro.stanford.edu and/or announce over the voice-loop as appropriate _____

7.2.35 Schedule full backups for science_pc. This can be done to tape
/home/sybase/sybdba/script/tape_backup.csh science_pc <YYMMDD> _____

7.2.36 Once full backups have run, schedule transaction dumps for science_pc in crontab
`05,35 * * * * /home/sybase/sybdba/script/tranlog_backup.csh science_pc` _____

7.2.37

You are now ready to do research or repairs on the off-lined science Sybase server. See section 9.2 below, "Maintenance of Level 0 and 1 Databases on the Offline Server – Database space is at a premium", for instructions on how to access the off-lined science server.

7.3 Switching Back to the Production Server After a Scheduled Switchover

This scenario is to be used when switching back to the original production server after a scheduled maintenance outage. This scenario may also be used in certain situations after an emergency switchover where the original Level 0 and 1 databases on the production server were not damaged.

7.3.1 The switchback time should have been scheduled when the switchover procedure was initiated. If not, do it now. See steps 7.1.1 through 7.1.5 above _____

<u>As soon as maintenance has completed on Level 0 and 1 databases</u>

7.3.2 If time permits, run dbcc's and update statistics on Level 0 and 1 databases on science_offline. This applies only to db's that had maintenance done. Fix errors _____

7.3.3 Start full backups of Level 0 and 1 databases on the off-lined production server (only db's where maintenance was done). Keep the trunc log option on while the data files are reprocessed.

```
ssh -l sybase science
cd /home/sybase/sybdba/script/
local_tape_backup.csh science_pc <YYMMDD> GPB_L0.GPB_L1.GPB_L1A
```
_____

<u>As soon as full backups have been made of Level 0 and 1 databases</u>

7.3.4 Tell the data processing team to start reprocessing data files against the off-lined science server and ask when they expect to be **completely** caught up. This includes both reprocessing the files that have come in up to that point in time as well as files that will be arriving while the reprocessing is still going on. By the time of the switchback, the level 0 and 1 databases on science_offline where maintenance was done and the standby server should have had the exact same set of data files processed. See section 9.2 below, "Maintenance of Level 0 and 1 Databases on the Offline Server – Database space is at a premium", for details on reprocessing the data files.

_____

7.3.5 Remind mission/flight director of the impending switchback. If necessary, adjust time depending on new needs or changed time estimates from data processing _____

7.3.6 Notify QA, SA and Data IPT Lead of the coming switchback, indicating estimated start time._____

7.3.7 Notify users of the planned outage that will occur while operation is switched back to the production system. The outage will last one-two hours plus the time required for maintenance of Level 2, MSS and OD databases if it is to be done here. Send notification to all@relgyro.stanford.edu _____

<u>Four – five hours before the switchover:</u>

7.3.8 Stop data processing cron. Login as tdp on moc-server, issue cron stop, cron status _____

<u>A couple of hours before the switchback:</u>

7.3.9 Verify with the data processing team that all data files have been reprocessed against the offline production server. _____

7.3.10 Broadcast a warning over the voice-loop that the science Sybase server will be coming down in a couple of hours for a switchback. _____

7.3.11 Backup the databases on science_pc that will be copied back over to the original, offline science server. Do NOT include level 0 and 1 databases that had maintenance done on science_offline

```
/home/sybase/sybdba/script/full_backups.csh science_pc \
    L2.orbit_determ.gpb_3_4_3.GPB_DP    #current MSS
```
_____

7.3.12 Preload the full backups generated in step 7.3.11 above to the offline science server. The switchback script automatically excludes Level 0 and 1 databases when "all" databases are loaded.

```
chmod u+x /home/sybase/sybdba/script/switchback_offline_db.csh
/home/sybase/sybdba/script/switchback_offline_db.csh  <list_of_dbs>
chmod a-x /home/sybase/sybdba/script/switchback_offline_db.csh
```

7.3.13 Run and check dbcc's on the preloaded databases if time permits

7.3.14 Keep preloading the tranlogs onto the off-lined production server as they are created. This is done to reduce the total outage time. The switchback script automatically excludes Level 0 and 1 databases when "all" databases are loaded

```
chmod u+x /home/sybase/sybdba/script/switchback_offline_log.csh
/home/sybase/sybdba/script switchback_offline_log.csh     <list of dbs>
```

7.3.15 Cancel sci-crunch crontab job against science_pc, for instance backing up the db and log

7.3.16 If time permits, do a full backup of the databases just copied over. (change param to current MSS db)

```
/home/sybase/sybdba/script science_offline L2.orbit_determ.gpb_3_4_3.GPB_DP
```

ACTUAL SWITCHOVER:

7.3.17 Prevent users from logging in to the Sybase science_pc server by locking all logins except "sa", "carin" and "probe". That is, log in to the science_pc server as a system administrator and run:

```
isql-U<user> -Sscience_pc
select 'exec sp_locklogin ' + name +', '"lock"'
from master..syslogins where name not in ("sa", "carin", "probe")
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 != 2
```

