Consistent State

Phone: 303-217-0961

E-Mail: kevink@consistentstate.com

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

Report Summary

Date: 2010-12-03

Stanford

ID: 1

Table of Contents

Summary

Status		
WARN	Cluster [hmidb0]; Service [base]: Cluster Checkpoints	Checkpoints are abnormal
INFO	Cluster [hmidb0]; Service [base]: Background Writer Stops	BGWriter Stops are slightly high
INFO	Cluster [hmidb0]; Service [base]: Block Stats	Very impressive block stats
WARN	Cluster [hmidb0]; Service [base]: Vacuum Activity	High analyzes and autovacuum numbers
WARN	Cluster [hmidb0]; Service [base]: Table Block Stats	hmi.lev0_seq_0011, hmi.lev0_isp_0021, and hmi.lev0_isp_0011 heavy disk reads
INFO	Cluster [hmidb0]; Service [base]: Table Scans Activity	All of the tables are using sequential scan
INFO	Cluster [hmidb0]; Service [base]: Index Scan Stats	hmi.lev0a.lev0a_pkey and aia_test.lev1p5.lev1p5_pkey could be more efficient
INFO	Cluster [hmidb2]; Service [slony]: SLONY Cluster Stats	November lag spike and updates to the cluster stats

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: Base

Base monitoring service

WARNING	**Checkpoints are abnormal**
	Checkpoints are high, and the foreground vs. background checkpoints are the opposite of what we would like to see.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.pdf

---

INFORMATION	**BGWriter Stops are slightly high**
	The numbers in this metric are slightly higher than we would like to see, so it might be worth your time to increase the number of writes allowed.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.pdf

---

INFORMATION	**Very impressive block stats**
	Your block stats are looking VERY good - most of your data is being read from memory, with very little requiring an expensive disk read. Keep it up!
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_db.pdf

---

WARNING	**High analyzes and autovacuum numbers**
	Autovacuum may be struggling to keep up with the database and the analyzes are being run very frequently. This may be normal, but it could be worth your time to choose more aggressive maintenance memory and autovacuum settings, and grouping your analyze statements could have a very large impact on performance.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.pdf

---

WARNING	**hmi.lev0_seq_0011, hmi.lev0_isp_0021, and hmi.lev0_isp_0011 heavy disk reads**
	For the tables hmi.lev0_seq_0011, hmi.lev0_isp_0021, and hmi.lev0_isp_0011, much of the data retrieved for these tables from September 2010 through November 2010 was pulled from disk, which may indicate that the size of the buffer pool, the table's indexing strategy, or the table's storage/partition strategy should be modified. However, we don't recommend changes until we can look at more data over time.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.pdf

---

INFORMATION	**All of the tables are using sequential scan**
	In all of the tables in our top ten list, most data is being retrieved via sequential scan. Each scan retrieves a lot of data, which is efficient, but performance could potentially be improved by a new index strategy.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.pdf

---

INFORMATION	**hmi.lev0a.lev0a_pkey and aia_test.lev1p5.lev1p5_pkey could be more efficient**
	We need more data, but the primary keys on these tables (hmi.lev0a and aia_test.lev1p5) may not be indexed in the most efficient way.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_idx.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

INFORMATION	**November lag spike and updates to the cluster stats**
	The lag time spiked in November to a value of nearly 500 minutes. While the problem was solved by December, we still thought it interesting to note. Also, we should mention here that our slony graphs will soon be changing. We will be breaking up lag_events and lag_time into separate graphs, and adding the set_count value to the SLONY Set Table Count graphs.
References	2010-12-03/1.423971c4-b256-11df-8c2d-0800274182f7.3.slony.cust_cluster.report.pdf