

From: [Horng-Liang Shih](#)
To: [Eva Dafonte Perez](#); [Maria Girone](#); [Pdb Service](#); [wlcg-scod \(WLCG Service Coordinator on duty\)](#); [wlcg-operations \(wlcg operations\)](#)
Cc: [ASGC - DB](#)
Subject: Process error report STRMADMIN_APPLY_ATLN@ASGC3D.GRID.SINICA.EDU.TW
Date: Tuesday, October 06, 2009 14:27:55

Dear Eva and Maria,

sorry for being this delay summarizing the severe incident observed at ASGC 3D database. and will convert into required SIR format later on.

Description

* Sept 27, Eva help escalating the ORA-12514 error code, and propagation cannot be re-started normally. verbose exception msg is 'TNS:listener does not currently know of service requested in connect descriptor', and ticket open earlier is:

Streams Monitor Error Report

Report date: 2009-09-27 02:21:49

Affected Site: CERN-PROD

Affected Database: ATLDSC.CERN.CH

Process Name: STRM_PROPAGATION_ASGC

Error Time: 27-09-2009 02:21:48

Error Message: ORA-12514: TNS:listener does not currently know of service requested in connect descriptor

ORA-06512: at "SYS.DBMS_AQADM_SYS", line 1087

ORA-06512: at "SYS.DBMS_AQADM_SYS", line 7642

ORA-06512: at "SYS.DBMS_AQADM", line 631

ORA-06512: at line 1

Impact

* central team confirm the cease of propagation and also fail to connect to ASGC database after db restored (point-in-time recovery):

Time line of the incident (all time stamp convert in UTC):

* 28-Sep-09 07:17 - CERN-PROD : Process error report STRM_PROPAGATION_ASGC@ATLDSC.CERN.CH, service degradation of listener is suspicious.

* 28-Sep-09 12:51 - ASGC confirm the lsnr service are up and running normally at two instances. waiting further profile from central team.

* 28-Sep-09 13:01 - local db admin confirm the situation remain and stream queue is empty as well.

* 28-Sep-09 13:35 - system clock screw observed in 2nd instance (irc02), fixed already. should have limited impact for the 3D propagation issues.

* 28-Sep-09 13:40 - Eva confirm the same error code "ORA-12528: TNS: listener: all appropriate instances are blocking new connections"

* 28-Sep-09 13:50 - we found instance one is down, with error code 'ERROR at line 1: ORA-01507: database not mounted', have fixed it shortly.

* 28-Sep-09 14:05 - all Isnr service are blocked, and confirm the db is not open properly except with the fs mounted. we're profiling further and consider switch the ctrl file. possible root cause: the problem seems affect by the recent power surge and cause hardware degradation of two blade chassis, where all the Oracle RAC cluster nodes located. Consider performing point-in-time restore for the 3D database, wait confirmation from central team. extend the profile in parallel in the same time.

* 29-Sep-09 05:24 - reply to central team: "the problem should have been resolved now, with assistant from expert on call. he will double check rest of details with you later on, while the propagation is disable now and we confirm LCRs activities already via stream monitor page"

* 29-Sep-09 08:01 - Eva indicate the restarted of the propagation to ASGC. it's then clarify shortly that the action was wrong took by operator due to the end of intervention and lack of proper sign for the site status after confirming the degradation situation.

* 29-Sep-09 14:45 - Eva outline the problem persist and suspect that the database is still corrupted, even after the point-in-time recovery. In fact, it is not possible to access to some old data so this makes impossible to try any kind of resynchronization using Streams with the current database status.

* 30-Sep-09 07:04 - Eva review the event and confirm "Streams propagation was working perfectly after Wednesday 23rd. We got an alarm and the person on shift restarted the propagation job successfully. Then, the propagation was aborted again after your scheduled intervention during the weekend"

* 30-Sep-09 - perform point-in-time restore to Sept 21, where we confirm to have full backup.

* 30-Sep-09 13:41 - Eva confirm with some errors when try to access the data and this is why we suspect there is data corrupted (12 entries being selected with owner like 'ATLAS_COOLOFL_DCS' and object_name like '%_F0028_%', and the missing object observed:

```
strmadmin@ASGC3D> select * from ATLAS_COOLOFL_DCS.COMP200_F0028_IOVS_SEQ;  
*
```

```
ERROR at line 1:  
ORA-08103: object no longer exists  
*
```

* 5-Oct-09 13:41 - confirm with Eva that the datafiles are no longer valid (the point-in-time restored refer to Sept 21 while streams might not be able to recover from a 2 weeks backlog or not having archived log files on disk), and wish to have transportable tablespaces and seeking best time slot if other T1 have scheduled intervention.