The only logins returned should be carin, sa and probe

7.3.18 Kill any active processes in Level 2, MSS and OD databases. Issue commands:

```
/home/sybase/sybdba/script/kill_spids science_pc L2
/home/sybase/sybdba/script/kill_spids science_pc gpb_3_4_2 #chg to current
/home/sybase/sybdba/script/kill_spids science_pc orbit_determ
```

7.3.19 Generate the last tranlog dumps from Level2, MSS and OD databases on science_pc (and any other db that will be copied from science_pc to science_offline)

```
/home/sybase/sybdba/script/tranlog_backup.csh  \
      science_pc  L2.gpb_3_4_3.orbit_determ.GPB_DP #current MSS
```

7.3.20 Load the last Level 2, MSS and OD tranlog dumps to the offline science server (and any other db that will be copied from science_pc to science_offline). The switchback script automatically excludes Level 0 and 1 databases when "all" databases are loaded

```
/home/sybase/sybdba/script/switchback_offline_log.csh <list_of_dbs>
chmod a-x /home/sybase/sybdba/script/switchback_offline_log.csh
```

7.3.21 Rename the science_offline backup server so it can be restarted as science_prod_back later

```
isql -Usa -Sscience_offline
shutdown SYB_BACKUP [with nowait]   -- shut down first to be safe
sp_helpserver
sp_dropserver SYB_BACKUP
sp_addserver   SYB_BACKUP, null, science_prod_back
sp_helpserver
```

7.3.22 Shut down both standby and off-lined production Sybase servers to sever all connections to them

```
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_stop
ssh -l sybase sci-crunch
/apps/licensed/sybase_local/etc/sybase_stop
```

7.3.23 Switch interfaces files on science to redirect users back to science

```
ssh -l sybase science
cd $SYBASE  ##/apps/licensed/sybase-12.5
cp interfaces.production interfaces
```

7.3.24 Start the original production Sybase server. It will now be known as science_prod

```
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_start
```

7.3.25 Verify you'll really be connected to science_prod when using science or science_prod

```
isql -U<user> -Sscience_prod
select @@servername   -- this should once again show science_prod
isql -U<user> -Sscience
select @@servername   -- this should once again show science_prod
```

7.3.26 Bring all user databases on science_prod online in order to remove the standby option. Substitute the current MSS databases for database gpb_3_4_3 below:

```
isql -U<user> -Sscience_prod
select @@servername           -- this should show science_prod
online database L2
online database gpb_3_4_3     -- change this to current MSS database
online database orbit_determ
online database GPB_DP
plus others that were loaded from science_pc
```

7.3.27 Reset database options for the science_prod server to make it function as a production server.

```
isql -U<user> -Sscience_prod
use master
go
exec sp_dboption GPB_L0, "trunc log", false
exec sp_dboption GPB_L0, "select into", false
exec sp_dboption GPB_L1, "trunc log", false
exec sp_dboption GPB_L1, "select into", false
exec sp_dboption GPB_L1A, "trunc log", false
exec sp_dboption GPB_L1A, "select into", false
go
use GPB_L0; GPB_L1; GPB_L1A etc
checkpoint
sp_helpdb
```

7.3.28 Unlock all the previously locked users on the production server.

```
isql -U<user> -Sscience_prod
select 'exec sp_locklogin ' + name +', '"unlock"'
      from master..syslogins where status & 2 = 2
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 = 2
```

No logins should be returned. (unless there are logins locked for other reasons)

7.3.29 Switch interfaces files on the moc-server to point users back to science_prod - Moc-authorization required

```
log on to the moc server as user database
cd /apps/licensed/sybase-12.5.0.3
cp interfaces.production interfaces
cd /apps/licensed/sybase-11.9.2
cp interfaces.production interfaces
```

7.3.30 Test connectivity for mission planning against science_prod (gpb_n_n_n, orbit_determ)

```
Log on to a moc-client with your personal user-id
MP; userid cmdops; pick MSS db; Param Gen; Application; GPS Ephemeris; Report
A valid report should be displayed
```

7.3.31 Test tcad connectivity on a science-client, verifying it is connecting to science_prod

```
Log on to a science-client with your personal user-id
tcad; select DISPLAY & TMdata; a few telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

7.3.32 Test tcad connectivity on a moc-client, verifying it is connecting to science_prod

```
Log on to a moc-client with your personal user-id
tcad; select DISPLAY & TMdata; pick telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

7.3.33 Test connectivity for the science application, L2. (The app dynamically reads the interfaces file)

```
Log on to a science-client (sybase user-id is fine)
matlab; ScienceDataLook('SIM5-part3'); click Separate Figures
A plot should appear in separate window (SQ_SciLPasFilt1)
```

7.3.34 Connectivity Signoff. Get signoffs by the dba/sa/engineer that tested each app for connectivity

Application	Tested on	Tested by dba/sysadmin/eng	Date/Time	Initial
Mission Planning	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.3.35 Optionally, if performing maintenance of Level 2, MSS, OD and/or GPB databases:

