



Consistent State

Phone: 303-217-0961

E-Mail: kevink@consistentstate.com

Web: <http://www.consistentstate.com/>

Report Summary

Date: 2011-04-22

Stanford

ID: 1

Table of Contents

Summary	3
Cluster: hmidb0	4
Subscribed Services	4
Service: Query	5
We will be using the pgfouine reports for query tuning	6
Service: Base	7
BGWriter Stops Spike	8
Awaiting approval of Explain Analyze schedule for query tuning	8
Average traffic per day	8
Tuning the buffers written by background processes	9
Settings to tune the BGWriter	9
aia_test.lev1p5 has a very high number of dead rows	9
Tables are continuing to be read more often	10
Requested Checkpoints May Show Problems	10
drms_sessions tables contain high amounts of dead rows	10
tables aia.lev0 and hmi.lev1 have high disk reads	11
Awaiting approval of DB Restart Schedule	11
Cluster: hmidb2	12
Subscribed Services	12
Service: Slony	13
HMIDB2 is low on disk space on root partition	14
Cluster: dcs2: dcs0 warm standby	15
Subscribed Services	15
Service: Warm Standby	16
WAL Archives Still Infrequent	17
Cluster: dcs2: dcs1 warm standby	18
Subscribed Services	18
Service: Warm Standby	19
WAL Archives Still Infrequent	20

Summary

Status		
WARN	Cluster [hmidb0]; Service [base]: BGWriter Stops	BGWriter Stops Spike
INFO	Cluster [hmidb0]; Service [base]: Explain Analyze schedule	Awaiting approval of Explain Analyze schedule for query tuning
INFO	Cluster [hmidb0]; Service [base]: Average Traffic	Average traffic per day
WARN	Cluster [hmidb0]; Service [base]: Buffers Written by Backend Processes	Tuning the buffers written by background processes
WARN	Cluster [hmidb0]; Service [base]: BGWriter Tuning	Settings to tune the BGWriter
WARN	Cluster [hmidb0]; Service [base]: Table Dead Rows	aia_test.lev1p5 has a very high number of dead rows
INFO	Cluster [hmidb0]; Service [base]: Tables reads climbing	Tables are continuing to be read more often
INFO	Cluster [hmidb0]; Service [base]: Checkpoints	Requested Checkpoints May Show Problems
WARN	Cluster [hmidb0]; Service [base]: Live vs Dead Rows	drms_sessions tables contain high amounts of dead rows
INFO	Cluster [hmidb0]; Service [base]: Table Block Stats	tables aia.lev0 and hmi.lev1 have high disk reads
WARN	Cluster [hmidb0]; Service [base]: DB Restart for Checkpoint tuning	Awaiting approval of DB Restart Schedule
INFO	Cluster [hmidb0]; Service [query]: Query Tuning	We will be using the pgfouine reports for query tuning
WARN	Cluster [hmidb2]; Service [slony]: HMIDB2 low on disk space	HMIDB2 is low on disk space on root partition
INFO	Cluster [dcs2: dcs0 warm standby]; Service [warm_standby]: Infrequent WAL Archives	WAL Archives Still Infrequent
INFO	Cluster [dcs2: dcs1 warm standby]; Service [warm_standby]: Infrequent WAL Archives	WAL Archives Still Infrequent

Cluster: hmidb0

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7

Cluster ID: 1

Host Name or IP: 192.168.0.49

Postgres Port: 5432

Subscribed Services

Service	Description
Query	pgFouine reporting
Base	Base monitoring service
System	ssh based systats & iostats

Service: Query

pgFouine reporting

INFORMATION	We will be using the pgfouine reports for query tuning
	We will be moving on query tuning, using up to half of your admin pack hours to do so. Watch the pgfouine html reports for changes in the coming weeks and months.
References	423971c4-b256-11df-8c2d-0800274182f7.1.query.cust_cluster.report.html

Service: Base

Base monitoring service

WARNING	BGWriter Stops Spike
	Over the past month, we've seen several huge spikes in the background writer stops. This usually means that there was too much data for the background writer to write to disk when it woke up, causing it to halt after hitting the maxpages value.
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf

INFORMATION	Awaiting approval of Explain Analyze schedule for query tuning
	We propose setting up a schedule where we can have queries set up to run EXPLAIN ANALYZE on the database during non-peak hours (say midnight till 6AM), to allow us to get timings and work on improving them. Once approved, this schedule will be used for ongoing query tuning when problematic queries arise.
References	423971c4-b256-11df-8c2d-0800274182f7.1.query.cust_cluster.report.html

