Report Summary

Date: June 25, 2011, midnight

Stanford

ID: 1

Contents

Feedback Summary

Status	Title	Summary
WARN	Cluster [hmidb0]; Service [base]: Index Hit Ratios	The index hit ratios dropped dramatically
INFO	Cluster [hmidb0]; Service [base]: Average Traffic	Average traffic per day
WARN	Cluster [hmidb0]; Service [base]: Tables reads climbing	Tables are continuing to be read more often
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
WARN	Cluster [hmidb0]; Service [base]: Buffers Written by Backend Processes	Tuning the buffers written by background processes
INFO	Cluster [hmidb0]; Service [base]: Checkpoints	Requested Checkpoints May Show Problems
WARN	Cluster [hmidb0]; Service [base]: Table Dead Rows	aia_test.lev1p5 has a very high number of dead rows
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
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 [hmidb0_sums]; Service [base]: Dead vs. Live Rows	public.sum_open in jsoc_sums improved dead row count
INFO	Cluster [hmidb0_sums]; Service [base]: Dead vs. Live Rows	public.sum_partn_avail in jsoc_sums bad dead vs. live ratio
WARN	Cluster [hmidb0_sums]; Service [base]: Database Block Stats	Database jsoc_sums has high disk reads
WARN	Cluster [hmidb0_sums]; Service [base]: Table Block Stats	Tables public.sum_main and public.sum_partn_alloc are being read from disk

Cluster: hmidb0

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7
Cluster ID: 1
Operating System: None
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 system and I/O stats

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	The index hit ratios dropped dramatically
	On the 24th of May, 2011, index hit ratios dropped from an average of about 99% to a maximum of 50%. This could be due to a sudden jump in the amount of data being read, resulting in more data being pushed out of memory and then being read from disk. We request some admin pack hours to investigate this and see what may have caused this.
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_idx.report.pdf
INFORMATION	Average traffic per day
	Your average traffic per day over the last week is approximately 819.023 MB/s
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf
WARNING	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
WARNING	BGWriter Stops Spike
	We saw a number of large bgwriter stops again, this report. We will continue to monitor these graphs to determine which values we need to increase or decrease to get the stops down to a minimum.
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
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
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	aia_test.lev1p5 has a very high number of dead rows
	Table aia_test.lev1p5 has an average of 4,000 deletes per hour, causing a high amount of deadspace to grow rapidly. Autovacuum doesn't vacuum the table until there are over 10,000 dead rows, which can happen once a week or even less. We suggest you use some admin pack hours to allow us to prepare a new vacuum strategy for this table.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.report.pdf
WARNING	drms_sessions tables contain high amounts of dead rows
	In our top ten tables with a high amount of dead rows, five of them are the drms_session tables from various schemas. 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

Cluster: hmidb2

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7
Cluster ID: 3
Operating System: None
Host Name or IP: 192.168.0.76
Postgres Port: 5432

Subscribed Services

Service	Description
Query	pgFouine reporting
Slony	slony monitoring
System	SSH-based system and I/O stats

Service: Slony

slony monitoring

WARNING	HMIDB2 is low on disk space on root partition
	HMIDB2 only has 3.0GB available disk space on the root partition. The partition is at94% 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: hmidb0_sums

Collector UUID: 423971c4-b256-11df-8c2d-0800274182f7
Cluster ID: 8
Operating System: None
Host Name or IP: 192.168.0.49
Postgres Port: 5434

Subscribed Services

Service	Description
Base	Base monitoring service

Service: Base

Base monitoring service

INFORMATION	public.sum_open in jsoc_sums improved dead row count
	Table public.sum_open in the jsoc_sums database had roughly equal dead rows and live rows on our report for 2011-06-18. This week, the dead row count dropped to about half of the live row count, which is a significant improvement. Still, we will be watching this table, and if improvements don't continue, we will likely be recommending a new vacuum strategy
References	423971c4-b256-11df-8c2d-0800274182f7.8.base.cust_tab.report.pdf
INFORMATION	public.sum_partn_avail in jsoc_sums bad dead vs. live ratio
	Table public.sum_partn_avail in the jsoc_sums database has consistently had more dead rows than live rows. Since it seems to be accessed often, this could cause a problem even though it has a low number of rows. We recommend using some of your admin pack hours to have us plan a new vacuum strategy for this table.
References	423971c4-b256-11df-8c2d-0800274182f7.8.base.cust_tab.report.pdf
WARNING	Database jsoc_sums has high disk reads
	We are seeing a very high number of disk reads vs. memory hits for the jsoc_sums database. If this continues, we may want to look into tuning these statistics.
References	423971c4-b256-11df-8c2d-0800274182f7.8.base.cust_db.report.pdf
WARNING	Tables public.sum_main and public.sum_partn_alloc are being read from disk
	Tables public.sum_main and public.sum_partn_alloc have very high disk reads and very low memory hits. These two tables look to be the main culprits for the database level block stats, so if and when we tune the database block stats, we will probably be looking at these two tables.
References	423971c4-b256-11df-8c2d-0800274182f7.8.base.cust_tab.report.pdf