- Make the databases you will do maintenance on unavailable to the end users

  For databases owned by sa:

  ```
  exec sp_dboption <dbname>, "dbo use", true
  use <dbname>
  checkpoint
  ```

  For databases owned by the actual user:

  ```
  use <dbname>
  sp_changedbowner sa
  ```

- Perform the maintenance as described in section 9.1, "Maintenance of Level2, MSS and OD Databases" below. Start maintenance with MSS and OD databases, then proceed to Level 2 database. As each database is done, make available to users by undoing the dbo-use only switch or reverting the ownership to the original owner

7.3.36 Restart Data processing cron. Login as tdp on moc-server, issue cron start, cron status

7.3.37 Application Signoff. Contact the leads for each application and ask them to verify that the databases are up-to-date and fully functional. Ask them to sign off in the table below as an acknowledgement and attach their check-lists, if used

Application	Verified on	Verified by team lead/designee	Date/Time	Initial
Mission Planning	Moc-client			
Orbit Determ	Moc-client			
Data Processing	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.3.38 Notify users that the science Sybase server is now available, running off the science server. Send email to all@relgyro.stanford.edu and/or announce over the voice-loop as appropriate. If access to a database is delayed due to maintenance, indicate this in the announcements and give an estimated time of availability for each database

_____

7.3.39 Start full backups of all databases.

```
/home/sybase/sybdba/script/full_backups.csh science_prod
```

_____

7.3.40 Reschedule the science crontab jobs that were previously put on hold. For example

```
/home/sybase/sybdba/script/full_backup.csh science_prod
/home/sybase/sybdba/script/tranlog_backup.csh science_prod
/home/sybase/sybdba/script/save_master_info.csh science_prod
/home/sybase/sybdba/script/space_stats.csh science_prod
/home/sybase/sybdba/script/cmp_disk_init.csh science_prod science_pc
/home/sybase/sybdba/script/cmp_create_db.csh science_prod science_pc
```

_____

A day or two after the successful switch back to the science_prod server:

7.3.41 Reschedule the preload jobs to the standby server that were previously put on hold

```
/home/sybase/sybdba/script/preload_standby.csh
/home/sybase/sybdba/script/preload_transaction.csh
```

_____

7.4 Switching Back to Production Server After an Emergency Switchover

This scenario is used when switching back to the original production server after an unscheduled emergency switchover. It is typically used when the original Level 0 and 1 databases on the original production server were damaged in some way. Depending on the circumstances, for instance if the Level 0 and 1 databases on the original production server are intact, scenario 7.3 above, "Switching Back to the Production Server After a Scheduled Switchover", may be more appropriate.

The procedure is pretty much an exact copy of section 7.1 above, "Scheduled Switchover to Standby Server", with the major difference being the direction of the actions, science_pc databases are being loaded onto science_offline.

Well before the scheduled switchover (one to several days):

7.4.1 Determine how long it will take to backup the current production (standby) server, load the full backups to the off-lined production server, run dbcc's afterwards and then load additional transaction logs to sync up the off-lined production server with the current production server____

7.4.2 Start full backups of the current production databases

```
/home/sybase/sybdba/script/full_backup.csh science_pc  OR
/home/sybase/local_tape_backup.csh science_pc <YYMMDD>
```

_____

7.4.3 Schedule the switchback with the flight director _____

7.4.4 Confirm the timing with the data processing team, their schedule may be affected. (No need to reprocess data against science_offline, Level 0 and 1 are copied from science_pc) _____

7.4.5 Notify QA, SA and Data IPT Lead of the coming switchback, indicating estimated start time____

7.4.6 Notify users of the planned one-hour outage that will occur while operation is switched back to the original production system. Send the email to all@relgyro.stanford.edu _____

7.4.7 Preload the databases to the off-lined production server. The switchback script automatically excludes Level 0 and 1 databases when "all" databases are loaded, so you need to list all the databases explicitely

```
chmod u+x /home/sybase/sybdba/script/switchback_offline_db.csh
/home/sybase/sybdba/script/switchback_offline_db.csh \
        GPB_L0.GPB_L1.GPB_L1A.L2.orbit_determ.gpb_3_4_3.GPB_DP
chmod a-x /home/sybase/sybdba/script/switchback_offline_db.csh >>
```

_____

Once the desired production databases have been preloaded to the standby server:

7.4.8 Run and check the dbccs on the off-lined production server. Correct problems found: first on the off-lined server, then on the production server, redo backup/switchback of affected databases___

7.4.9 If time permits, start full backups of the newly loaded databases on science_offline. The easiest is probably to do a tape backup

```
/home/sybase/sybdba/script/local_tape_backup.csh science_offline <YYMMDD>
```

