


DB2's Got Talent - 2016

Nadir Doctor


Consultant – Database and Infrastructure Expertise

- Focus on power of sql to generate a couple of reports and avoid complex parsing via DB2/Unix command-line tools like db2pd, awk, sed
- Minimize anxiety, headaches, heartburn by using a handy tool for retaining smart metadata to aid in performance comparison before and after database upgrade




Sql to generate report for database log archival duration

```
2016-02-13-02.19.55.808963+000 E7363718A377      LEVEL: Info
PID      : 11010050          TID   : 76590          PROC  : db2sysc 0
INSTANCE: db2<sid>          NODE  : 000
EDUID   : 76590             EDUNAME: db2logmgr <SID> 0
FUNCTION: DB2 UDB, data protection services, sqlpgArchiveLogFile, probe:3108
DATA #1 : <preformatted>
Started archive for log file S0117278.LOG.
...
...
...
2016-02-13-02.21.24.131298+000 E7364096A441      LEVEL: Info
PID      : 11010050          TID   : 76590          PROC  : db2sysc 0
INSTANCE: db2<sid>          NODE  : 000
EDUID   : 76590             EDUNAME: db2logmgr <SID> 0
FUNCTION: DB2 UDB, data protection services, sqlpgArchiveLogFile, probe:3180
DATA #1 : <preformatted>
Completed archive for log file S0117278.LOG to VENDOR chain 3 from
/db2/<SID>/log_dir/NODE0000/NODE0000/.
```



Sql to generate report for database log archival duration

```
select timestamp,  
       substr(fullrec, 335, 53) as msg  
from table (pd_get_diag_hist(null, null, null, null, null)) as db2_diag_info  
where eduname like 'db2logmgr%'  
       and function = 'sqlpgArchiveLogFile'  
order by timestamp desc  
fetch first 100 rows only  
with ur;  
  
--Good for DB2 9.7 and higher; easily extends to DPF [large partitioned database]  
--For DB2 10.5, specify substr(fullrec, 384, 53) due to 2 added clauses for db and  
  hostname  
--Note that sysibmadm.db_history.start_time can't be used as that is the instant  
  a database log was allocated and not when archive activity commenced  
--fullrec like '%Started archive for log file%' or  
--fullrec like '%Completed archive for log file%'  
--locate('Started archive for log file', fullrec, octets)  
--locate('Completed archive for log file', fullrec, octets)
```




Sql to generate report for database log archival duration

TIMESTAMP	MSG
2016-02-13-03.14.51.024833	Completed archive for log file S0117317.LOG to VENDOR
2016-02-13-03.14.02.058275	Started archive for log file S0117317.LOG.
2016-02-13-03.14.00.882601	Completed archive for log file S0117316.LOG to VENDOR
2016-02-13-03.11.36.380949	Started archive for log file S0117316.LOG.
2016-02-13-03.11.35.193550	Completed archive for log file S0117315.LOG to VENDOR
2016-02-13-03.10.08.196507	Started archive for log file S0117315.LOG.
2016-02-13-03.10.06.629295	Completed archive for log file S0117314.LOG to VENDOR
2016-02-13-03.08.23.342321	Started archive for log file S0117314.LOG.
2016-02-13-03.08.22.115648	Completed archive for log file S0117313.LOG to VENDOR
2016-02-13-03.07.21.492110	Started archive for log file S0117313.LOG.
2016-02-13-03.07.20.281398	Completed archive for log file S0117312.LOG to VENDOR
2016-02-13-03.05.37.658580	Started archive for log file S0117312.LOG.
2016-02-13-03.05.35.853138	Completed archive for log file S0117311.LOG to VENDOR
2016-02-13-03.02.44.733400	Started archive for log file S0117311.LOG.

...

...

...




