

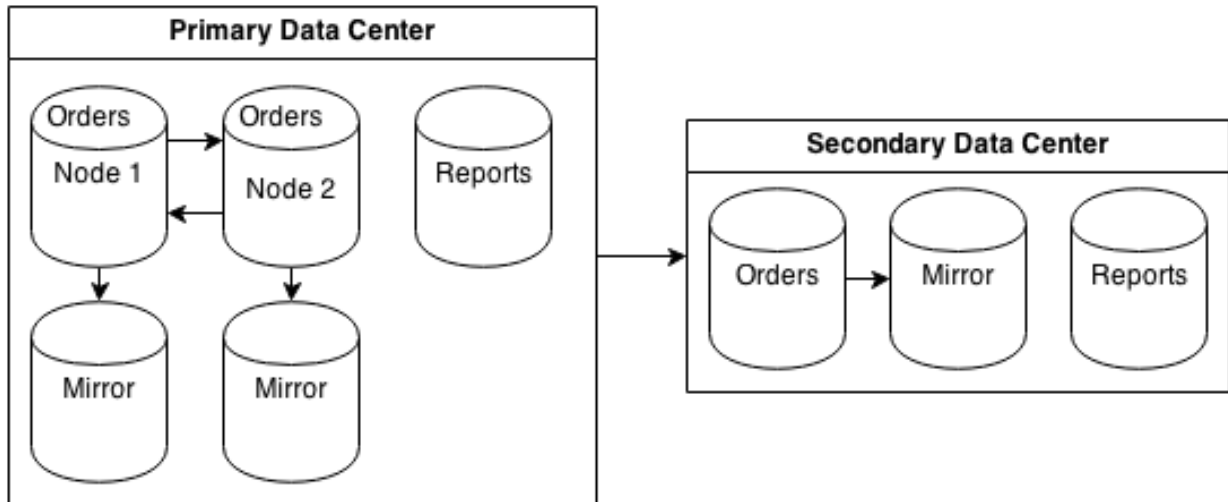


Chapter 1: Hardware planning

	A	B	C	D	E	F
1	Type	Capacity	Supplier	Price	Count	Total Cost
2	Chassis					0
3	CPU					0
4	Hard Drive (3.5")					n





	A	B	C	D	E
1	Type	Capacity	Supplier	Price	Count
2	CPU	10-core			3
3	Network Card	10GbE			3
4	RAID Controller	1GB, RAID 10			3
5	RAM	16GB			10



Chapter 2: Handling and avoiding Downtime

```
pid | client_port | state | duration | query
-----+-----+-----+-----+-----
5766 | -1 | idle in transaction | 00:03:53.894828 | BEGIN;
(1 row)
```

Chapter 3: Pooling Resources



Properties	SSL	SSH Tunnel	Advanced
Name	<input type="text" value="Primary Cluster"/>		
Host	<input type="text" value="192.168.56.30"/>		
Port	<input type="text" value="6432"/>		

Chapter 4: Troubleshooting

```

----total-cpu-usage---- -dsk/total- -net/total- ---paging-- ---system--
usr  sys  idl  wai  hiq  siq| read  writ| recv  send| in  out| int  csw
  1   0  97   1   0   0|  81k  229k|    0    0|    0    0| 100  479
 17   4   2  75   0   2| 376k 1920k|    0    0|    0    0| 440 4335
 16   2  10  70   0   2| 320k 1344k|    0    0|    0    0| 382 3371
 19   3   1  73   0   3| 496k 1956k|    0    0|    0    0| 502 5574
 15   3   3  77   0   2| 320k 2320k|    0    0|    0    0| 449 3936
 17   2   9  71   0   1| 304k 1248k|    0    0|    0    0| 361 3481
 19   3   3  73   0   2| 496k 1816k|    0    0|    0    0| 513 6388
 18   5   0  74   0   2| 376k 2112k|    0    0|    0    0| 481 4988