_____

7.4.10 Prep the backup directories for storing future log dumps. That is, remove the oldest generation of dumps and move everything else down one generation and making sure gen_0 is empty

```
cd /sybackup1/carin/sybbackups/science_prod
/bin/rm -rf <db>/<oldest_gen>
mv <db>/<next_to_oldest_gen> <db>/<oldest_gen> til gen_0 is now gen_1
mkdir gen_0
cd /sybackup2/carin/sybbackups/science_prod
/bin/rm -rf <db>/<oldest_gen>
mv <db>/<next_to_oldest_gen> <db>/<oldest_gen> til gen_0 is now gen_1
```

```
mkdir gen_0
cd /sybackup/carin/syblogs/science_prod
/bin/rm -rf <db>/<oldest_gen>
mv <db>/<next_to_oldest_gen> <db>/<oldest_gen> til gen_0 is now gen_1
mkdir gen_0
```

7.4.11 Add a crontab job that preloads the tranlogs to science_offline. Use no time-lag and run frequently

```
chmod u+x /home/sybase/sybdba/script/switchback_offline_log.csh
in crontab:
15 * * * * /home/sybdba/script/switchback_offline_log.csh all 0
15 1,3,5,etc * * * /home/sybdba/script/switchback_offline_log.csh \
      GPB_L0.GPB_L1.GPB_L1A.L2.orbit_determ.gpb_3_4_3.GPB_DP 0
```

Four – five hours before the switchover:

7.4.12 Stop data processing cron. Login as tdp on moc-server, issue cron stop, cron status

Once the last data file has finished processing or two hours before the switchover at the latest:

7.4.13 Broadcast a warning over the voice-loop that the science Sybase server will be coming down in a couple of hours for a switchback

7.4.14 Verify with the data processing team that all data processing has indeed completed.

7.4.15 Remove the currently active switchback_offline_log job from the crontab schedule.

7.4.16 Run tranlog dumps against the standby server

```
/home/sybase/sybdba/script/tranlog_backup.csh science_pc all
```

7.4.17 Load the tranlogs to the off-lined production server

```
/home/sybdba/script/switchback_offline_log.csh
   GPB_L0.GPB_L1.GPB_L1A.L2.orbit_determ.gpb_3_4_3.GPB_DP
```

7.4.18 Put all backup jobs that run against the standby sybase server on hold by commenting out in crontab on sci-crunch (full_backup, tranlog_backup, etc.)

7.4.19 Review all other dba jobs that run against the standby and off-lined sybase servers and put on hold as appropriate by commenting out in crontab on science_pc and science respectively

7.4.20 Keep running tranlog dumps against science_pc and preloads on science_offline if time allows.

```
/home/sybase/sybdba/script/tranlog_backup.csh science_pc all
/home/sybdba/script/switchback_offline_log.csh
      GPB_L0.GPB_L1.GPB_L1A.L2.orbit_determ.gpb_3_4_3.GPB_DP
```

ACTUAL SWITCHOVER:

7.4.21 Prevent users from logging in to the standby sybase server by locking all logins except “sa”, “carin” and “probe”:

```
isql -U<user> -Sscience_pc
select 'exec sp_locklogin ' + name +', '"lock"'
      from master..syslogins where name not in ("sa", "carin", "probe")
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 != 2
```

The only logins returned should be carin, sa and probe

7.4.22 Kill any active Sybase processes in at least Level 2, MSS and OD databases.

```
/home/sybase/sybdba/script/kill_spids science_pc L2
/home/sybase/sybdba/script/kill_spids science_pc gpb_3_4_2 ##or current
/home/sybase/sybdba/script/kill_spids science_pc orbit_determ
```

7.4.23 Manually run the last transaction log dump on the standby server and preload on off-lined server

```
/home/sybase/sybdba/script/tranlog_backup.csh science_pc all
/home/sybdba/script/switchback_offline_log.csh
   GPB_L0.GPB_L1.GPB_L1A.L2.orbit_determ.gpb_3_4_3.GPB_DP
chmod a-x /home/sybdba/script/switchback_offline_log.csh
```

7.4.24 Rename the science_offline backup server so it can be restarted as science_prod_back later

```
isql -Usa -Sscience_offline
shutdown SYB_BACKUP [with nowait]   -- shut down first to be safe
sp_helpserver
sp_dropserver SYB_BACKUP
sp_addserver   SYB_BACKUP, null, science_prod_back
sp_helpserver
```

7.4.25 Shut down both the production and standby Sybase servers to force all connections to be severed.

```
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_stop
ssh -l sybase sci-crunch
/apps/licensed/sybase_local/etc/sybase_stop
```

7.4.26 Switch interfaces files on science to redirect users back to science

```
ssh -l sybase science
cd $SYBASE  ##/apps/licensed/sybase-12.5
cp interfaces.production interfaces
```

7.4.27 Start the original production Sybase server. It will now be known as science_prod and science

```
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_start
```

7.4.28 Verify you'll really be connected to science_prod when using science_prod or science

```
isql -U<user> -Sscience_prod
select @@servername   -- this should once again show science_prod
isql -U<user> -Sscience
select @@servername   -- this should once again show science_prod
```