Sql to generate database backup status report including retries [in case of initial failure and success thereafter] for a DPF db – large partitioned database

- Obtain a bird's eye view of a cumulative summary report for multiple partitions in an immense enterprise data warehouse environment
- Utilize `sysproc.db_partitions()` table function to get server information for resident partitions via sql
- Avoid complex parsing of db2pd output or with Unix tools like awk or sed
- Good for DB2 9.7 and higher


sysibmadm.db_history view

- Note odd datatype usage – `start_time` is a string!
- `objecttype='D'` for full database; `operation='B'` for backup
- `operationtype='N'` for online and `'O'` for incremental online
- Each of `sqlcode` and `sqlcaid` are null for successful partition backup
- if `sqlcaid` clause is omitted in where section of sql statement, then we will get status of backup retries as well
- Code snippet on next slide is a helpful starter to generate informative reports
- also use cte's



Sql to generate database backup status report including retries [in case of initial failure and success thereafter] for a DPF db – large partitioned database

```
select dbpartitionnum,  
       substr(dbpart.host_name, 1, position('.', dbpart.host_name, octets) - 1)  
       as host_name,  
       sqlcode,  
       sqlerrmc,  
       sqlerrp  
  
-- distinct seqnum for each session: if 8 sessions, seqnum will range from 1-8  
from sysibmadm.db_history  
     inner join table (sysproc.db_partitions()) as dbpart  
           on dbpartitionnum = dbpart.partition_number  
where start_time > '20130516%'  
     and operation = 'B'  
     and operationtype = 'N'  
     and sqlcaid = 'SQLCA'  
group by dbpartitionnum, host_name, sqlcode, sqlerrmc, sqlerrp  
order by dbpartitionnum with ur;
```



Handy tool to retain performance comparison metrics pre & post database upgrade

`$DB2DIR/bin`

`db2cos`

`db2cos_clp`

`db2cos_datacorruption`

`db2cos_hadr`

`db2cos_hang`

`db2cos_lib`

`db2cos_perf`


`db2cos_preupgrade`

`db2cos_threshold`

`db2cos_trap`

Available in DB2 9.7 and higher [validated presence in DB2 9.7 fixpack 7 release]

New addition in DB2 10.5 fixpack 4 and higher - `db2cos_member`



Handy tool to retain performance comparison metrics pre & post database upgrade

```
#!/bin/sh
```

```
#####
```

```
# This is a data collection script for database upgrade related issues. It collects  
# explain plan,db2pd output,db2 snapshots and OS performance data. In some  
# cases it might be needed to modify this script to collect extra information such  
# as (change vmstat or iostat command specific to OS). To do so place a copy of  
# this file in sqllib/adm and modify it there. The script in sqllib/bin will be  
# executed only when the customized script is not found in sqllib/adm.
```

```
# Script Name: db2cos_preupgrade Argument: DB name is a required option
```

```
...
```

```
...
```


```
...
```

```
echo "*****WARNING*****"
```

```
echo "* This tool should be run with caution. It can cause significant performance  
degradation, especially on busy system with high number of active connections *"
```

```
echo "* Run times indicated below are estimates based on test runs. Your actual  
run times might vary, dependent on factors such as your hardware specifications,  
operating system configuration and current workload *"
```

```
echo "*****"
```

Handy tool to retain performance comparison metrics pre & post database upgrade


Recommendations:

- Execute the db2cos_preupgrade script **before** database upgrade at a couple of off-peak and peak periods [weekday, weekend, special month-end processing]
- Execute the db2cos_preupgrade script **after** database upgrade at a couple of **similar** off-peak and peak periods [weekday, weekend, special month-end processing]

This will make it easier to compare any access path changes in case significant deterioration is noted for any sql post database upgrade

Summary:

- Lay emphasis on using DB2's rich function library and sql instead of complex parsing via DB2/Unix command-line tools like db2pd, awk, sed
- Review all DB2-supplied scripts for maximizing user productivity and system availability



A Big Thank You

Audience and Judges

Sponsors – IDUG, DBI, Triton, KBCE, Responsive Systems

Questions: Nadir Doctor <nadird3 @ gmail . com>



IDUG DB2 North America Tech Conference
Austin, Texas | May 2016



Olympic.org

Official website of the Olympic Movement

Search

Home

Olympic Games

YOG

Sports

Athletes

Countries

Olympism

Photos

Videos

Home > Olympic Games > Rio de Janeiro 2016



RIO DE JANEIRO 2016

5th August - 21st August