```

```


-net/total- ---load-avg--- ----total-cpu-usage---- --io/total- sda--sr0-
recv  send| 1m  5m 15m  usr  sys  idl  wai  hiq  siq| read  writ|util:util
  0     0|2.06 0.79 0.36  2   2  89   6   0   0|27.6  12.1|8.56:0.00
238k  201k|2.06 0.79 0.36  2  26   0  62   0   9| 156  169|95.8:  0
218k  186k|2.13 0.83 0.37  4  29   0  59   0   7| 147  190|96.2:  0
265k  219k|2.13 0.83 0.37  4  32   0  54   0  10| 134  203|95.4:  0
176k  157k|2.13 0.83 0.37  3  30   1  55   0  11| 134  249|95.4:  0
120k  117k|2.13 0.83 0.37  4  29   1  55   0  11| 119  258|94.7:  0


```

```

----system----- -----memory-usage----- ---system-- --dsk/sda-- -----most-expensive-----
time          | used  buff  cach  free| int  csw | read  writ| i/o process
15-10 18:08:11| 383M 9028k 298M 1311M| 561 3360| 412k 1307k| postgres 1812k 1534k
15-10 18:08:12| 384M 9028k 301M 1306M| 627 6266| 3536k  0 | pgbench  0  424k
15-10 18:08:13| 384M 9028k 306M 1302M| 625 5973| 3560k  32k| pgbench  0  389k
15-10 18:08:14| 384M 9028k 309M 1298M| 776 6745| 3392k  0 | pgbench  0  397k
15-10 18:08:15| 384M 9028k 313M 1294M| 599 5670| 3720k  0 | pgbench  0  384k
15-10 18:08:16| 384M 9036k 317M 1291M| 561 5596| 3192k  24k| pgbench  0  354k
15-10 18:08:17| 384M 9036k 321M 1286M| 671 6438| 4128k  0 | pgbench  0  433k

```





% time	seconds	usecs/call	calls	errors	syscall
0.00	0.000000	0	93		lseek
0.00	0.000000	0	63		sendto
0.00	0.000000	0	17		brk
0.00	0.000000	0	2	1	recvfrom
0.00	0.000000	0	1		epoll_wait
100.00	0.000000_		176	1	total

Chapter 5: Monitoring

	A	B	C	D	E	F
1	Monitor	Importance	Frequency	Warning Level	Critical Level	Action
2	Disk Space of /db	major	1 hour	1.5TB	2TB	email support
3	PostgreSQL online	critical	10 seconds	N/A	no	email DBAs
4	Server Ping	critical	10 seconds	100ms	500ms	email support
5	OS User Count	minor	1 minute	10	20	ignore
6	/db Mount	critical	10 seconds	N/A	missing	panic



Host Status Details For All Host Groups

Host	Status	Last Check	Duration	Status Information
localhost	UP	2014-02-09 16:41:34	0d 2h 3m 50s	PING OK - Packet loss = 0%, RTA = 0.05 ms

PostgreSQL Servers ([pg-servers](#))


Host	Status	Services	Actions
pg-1	UP	1 OK	

Host	Service	Status	Last Check
pg-1	Current Load	OK	2014-02-09 18:58:32
	PostgreSQL Status	OK	2014-02-09 18:59:06