7.4.29 Bring all user databases on science online in order to remove the standby option. Substitute the current MSS databases for database gpb_3_4_3 below:

```
isql -U<user> -Sscience_prod
select @@servername           -- this should show science
online database GPB_L0
online database GPB_L1
online database GPB_L1A
online database L2
online database gpb_3_4_2     -- change this to current MSS database
online database orbit_determ
online database GPB_DP
```

7.4.30 Unlock all the previously locked users on the production server.

```
isql -U<user> -Sscience_prod
select 'exec sp_locklogin ' + name +', '"unlock"'
      from master..syslogins where status & 2 = 2
```

execute the generated sql commands

```
select name from master..syslogins where status & 2 = 2
```

No logins should be returned. (unless there are logins locked for other reasons)

7.4.31 Switch interfaces files on the moc-server to point users back to science - Moc-authorization required

```
log on to the moc server as user database
cd /apps/licensed/sybase-12.5.0.3
cp interfaces.production interfaces
cd /apps/licensed/sybase-11.9.2
cp interfaces.production interfaces
```

7.4.32 Test connectivity for mission planning against science_prod (gpb_n_n_n, orbit_determ)

```
Log on to a moc-client with your personal user-id
MP; userid cmdops; pick MSS db; Param Gen; Application; GPS Ephemeris; Report
A valid report should be displayed
```

_____

7.4.33 Test tcad connectivity on a science-client, verifying it is connecting to science_prod

```
Log on to a science-client with your personal user-id
tcad; select DISPLAY & TMdata; a few telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

_____

7.4.34 Test tcad connectivity on a moc-client, verifying it is connecting to science_prod

```
Log on to a moc-client with your personal user-id
tcad; select DISPLAY & TMdata; pick telemetry items; deselect 1K/2K; New Plot
Graphs should be displayed
```

_____

7.4.35 Test connectivity for the science application, L2. (The app dynamically reads the interfaces file)

```
Log on to a science-client (sybase user-id is fine)
matlab; ScienceDataLook('SIM5-part3'); click Separate Figures
A plot should appear in separate window (SQ_SciLPasFilt1)
```

_____

7.4.36 Connectivity Signoff. Get signoffs by the dba/sa/engineer that tested each app for connectivity

Application	Tested on	Tested by dba/sysadmin/eng	Date/Time	Initial
Mission Planning	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.4.37 Restart Data processing cron. Login as tdp on moc-server, issue cron start, cron status _____

7.4.38 Application Signoff. Contact the leads for each application and ask them to verify that the databases are up-to-date and fully functional. Ask them to sign off in the table below as an acknowledgement and attach their check-lists, if used

Application	Verified on	Verified by team lead/designee	Date/Time	Initial
Mission Planning	Moc-client			
Orbit Determ	Moc-client			
Data Processing	Moc-client			
TCAD	Science-client			
TCAD	Moc-client			
Science	Science-client			

7.4.39 Notify users that the science Sybase server is now available, running off the science server. Send email to all@relgyro.stanford.edu and/or announce over the voice-loop as appropriate. _____

7.4.40 Reschedule the science crontab jobs that were previously put on hold. For example

```
/home/sybase/sybdba/script/full_backup.csh science_prod
/home/sybase/sybdba/script/tranlog_backup.csh science_prod
/home/sybase/sybdba/script/save_master_info.csh science_prod
/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/space_stats.csh science_prod
```

_____

A day or two after the successful switch back to the science_prod server:

7.4.41 Reschedule the preload jobs that were previously put on hold in crontab on sci-crunch

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

_____

8 Accessing science_prod while science_pc is production server

While science_pc, running on sci-crunch, is the production dataserver of record, the original science_prod dataserver, running on sci-base, must be made available for testing, repairs and maintenance. This is done by starting the original science_prod server under the science_offline alias. Internally the server is still known as science_prod, just like the server now running as science_prod is internally known as science_pc.

8.1 Starting the data and backup servers on sci-base

8.1.1 Start the original science_prod dataserver under the aliases science_offline. This is where Level 0 and 1 maintenance will take place.

```
ssh -l sybase sci-base
/apps/licensed/sybase_local/etc/sybase_start
```

_____

8.1.2 Unless already done, fix the SYB_BACKUP setting so the backup server can be started as science_offline_back

```
isql -U<user> -Sscience_offline
sp_helpserver
sp_dropserver  SYB_BACKUP
sp_addserver   SYB_BACKUP, null, science_offline_back
sp_helpserver
```

_____

8.1.3 Start the original science_prod backupserver as science_offline_back

```
ssh -l sybase science
cd $SYBASE/ASE-12_5/install
./startserver -f RUN_science_offline_back
```

_____

8.1.4 Verify you can connect to the off-lined science_prod server

```
isql -U<user> -P<password> -Sscience_offline
select @@servername         -- should display science_prod!!
```

_____

8.1.5 Verify you can still connect to the production science_prod server