Analysis

* back trace of the db log files (alert, trace, and lmon), we confirm

the db have been shutdown successfully 30min before the scheduled intervention at Sept 26 23:30 (SD announced at 24:00). the root cause remain unclear why the DB service couldn't startup normally.

* the final (2nd time) point-in-time restore performed last Fri refer to full backup at Sept 21th. and seems the full backup from Sept 21st is not longer valid in order to recover streams or might the streams not be able to recover from a two weeks backlogs.

Follow up

* to avoid the same situation, as proposed internally (the better one), (the other option is to "simulate" backups with RMAN using the option: VALIDATE), is performing real restores in tier server from scratch, just with backup files. In this case, if you are able to restore db, then we are confirm that the backup is 100% full usable.

* consider asking help from other T1 to extend two hours performing table space replication after the maintenance window.

* Eva also recommend: "the only possible solution is to perform a complete re-instantiation using transportable tablespaces. You need to agree with another Tier1 site to help you transferring the files."

* hereafter the extract of event log regarding the restore of database perform by DB admin earlier:

OLDER CONTROLFILE BACKUP

```
BS Key Type LV Size Device Type Elapsed Time Completion Time
-----
7888 Full 16.73M DISK 00:00:01 21-SEP-09
BP Key: 8266 Status: AVAILABLE Compressed: NO Tag:
TAG20090921T153015
Piece Name:
+ATLAS_DATA/asgc3d/autobackup/2009_09_21/s_698167815.1494.698167829
Control File Included: Ckp SCN: 6077891093186 Ckp time: 21-SEP-09
SPFILE Included: Modification time: 01-SEP-09
```

```
RMAN> restore controlfile from
'+ATLAS_DATA/asgc3d/autobackup/2009_09_21/s_698167815.1494.698167829';
```

```
Starting restore at 29-SEP-09
using target database control file instead of recovery catalog
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=2182 instance=asgc3d1 devtype=DISK
```

```
channel ORA_DISK_1: restoring control file
channel ORA_DISK_1: restore complete, elapsed time: 00:00:08
output filename=+ATLAS_DATA/asgc3d/controlfile/current.1512.698871999
output filename=+ATLAS_DATA/asgc3d/controlfile/current.1423.698871999
Finished restore at 29-SEP-09
```

```
RMAN> alter database mount;
```

```
RMAN-00571:
```

```
=====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
```

```
RMAN-00571:
```

```
=====
RMAN-03002: failure of alter db command at 09/29/2009 19:06:56
ORA-00203: using the wrong control files
ORA-00202: control file:
```

'+ATLAS_DATA/asgc3d/controlfile/current.1512.698871999'

in the init.ora file I remove the first one and start db with the second:
output filename=+ATLAS_DATA/asgc3d/controlfile/current.1423.698871999

Database mounted in Shared Mode (CLUSTER_DATABASE=TRUE)
Completed: ALTER DATABASE MOUNT

Then start restore with this command:

```
RMAN> run {  
SET UNTIL SCN 6077727611290;  
#SET UNTIL TIME "TO_DATE('Sep 21 2009 18:00:00','Mon DD YYYY HH24:MI:SS)";  
restore database;  
recover database;  
}
```

RMAN-03002: failure of restore command at 09/29/2009 19:21:45
RMAN-20207: UNTIL TIME or RECOVERY WINDOW is before RESETLOGS time

RMAN> list incarnation;

List of Database Incarnations

| DB Key | Inc Key | DB Name | DB ID | STATUS | Reset SCN | Reset Time |
|--------|---------|---------|------------|---------|---------------|------------|
| 1 | 1 | ASGC3D | 1409391450 | PARENT | 1 | 27-JAN-09 |
| 2 | 2 | ASGC3D | 1409391450 | PARENT | 6048726847478 | 24-MAR-09 |
| 3 | 3 | ASGC3D | 1409391450 | CURRENT | 6077727611291 | 28-SEP-09 |

RMAN> reset database to incarnation 2;

database reset to incarnation 2

To control the log, the time was: Tue Sep 29 19:22:48 2009

channel ORA_DISK_1: restored backup piece 1
piece handle=+ATLAS_BACKUP/asgc3d/datafiles/uakpn18k_1_1
tag=TAG20090920T090003
channel ORA_DISK_1: restore complete, elapsed time: 00:00:46