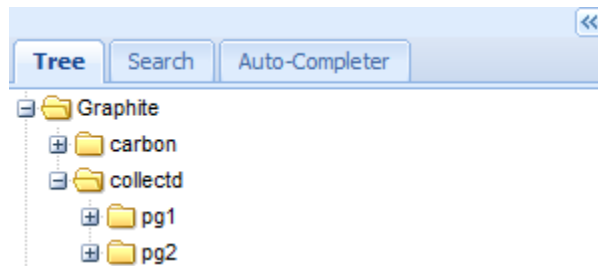
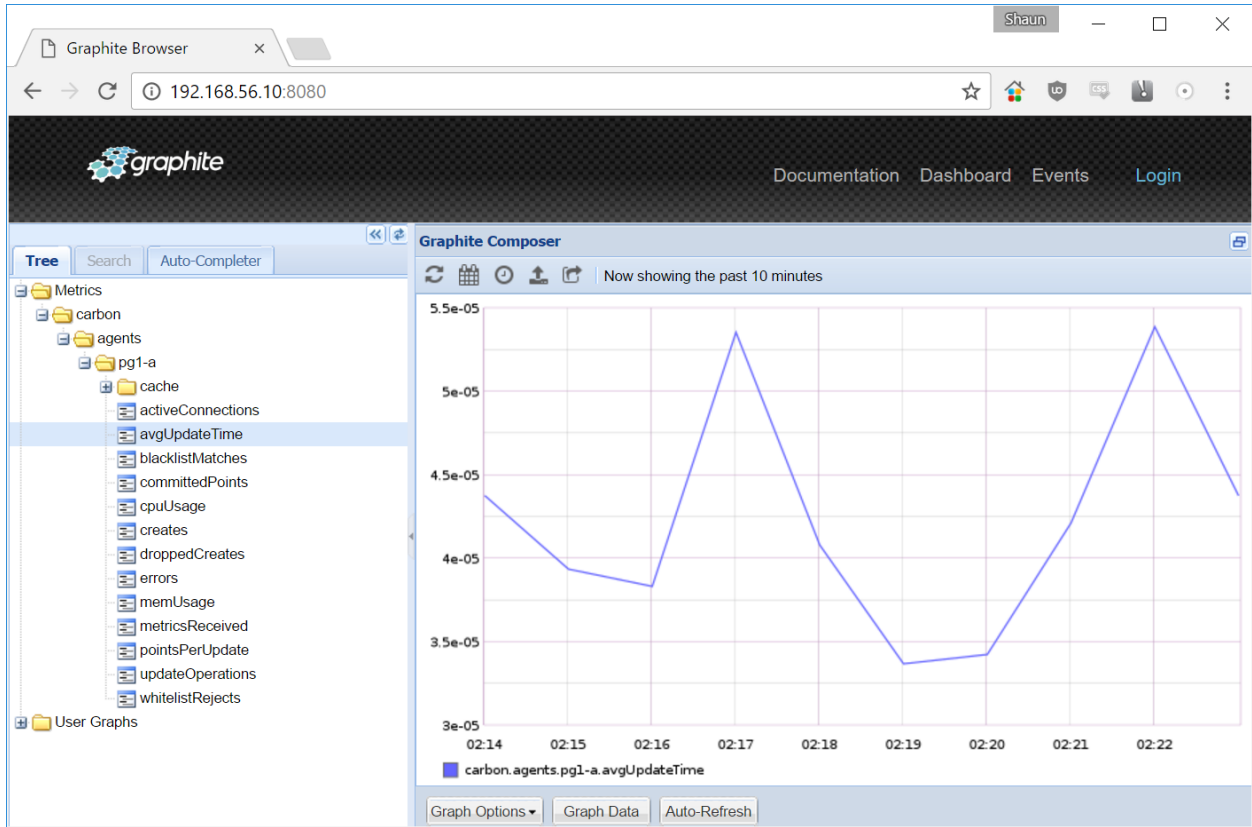
Now you should do an [inventory](#) in order to auto-configure all services