```
isql -U<user> -P<password> -Sscience_prod
select @@servername         -- should display science_pc!!
```

_____

8.2 Accessing the Sybase data and backup servers

When the standby dataserver is the production dataserver of record

isql –Usa –Sscience –P<science_pc_password>
select @@servername → returns science_pc

isql –Usa –Sscience_prod –P<science_pc_password>
select @@servername → returns science_pc

isql –Usa –Sscience_pc –P<science_pc_password>
select @@servername → returns science_pc

isql –Usa –Pscience_offline –P<science_prod_password>
select @@servername → returns science_prod

9 Performing Maintenance

The approach to doing maintenance or effecting database repairs differs between Level 0 and 1 databases on the one hand and Level 2, MSS and OD databases on the other.

Since Level 2, MSS and OD databases are directly updated by the users, database maintenance and repairs must be made while the database in question is taken offline and unavailable to the end users. It is not possible to maintain a parallel set of databases with simultaneous updates.

Level 0 and 1 databases are updated only by the data processing team, through the processing of data files as they arrive from different sources. This can easily be recreated by the operational staff. These databases are very large, requiring unacceptably long outages if maintenance was done while the databases were unavailable to the end users. Maintenance on these databases is thus done on the original databases, on the offline science server, while users are accessing the same databases on the standby server. Before the switchover is reversed, the files that were processed against the standby server are reprocessed against the off-lined science server. Any other changes made on the temporary production server are lost. They are not automatically replicated to the offline science server.

9.1 Maintenance of Level2, MSS and OD Databases

Maintenance on these databases must be made in real time, while the databases are taken offline and inaccessible to the end users. There are two different approaches to doing this. What alternative is used depends on the tasks to be performed but the author's preference is to use the first.

1. Perform maintenance during the switch to the standby server
   Before allowing the users access in step 7.1.41 above, perform all the necessary maintenance tasks. When it is time to switch back to the original science server, simply copy the databases from the standby server using backup and recovery tools.
2. Perform maintenance during the switch back to production, after actual move of the databases
   Perform maintenance after all the databases have been copied back to the off-lined production server and before the users are allowed back on in step 7.3.38 above.

The time required to do maintenance on the Level 2, MSS and OD databases depends of course on what specific activities will take place and the size of the database and tables. A typical maintenance scheme would be to compress tables with excessive amounts of gaps to both improve speed and space utilization. Typically you would only reorganize the larger tables in the Level2 and possibly OD databases.

The fastest way to compress a clustered table, if space is available, is to drop and recreate the clustered index. If there is not enough db space to recreate the index, bcp the table out, truncate the table, drop the indexes, run a fast bcp in and recreate the clustered index with the sorted_data option. Add any non-clustered indexes. The disadvantage is that it requires full backups to be performed after the reload. If compressing smaller tables, just do a slow bcp in.

Performing maintenance on Level 2, MSS and OD databases:

The following is an example of how to do maintenance and can be used as a **guideline**. Each specific situation may warrant other and completely different steps to be taken.

9.1.1 While verifying that the switchover or switchback was done properly, all databases were taken online and thus 'inadvertently' made available to the end users. Make the Level 2, MSS and OD databases that you will be doing maintenance on inaccessible again to avoid eager end users logging in prematurely.

For databases owned by sa:

```
exec sp_dboption <dbname>, "dbo use", true
use <dbname>
checkpoint
```

For databases owned by the actual user:

```
use <dbname>
sp_changedbowner sa
```

_____

For each Level 2, MSS and OD database to perform maintain on:

9.1.2 If there is enough space to drop & recreate the clustered index (i.e. 110% of the compressed table size is unallocated in the database (See section 9.3.1 below for details):

9.1.2.1 Add a new, temporary segment onto the segments where you'll want to place the compressed table. This temporary segment typically overlaps the original segment. This is all necessary to actually move the data when 'with sorted data' index option is used

_____

9.1.2.2 Drop nonclustered and clustered index

_____

9.1.2.3 Recreate clustered index on the temporary segment, using "with sorted data" _____

9.1.2.4 Recreate nonclustered index _____

9.1.2.5 Point the table back to the segment where you want future table allocations.

9.1.2.6 Drop the temporary segment _____

9.1.2.7 Verify the new histogram for the table in optdiag looks ok .

_____

9.1.2.8 Drop the temporary segment _____

9.1.3 Compress larger, clustered tables using fast bcp in the following fashion

9.1.3.1 Start the bcp out. For every table to compress

```
bcp <db>..<table> out <filename> -n -U<user> -S<srv>
```
_____

9.1.3.2 Turn on fast bcp & trunc log on checkpoint

```
exec sp_dboption <dbname>, 'select into', true
exec sp_dboption <dbname>, 'trunc log', true
use <dbname>
checkpoint
```
_____

9.1.3.3 For every table, truncate table and drop index

```
truncate table <table>
drop index <table>.<clustered_index>
sp_helpindex <table>    -- verify all indexes dropped
```
_____

9.1.3.4 Bcp all the tables in. Be sure to use the blocksize –b option. The –A denotes packet size. For every table to compress

```
bcp <db>..<table> in <filename> -n -U<user> -S<srv> -b100000 -A8192
```
_____

9.1.3.5 Recreate all the clustered indexes with the sorted_data option to avoid moving the data again

