Consistent State

Phone: 303-217-0961

E-Mail: kevink@consistentstate.com

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

Report Summary

Date: 2011-06-04

Stanford

ID: 1

Table of Contents

Summary

Status		
WARN	Cluster [hmidb0]; Service [base]: Minimal index usage on public.sum_partn_alloc	The table public.sum_partn_alloc has little index usage.
WARN	Cluster [hmidb0]; Service [base]: Rebuild public.sum_partn_avail	The table public.sum_partn_avail has a very high amount of dead rows
WARN	Cluster [hmidb0]; Service [base]: BGWriter maxpages tuning	Still higher checkpoint pages than bgwriter pages.
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]: SLONY lag on 6/3	SLONY lags peak to over 20,000 seconds on 6/3
WARN	Cluster [hmidb2]; Service [slony]: HMIDB2 low on	HMIDB2 is low on disk space on

Status		
	disk space	root partition

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 sysstats & 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	**The table public.sum_partn_alloc has little index usage.**
	The table public.sum_partn_alloc has minimal index usage, we recommend use of admin pack hours to tune indexes for this table.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_idx.report.pdf

---

WARNING	**The table public.sum_partn_avail has a very high amount of dead rows**
	The table public.sum_partn_avail has a very high amount of dead rows, we recommend a full rebuild of the table.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.report.pdf

---

WARNING	**Still higher checkpoint pages than bgwriter pages.**
	We are still seeing more checkpoint pages being writtern than bgwriter pages, we will be increase bgwriter_lru_maxpages to 500 from 250, and set checkpoint_timeout to 40min for the same reason.
References	423971c4-b256-11df-82d-0800274182f7.1.base.cust_cluster.report.pdf

---

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 744MB/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**
	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

---

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 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

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

Host Name or IP: 192.168.0.76

Postgres Port: 5432

Subscribed Services

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

Service: Slony

slony monitoring

WARNING	**SLONY lags peak to over 20,000 seconds on 6/3**
	SLONY lag peaks to more than 20,000 seconds, recommend use of admin pack hours to investigate the cause.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_idx.re

---

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

---