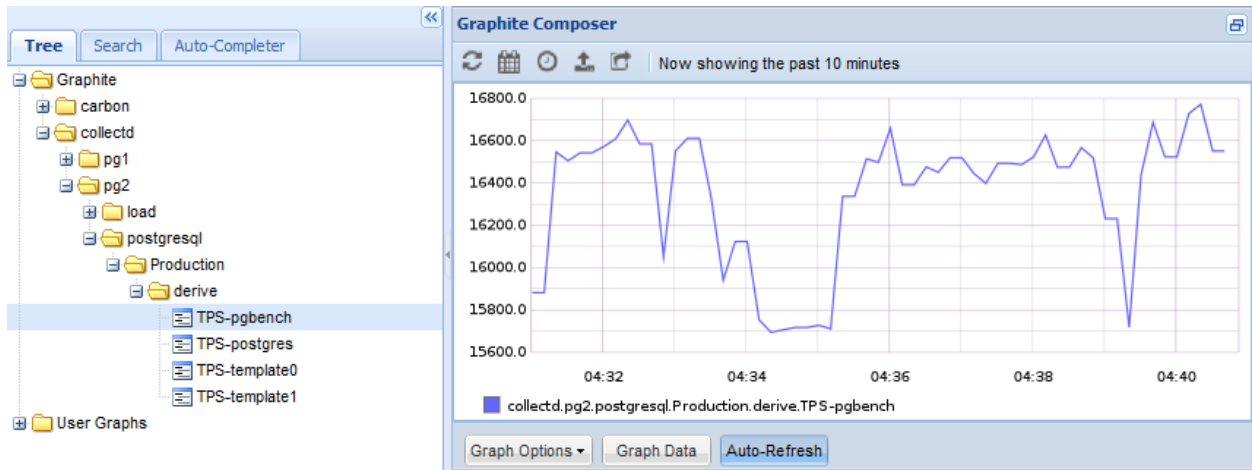
pg-1						
State	Service	Status detail	Icons	Age	Checked	Perf-O-Meter
OK	Check_MK	OK - Agent version 1.2.4, execution time 0.2 sec		3 min	17 sec	0.2s
OK	CPU load	OK - 15min load 0.13 at 4 CPUs		3 min	17 sec	0.1
OK	CPU utilization	OK - user: 2.2%, system: 0.7%, wait: 0.1%		3 min	17 sec	2%



```

drwxr-xr-x 2 root root 4096 Oct 18 18:04 ./
drwxr-xr-x 4 root root 4096 Oct 18 18:03 ../
-rw-r--r-- 1 root root 267 Oct 18 18:07 derive-TPS-pgbench-2016-10-18
-rw-r--r-- 1 root root 411 Oct 18 18:07 derive-TPS-postgres-2016-10-18
-rw-r--r-- 1 root root 369 Oct 18 18:07 derive-TPS-template0-2016-10-18
-rw-r--r-- 1 root root 369 Oct 18 18:07 derive-TPS-templatel-2016-10-18
  
```







login

username

password



My Graphs

- My Graphs
- Trading
- Production TPS
- Database Write Activity

Buttons: Graph Options, Graph Data

Graph Data [X]

collectd.pg2.postgresql.Production.derive.TPS-pgbench


Add

Edit

Remove

Apply Function ▾

Undo Function



Graphs ▾ Share Relative Time Range Absolute Time Range

New Graph ▸


Edit Default Parameters

Resize

Remove All

From URL

From Saved Graph



Dashboard Finder [X]

Trading System Overview

filter dashboard listing

Open Delete Close

Select Relative Time Range [X]

Show the past: 30

minutes [v]

Until:

now [v]

Ok Cancel

Chapter 6: Replication

	A	B	C	D	E	F
1	Source Server	Target Server	Type	DB Name	Tables	Set
2	Trading	Trading DR	Replica	All	All	N/A
3	Trading	Trading Ad Hoc	Replica	All	All	N/A
4	Trading	Reporting	Logical	maindb	customer	orders
5	Trading	Reporting	Logical	maindb	order	orders
6	Trading	Reporting	Logical	maindb	product	orders

```
$> sudo service rsync stop
* Stopping rsync daemon rsync [ OK ]
```


```
drwx----- 2 postgres postgres 4096 Oct 22 15:25 pg_archived
drwx----- 19 postgres postgres 4096 Oct 22 15:33 pgdata
```

```
cp: cannot stat `00000004000000010000007E': No such file or directory
cp: cannot stat `00000004000000010000007E': No such file or directory
cp: cannot stat `00000004000000010000007E': No such file or directory
```

```
client_addr | username | state
-----+-----+-----
192.168.56.20 | rep_user | streaming
```

```
client_addr | state | sync_state | application_name
-----+-----+-----+-----
192.168.56.20 | streaming | sync | node2
```

```
-rw----- 1 postgres postgres 16777216 Oct 22 16:19 0000000100000000000000007
-rw----- 1 postgres postgres 16777216 Oct 22 16:20 0000000100000000000000008
-rw----- 1 postgres postgres 16777216 Oct 22 16:20 0000000100000000000000009
-rw----- 1 postgres postgres 16777216 Oct 22 16:20 000000010000000000000000A.partial
-rw-rw-r-- 1 postgres postgres 242 Oct 22 16:20 wal_archive.log
```



schemaname	tablename
public	pgbench_accounts
public	pgbench_branches
public	pgbench_history
public	pgbench_tellers

```

postgres=# SELECT count(*) FROM pgbench_accounts;
count
-----
100000
  