Tue Sep 29 19:57:47 2009
alter database recover logfile
'+ATLAS_DATA/asgc3d/archivelog/2009_09_28/thread_2_seq_4571.879.698789657'
Tue Sep 29 19:57:47 2009
Media Recovery Log
+ATLAS_DATA/asgc3d/archivelog/2009_09_28/thread_2_seq_4571.879.698789657
Tue Sep 29 19:57:48 2009
Errors in file /u01/app/oracle/admin/asgc3d/bdump/asgc3d1_p000_11562.trc:
ORA-00600: internal error code, arguments: [4552], [2], [0], [], [], [],
[], []
Tue Sep 29 19:57:49 2009
Errors in file /u01/app/oracle/admin/asgc3d/bdump/asgc3d1_p000_11562.trc:
ORA-00308: cannot open archived log
'+ATLAS_DATA/asgc3d/onlinelog/group_15.291.677251137'

```

ORA-17503: ksfopn:2 Failed to open file
+ATLAS_DATA/asgc3d/onlinelog/group_15.291.677251137
ORA-15012: ASM file
'+ATLAS_DATA/asgc3d/onlinelog/group_15.291.677251137' does not exist
ORA-00600: internal error code, arguments: [4552], [2], [0], [], [], [],
[], []
Tue Sep 29 19:57:49 2009
Errors in file /u01/app/oracle/admin/asgc3d/bdump/asgc3d1_p000_11562.trc:
ORA-10562: Error occurred while applying redo to data block (file# 10,
block# 7622242)
ORA-10564: tablespace ATLAS_COOLOFL_DCS_DATA01
ORA-01110: data file 10:
'+ATLAS_DATA/asgc3d/datafile/atlas_coolofl_dcs_data0001.dbf'
ORA-10560: block type '0'
ORA-00600: internal error code, arguments: [4552], [2], [0], [], [], [],
[], []
Errors with log
+ATLAS_DATA/asgc3d/archivelog/2009_09_28/thread_2_seq_4571.879.698789657
Recovery interrupted!
Tue Sep 29 19:58:02 2009
Recovered data files to a consistent state at change 6077727611290

```

I restart the process with this parameter:

```

run {
SET UNTIL SCN 6077727611290;
restore database;
recover database;
}

```

```

archive log
filename=+ATLAS_DATA/asgc3d/archivelog/2009_09_28/thread_1_seq_4600.753.698789657
thread=1 sequence=4600
media recovery complete, elapsed time: 00:00:21
Finished recover at 29-SEP-09

```

```

SQL> alter database open resetlogs;
Database altered.

```

```

Completed: alter database open resetlogs
Tue Sep 29 20:45:01 2009
ALTER SYSTEM SET
service_names='SYS$STRMADMIN.STRM_QUEUE_ATL_APN.ASGC3D.GRID.SINICA.EDU.TW'
SCOPE=MEMORY SID='asgc3d1';

```

```

Service "SYS$STRMADMIN.STRM_QUEUE_ATL_APN.ASGC3D.GRID.SINICA.EDU.TW" has
1 instance(s).
Instance "asgc3d1", status READY, has 2 handler(s) for this service...

```

```

[oracle@irc02 ~]$ crs_stat -t
Name          Type          Target    State    Host
-----
ora....d1.inst application  ONLINE   ONLINE   irc01
ora....d2.inst application  ONLINE   ONLINE   irc02
ora.asgc3d.db application  ONLINE   ONLINE   irc02
ora....SM1.asm application  ONLINE   ONLINE   irc01
ora....01.lsnr application  ONLINE   ONLINE   irc01

```

| | | | | |
|----------------|-------------|--------|--------|-------|
| ora.irc01.gsd | application | ONLINE | ONLINE | irc01 |
| ora.irc01.ons | application | ONLINE | ONLINE | irc01 |
| ora.irc01.vip | application | ONLINE | ONLINE | irc01 |
| ora....SM2.asm | application | ONLINE | ONLINE | irc02 |
| ora....02.lsnr | application | ONLINE | ONLINE | irc02 |
| ora.irc02.gsd | application | ONLINE | ONLINE | irc02 |
| ora.irc02.ons | application | ONLINE | ONLINE | irc02 |
| ora.irc02.vip | application | ONLINE | ONLINE | irc02 |

BR,
J

--
Jason Shih
ASGC/OPS
Tel: +886-2-2789-8374
Fax: +886-2-2783-5434