INFORMATION	Average traffic per day
	Your average traffic per day over the last week is about 1637.31118 MB per second
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf

WARNING	Tuning the buffers written by background processes
	After doing some research, we found that the metric, 'Buffers Written by Background Processes (NOT the BGWriter)', shows us the number of buffers that are written directly to disk because the buffer pool is not set large enough to contain the data from an insert or update statement. We are requesting that you allow us to use some of your admin pack hours to look at your server and see if we can increase your buffer pool enough to reduce the numbers on this graph significantly.
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf

WARNING	Settings to tune the BGWriter
	We are currently waiting on the monthly restart schedule (per our discussion) before moving on checkpoint tuning. Look for changes soon.
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf

WARNING	aia_test.lev1p5 has a very high number of dead rows
	Table aia_test.lev1p5 has continually had in excess of 1 million dead rows. This is quite excessive, and we highly recommend you use your admin pack hours to let us create a new vacuum strategy for this table.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.report.pdf

INFORMATION	Tables are continuing to be read more often
	Table hit ratios are decreasing, dead space is climbing, and block stats are showing significantly more disk reads than before. We highly recommend using some of your admin pack hours to have us analyze the source of this and discover some way of bringing your hit ratios back to where they used to be.
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_tab.report.pdf

INFORMATION	Requested Checkpoints May Show Problems
	We would like to know why we are occasionally seeing requested checkpoints. While requesting checkpoints may not be a bad thing, it depends on the reason - if you are requesting checkpoints because you have to for performance, stability, or other similar reasons, there may be some other problem. Otherwise, it's probably not a big deal.
References	423971c4-b256-11df-82d-0800274182f7.1.base.cust_cluster.report.pdf

WARNING	drms_session tables contain high amounts of dead rows
	In our top ten tables with a high amount of dead rows, six of them are the drms_session tables from various schema's. When looking behind these tables, they all receive a high number of UPDATES, but low INSERTs and DELETEs. Suggestion: Review vacuum strategy for the drms_session tables. We suggest using admin pack hours to allow us to 'rebuild' these tables to eliminate the dead space.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.report.pdf

INFORMATION	tables aia.lev0 and hmi.lev1 have high disk reads
	Two tables have a high number of reads from disc, they are aia.lev0 and hmi.lev1. Previous months data shows that this is unusual, as normally the data is retrieved from memory. Suggestion: We are requesting that you allow us to use some of your admin pack hours to do some buffercache stats research on these tables.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.report.pdf

WARNING	Awaiting approval of DB Restart Schedule
	We will be submitting a database restart request soon, a week ahead of the requested date
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf

Cluster: hmidb2

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7

Cluster ID: 3

Host Name or IP: 192.168.0.76

Postgres Port: 5432

Subscribed Services

Service	Description
Query	pgFouine reporting
Slony	slony monitoring
System	ssh based systats & iostats

Service: Slony

slony monitoring

WARNING

HMIDB2 is low on disk space on root partition

HMIDB2 only has 3.2GB available disk space on the root partition. The partition is at 94% full and hasn't shown signs of free space decreasing, but it could pose a problem again.

References

[423971c4-b256-11df-8c2d-0800274182f7.3.system_stats.cust_disk.report.pdf](#)

Cluster: dcs2: dcs0 warm standby

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7

Cluster ID: 6

Host Name or IP: 192.168.0.12

Postgres Port: 5430

Subscribed Services

Service	Description
Warm Standby	Warm Standby monitoring service

Service: Warm Standby

Warm Standby monitoring service

INFORMATION	WAL Archives Still Infrequent
	You may want to use some of your admin pack hours to have us look at your archive_timeout values on your warm standby instances to make sure that you won't be losing large amounts of data if a failover happens
References	423971c4-b256-11df-8c2d-0800274182f7.6.warm_standby.cust_cluster.report.pdf

Cluster: dcs2: dcs1 warm standby

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7

Cluster ID: 7

Host Name or IP: 192.168.0.12

Postgres Port: 5431

Subscribed Services

Service	Description
Warm Standby	Warm Standby monitoring service

Service: Warm Standby

Warm Standby monitoring service

INFORMATION	WAL Archives Still Infrequent
	You may want to use some of your admin pack hours to have us look at your archive_timeout values on your warm standby instances to make sure that you won't be losing large amounts of data if a failover happens
References	423971c4-b256-11df-8c2d-0800274182f7.7.warm_standby.cust_cluster.report.pdf