```


```

Current connection settings:
1. Host:          192.168.56.10
2. Port:          5432
3. User:          bucardo
4. Database:      bucardo
5. PID directory: /var/run/bucardo
Enter a number to change it, P to proceed, or Q to quit:
  
```

```

          PID of Bucardo MCP: 18941
          No syncs have been created yet.
1. Table: public.pgbench_accounts  DB: pg1  PK: aid (integer)
2. Table: public.pgbench_branches  DB: pg1  PK: bid (integer)
3. Table: public.pgbench_tellers   DB: pg1  PK: tid (integer)
  
```

Name	State	Last good	Time	Last I/D	Last bad	Time
pgbench	Good	Oct 22, 2016 19:27:43	21h 58m 57s	2662/2662	none	



```

postgres@pg1:/etc/skytools$ pgrep -alf londiste
17143 /usr/bin/python /usr/local/bin/londiste3 -d primary.ini worker
17149 /usr/bin/python /usr/local/bin/londiste3 -d subscriber.ini worker

```

```

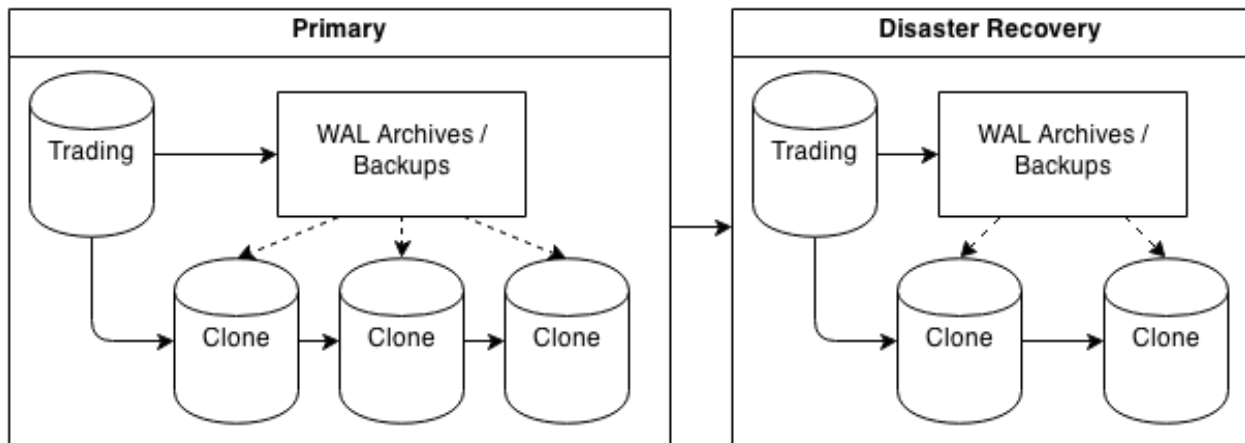
2016-10-23 11:51:51,671 30313 INFO Locking public.pgbench_accounts
2016-10-23 11:51:51,673 30313 INFO Syncing public.pgbench_accounts
2016-10-23 11:51:54,197 30313 INFO Counting public.pgbench_accounts
2016-10-23 11:51:54,377 30313 INFO srcdb: 100000 rows, checksum=39460277388
2016-10-23 11:51:54,610 30313 INFO dstdb: 100000 rows, checksum=39460277388

```

subscription_name	status	provider_node	replication_sets
pgbench	replicating	origin	{pgbench}

status	latest_end_lsn	latest_end_time	slot_name
streaming	0/F0012C0	2016-10-25 20:40:49.485761-05	pg2_slot

