



Consistent State

Phone: 303-217-0961

E-Mail: [kevink@consistentstate.com](mailto:kevink@consistentstate.com)

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

## Report Summary

Date: 2011-03-04

Stanford

ID: 1

# Table of Contents

<a href="#">Summary</a> .....	3
<a href="#">Cluster: hmidb0</a> .....	4
<a href="#">Subscribed Services</a> .....	4
<a href="#">Service: Base</a> .....	5
<a href="#">BGWriter Stops Slightly High</a> .....	6
<a href="#">We suggest these settings to tune the BGWriter</a> .....	6
<a href="#">Average traffic per day</a> .....	6
<a href="#">Tables are continuing to be read more often</a> .....	7
<a href="#">aia_test.synoptic2 and aia_test.lev1p5 have very large numbers of dead rows</a> .....	7
<a href="#">Cluster: hmidb2</a> .....	8
<a href="#">Subscribed Services</a> .....	8
<a href="#">Service: Slony</a> .....	9
<a href="#">Slony Lag Spike</a> .....	10
<a href="#">Cluster: dcs2: dcs0 warm standby</a> .....	11
<a href="#">Subscribed Services</a> .....	11
<a href="#">Service: Warm Standby</a> .....	12
<a href="#">WAL Archives Still Infrequent</a> .....	13
<a href="#">Cluster: dcs2: dcs1 warm standby</a> .....	14
<a href="#">Subscribed Services</a> .....	14
<a href="#">Service: Warm Standby</a> .....	15
<a href="#">WAL Archives Still Infrequent</a> .....	16

## Summary

Status		
INFO	Cluster [hmidb0]; Service [base]: BGWriter Stops	BGWriter Stops Slightly High
WARN	Cluster [hmidb0]; Service [base]: BGWriter Tuning	We suggest these settings to tune the BGWriter
INFO	Cluster [hmidb0]; Service [base]: Average Traffic	Average traffic per day
INFO	Cluster [hmidb0]; Service [base]: Tables reads climbing	Tables are continuing to be read more often
CRIT	Cluster [hmidb0]; Service [base]: Table Dead Rows	aia_test.synoptic2 and aia_test.lev1p5 have very large numbers of dead rows.
INFO	Cluster [hmidb2]; Service [slony]: Slony Stats	Slony Lag Spike
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***

<b>Service</b>	<b>Description</b>
Query	pgFouine reporting
Base	Base monitoring service
System	ssh based systats & iostats

***Service: Base***

Base monitoring service

INFORMATION	<b>BGWriter Stops Slightly High</b>
	We want to watch the trends for this based on the recent checkpoint tuning for a week or so; if little or no improvement is noticed, we will likely recommend some additional BGWriter-specific changes.
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_tab.report.pdf

---

WARNING	<b>We suggest these settings to tune the BGWriter</b>
	The following settings should help tune the BGWriter to more impressive values: bgwriter_delay = 50ms; bgwriter_lru_maxpages = 250; bgwriter_lru_multiplier = 5.0
References	423971c4-b256-11df-8c2d-0800274182f7.1.base.cust_cluster.report.pdf

---

INFORMATION	<b>Average traffic per day</b>
	Your average traffic per day over the last ten days is about 1746.9 MB per second
References	423971c4-b256-11df-8c2d-0800274182f7.1.jsoc.base.cust_db.report.pdf

---

<b>INFORMATION</b>	<b>Tables are continuing to be read more often</b>
	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

---

<b>CRITICAL</b>	<b>aia_test.synoptic2 and aia_test.lev1p5 have very large numbers of dead rows.</b>
	Both aia_test.synoptic2 and aia_test.lev1p5 have 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

---

## **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***

<b>Service</b>	<b>Description</b>
Query	pgFouine reporting
Slony	slony monitoring
System	ssh based systats & iostats



***Service: Slony***

slony monitoring

<b>INFORMATION</b>	<b>Slony Lag Spike</b>
References	423971c4-b256-11df-8c2d-0800274182f7.3.slony.cust_cluster.report.pdf

---

## Slony Lag Spike

We are noticing that there seems to be a lag time spike of 100+ minutes every week on your slony slave (hmidb2). If you are unsure why this is happening, you should probably have us investigate using your admin pack hours.

References

423971c4-b256-11df-8c2d-0800274182f7.3.slony.cust\_cluster.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***

<b>Service</b>	<b>Description</b>
Warm Standby	Warm Standby monitoring service

***Service: Warm Standby***

Warm Standby monitoring service

<b>INFORMATION</b>	<b>WAL Archives Still Infrequent</b>
	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	<a href="#">423971c4-b256-11df-8c2d-0800274182f7.6.warm_standby.cust_cluster.report.pdf</a>

---

## **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***

<b>Service</b>	<b>Description</b>
Warm Standby	Warm Standby monitoring service

***Service: Warm Standby***

Warm Standby monitoring service

<b>INFORMATION</b>	<b>WAL Archives Still Infrequent</b>
	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	<a href="#">423971c4-b256-11df-8c2d-0800274182f7.7.warm_standby.cust_cluster.report.pdf</a>

---