```
create unique clustered index …… with ignore_dup_key, sorted_data
```
_____

9.1.3.6 If there were additional indexes, create them _____

9.1.3.7 When all the fast bcp's have been run, reset database options

```
exec sp_dboption <dbname>, 'select into', false
exec sp_dboption <dbname>, 'trurnc log', false
use <dbname>
checkpoint
```
_____

9.1.4 Compressing smaller tables using slow bcp

```
bcp <dbname>..<tablename> out <filename> -n -U<user> -S<srv>
bcp <dbname>..<tablename> in <filename> -n -b 100000 -U<user> -S<srv>
```
_____

9.1.5 Make database available to end users

For databases owned by sa:

```
exec sp_dboption <dbname>, "dbo use", false
use <dbname>
checkpoint
```
For databases owned by the actual user:

```
use <dbname>
```
sp_changedbowner <original owner> _____

9.1.6 Initiate a full backup of the database in question only if fast bcp's were run.

```
/home/sybase/sybdba/script/full_backup.csh <server> <database>
```
_____

9.1.7 Notify users that the database is now available. (Try to cut down the number of emails sent when many databases are worked on.) _____

9.1.8 Update statistics for all databases

```
/home/sybase/sybdba/script/update_stats.csh <server> <database>
```
_____

9.2 Maintenance of Level 0 and 1 Databases on the Offline Server – Database space is at a premium

Performing maintenance on Level 0 and 1 databases when database space is at a premium:

Use this scenario when there is not enough database space to drop and recreate a clustered index.

The fastest way to compress a table when there is not enough db space to recreate the clustered index is to bcp out the table in the order of the clustered index key, drop the clustered index, run a fast bcp in, while keeping the rows in the same order, and finally creating the clustered index using the with sorted_data option. It is not possible to bcp out more than 2 billion rows from a table. This problem is solved by temporarily creating a few views on the table based on the leading key in the clustered index.

The following is an example of how to do maintenance and can be used as a guideline. Each specific situation may warrant other and completely different steps to be taken.

9.2.1 To facilitate and speed up the maintenance for the Level 0 and 1 databases transaction logging should be stopped and select into and fast bcp's allowed.

```
isql –U<user> –Sscience_offline
use master
go
exec sp_dboption GPB_L0, “trunc log”, true
exec sp_dboption GPB_L0, “select into”, true
exec sp_dboption GPB_L1, “trunc log”, true
exec sp_dboption GPB_L1, “select into”, true
exec sp_dboption GPB_L1A, “trunc log”, true
exec sp_dboption GPB_L1A, “select into”, true
go
use GPB_L0
go
checkpoint
go
use GPB_L1
go
checkpoint
go
use GPB_L1A
go
checkpoint
go
sp_helpdb
```

9.2.2 Compress TManalog in the following fashion

9.2.2.1 Create views on the table, making sure no one view contains more than 2 billion rows. For example

```
create view T1 as select * from TManalog where TMID <= 5000
create view T2 as select * from TManalog where TMID > 5000 and <=8000
create view T3 as select * from TManalog where TMID > 8000
```

9.2.2.2 Bcp out the views. As a guideline 200 millions can be bcp'ed out in one hour. Run two or three bcp out's in parallel, the multiple bcp's will not impede each other noticeably.

```
bcp GPB_L1A…T1 out /carin/maint/T1.n –n –Usa –Sscience_offline etc
```

9.2.2.3 Drop the views, indexes and truncate the table

```
drop view T1 etc
drop index TManalog.TManalog_Pkey
```

```
truncate table TManalog
```

9.2.2.4 Run fast bcp's in. Be extremely careful to bcp in the files in the order of clustered key. You may only run one bcp at a time. Don't forget the batch size parameter. About 225 million rows can be fast bcp'ed in per hour.

```
sp_helpdb GPB_L1A     -- verify trunc log, select into turned on
sp_helpindex TManalog -- verify no indexes
```

bcp GPB_L1A..TManalog in /carin/maint/T1.n -n –b100000 –A8192 –Usa –P –Sscience_offline

bcp GPB_L1A..TManalog in /carin/maint/T2.n -n –b100000 –A8192 –Usa –P –Sscience_offline

bcp GPB_L1A..TManalog in /carin/maint/T3.n -n –b100000 –A8192 etc

9.2.2.5 Recreate all the clustered indexes. Use the sorted_data option to avoid moving the data. Estimate one hour for every 650 million rows.

```
create unique clustered index TManalog_Pkey  ……
    with ignore_dup_key, sorted_data on seg_analog
```

9.2.2.6 You shouldn't have to update the statistics for the table, check histogram in optdiag

```
Update statistics TManalog    -- if necessary
```

9.2.3 Compress other large tables by running a fast bcp in and reindexing with sorted data. Follow steps 9.2.1 above but skip the step to create a view. Create nonclustered indexes as needed.

9.2.4 Compress smaller tables by simply running a bcp out, truncate the table and bcp in.

<u>Reprocessing data files on the off-lined science server:</u>

9.2.5 It is critical that the data processing team keeps track of the command files that are processed against science_pc and that they are kept so they can be rerun against science_offline once all maintenance and database repairs have been done

9.2.6 To reprocess data, simply execute the tdp program with a –S<server> parameter. If no –S parameter is given, tdp will access the server instance known as GPB_DATA/GPB in the interfaces file. For example

```
tdp                   -- uses GPB_DATA/GPB server in interfaces file,
                      normally science_prod, science_pc when operating in
                      standby mode
tdp -Sscience_prod    -- uses science in interfaces file
tdp -Sscience_offline -- uses science_offline in interfaces file, science
```

9.3 Maintenance of Level 0 and 1 Databases on the Offline Server – Database space is available

Performing maintenance on Level 0 and 1 databases when there is lots of database space:

Use the following instructions when it is possible to allocate enough space to the database to drop and recreate the clustered index. The instructions below shows how to calculate the space needed. The following is an example of how to do maintenance and should be used as a **guideline**. Each specific situation may warrant other and completely different steps to be taken.

9.3.1 Compress TManalog in the following fashion

9.3.1.1 Calculate empty database space needed
Run optdiag binary statistics to get the "Space utilization" number
use binary option to allow for loading the histogram if it's incorrect after creating the index
Execute sp_spaceused TManalog to get KB used for data and index
Free space needed = (KB_data + KB_index) * Space utilization * 110%
Round up to be **extra** sure you don't run out of space _____

9.3.1.2 Expand the database to allow for free space as calculated in step 9.3.1.1 above _____

9.3.1.3 Adjust the segments for the newly allocated database devices

```
Exec sp_extendsegment seg_analog, GPB_L1A, <new_device>
Exec sp_dropsegment   "default", GPB_L1A, <new_device>
Exec sp_dropsegment   "system", GPB_L1A, <new_device>
```
_____

9.3.1.4 Add a new segment to the devices where you want the reindex process to write the new, compressed table. I.e, these are the disk devices that have free space. Use the new devices plus any old devices with a large chunk of space

```
Exec sp_addsegment seg_new, GPB_L1A, <devices with free space>
Exec sp_extendsegment seg_new, GPB_L1A, <devices with free space>
```
_____

9.3.1.5 Drop the index.

```
drop index TManalog.TManalog_PKey
```
_____

9.3.1.6 Recreate the index.
"with sorted_data" skips the sorting step (which requiresa lot of time and space)
"on seg_new" forces the data rows to be rewritten as they are moved to the new segment

```
Create unique clustered index TManalog_PKey
  on TManalog  (TMID, SCT_Cycle, SCT_VTCW)
  with ignore_dup_key, sorted_data
  on 'seg_new'
```
_____

9.3.1.7 Steer the table back to the original segment.

```
Exec sp_placeobject seg_analog, TManalog
Exec sp_dropsegment "seg_new", GPB_L1A, <all the devices>
```
_____

9.3.2 Compress other large tables by using the same steps above _____

Backfilling a single table with data that's been processed against the production server

If, for instance, only the Level 1A database is being maintained it might be easier to generate the missing data from level 0, reprocess that file and then bcp in the data file generated. The Data Processing Team can do this using the l02l1 and TDP programs

9.3.3 Backfilling data to Level 1 by generating from Level 0

9.3.3.1 Login as tdp

```
cd /apps/supported/lasp/tdp
./l02l1
choose "include 1k and 2k data" checkbox
choose cycle 12, enter time range and hit ok
< processing takes about 15 min /day generated>
cd /apps/supported/lasp/data/
look for the packet100.tmp file, this is the file you'll be processing
```

TDP

At the IDL prompt, type
.run gpb_tdp_L1
process_ptm, 'packet100.tmp', /CHECKONLY , /SCIONLY

< generating the bcp file takes 1 hr/day generated> _____

9.3.3.2 Bcp in the bcp file
bcp GPB_L1A..TManalog in <dir>/tmanalog.tmp –n –Usa –S<srv> -A8192 –b1000000 _____

10 Glossary

ARB Anomaly Review Board

bcp Bulk copy – A Sybase utility to quickly transfer large amounts of data between a database table or view and an operating system file.

DBA Database Administrator

dbcc Database Consistency Check.- A set of Sybase commands that checks the referential integrity of a database or table.

FD Flight Director

MD Mission Director

MOM Moc Manager

MSS Mission Support Software

OD Orbit Determination

11 Certification

I certify that the database switchover and switchback was performed in whole and that the data recorded above is complete and accurate.

Database Administrator _____ Date _____

Database switchover and switchback has been completed as certified by the signatures above

Moc Manager/
Flight Director _____ Date _____

This is to certify that the information obtained under this database switchover and switchback procedure is as represented and the documentation is complete and correct.

Quality Assurance _____ Date _____

Mark the signed off hardcopy with the words “As run <date> and file in binder located in dba’s office
Attach check-lists from application owners