Chapter 7: Replication Management Tools



	A	B	C	D	E	F
1	Server Name	Source	Environment	Streaming	Promotion	Backup
2	trading		Production	TRUE	TRUE	FALSE
3	clone-1	trading	Production	TRUE	FALSE	FALSE
4	wal-archive	trading	Production	FALSE	FALSE	TRUE
5	trading-dr	trading	DR	TRUE	TRUE	TRUE
6	dr-clone-1	trading-dr	DR	TRUE	FALSE	FALSE

Server primary:

```

PostgreSQL: OK
superuser: OK
wal_level: OK
directories: OK
retention policy settings: OK
backup maximum age: OK (no last_backup_maximum_age provided)
compression settings: OK
failed backups: OK (there are 0 failed backups)
minimum redundancy requirements: OK (have 0 backups, expected at least 0)
ssh: OK (PostgreSQL server)
not in recovery: OK
archive_mode: OK
archive_command: OK
continuous archiving: OK
archiver errors: OK

```

```

barman@pg2:~$ barman list-backup primary
primary 20161026T195808 - Wed Oct 26 19:56:20 2016 - Size: 68.3 MiB - WAL Size: 0 B

```



```
Starting remote restore for server primary using backup 20161026T195808
Destination directory: /db/pgdata
Copying the base backup.
Copying required WAL segments.
Generating archive status files
Identify dangerous settings in destination directory.
```

IMPORTANT

These settings have been modified to prevent data losses

postgresql.conf line 217: archive_command = false

WARNING

You are required to review the following options as potentially dangerous

```
postgres@pg-primary:~$ sanity-check.sh
Checking:
- /usr/local/bin
- /usr/local/lib
9 programs, 31 libraries.
Tar version
All checked, and looks ok.
```

```
postgres@pg-primary:~$ tail -f /var/log/postgresql/omnipitr.log
2016-10-27 19:45:01.045352 -0500 : 8062 : omnipitr-archive : LOG : Segment
/db/pgdata/pg_xlog/0000000100000000000000013 successfully sent to all destinations.
2016-10-27 19:46:47.194793 -0500 : 8111 : omnipitr-archive : LOG : Segment
/db/pgdata/pg_xlog/0000000100000000000000014 successfully sent to all destinations.
```

```
postgres@pg-primary:~$ repmgr -f /etc/repmgr.conf master register
[2016-10-27 20:00:41] [NOTICE] master node correctly registered for cluster
pgnet with id 1 (conninfo: host=pg-primary dbname=postgres)
```

```
[2016-10-27 20:24:32] [INFO] connecting to database 'host=pg-primary dbname=postgres'
[2016-10-27 20:24:32] [INFO] connected to database, checking its state
[2016-10-27 20:24:32] [INFO] checking cluster configuration with schema 'repmgr_pgnet'
[2016-10-27 20:24:32] [INFO] checking node 1 in cluster 'pgnet'
[2016-10-27 20:24:32] [INFO] reloading configuration file and updating repmgr tables
[2016-10-27 20:24:32] [INFO] starting continuous master connection check
```

```
[2016-10-27 20:41:53] [NOTICE] destination directory '/db/pgdata' provided
[2016-10-27 20:41:53] [NOTICE] starting backup (using pg_basebackup)...
[2016-10-27 20:41:53] [HINT] this may take some time; consider using the
-c/--fast-checkpoint option
[2016-10-27 20:41:54] [NOTICE] standby clone (using pg_basebackup) complete
[2016-10-27 20:41:54] [NOTICE] you can now start your PostgreSQL server
[2016-10-27 20:41:54] [HINT] for example : pg_ctl -D /db/pgdata start
```

standby_node	standby_name	replication_lag
2	child1	0 bytes

(1 row)

```
postgres@pg-clone:~$ repmgr -f /etc/repmgr.conf cluster show
```

Role	Name	Upstream	Connection String
* master	parent		host=pg-primary dbname=postgres
standby	child1	parent	host=pg-clone dbname=postgres

```
postgres@pg-clone:~$ repmgr -f /etc/repmgr.conf cluster show
```

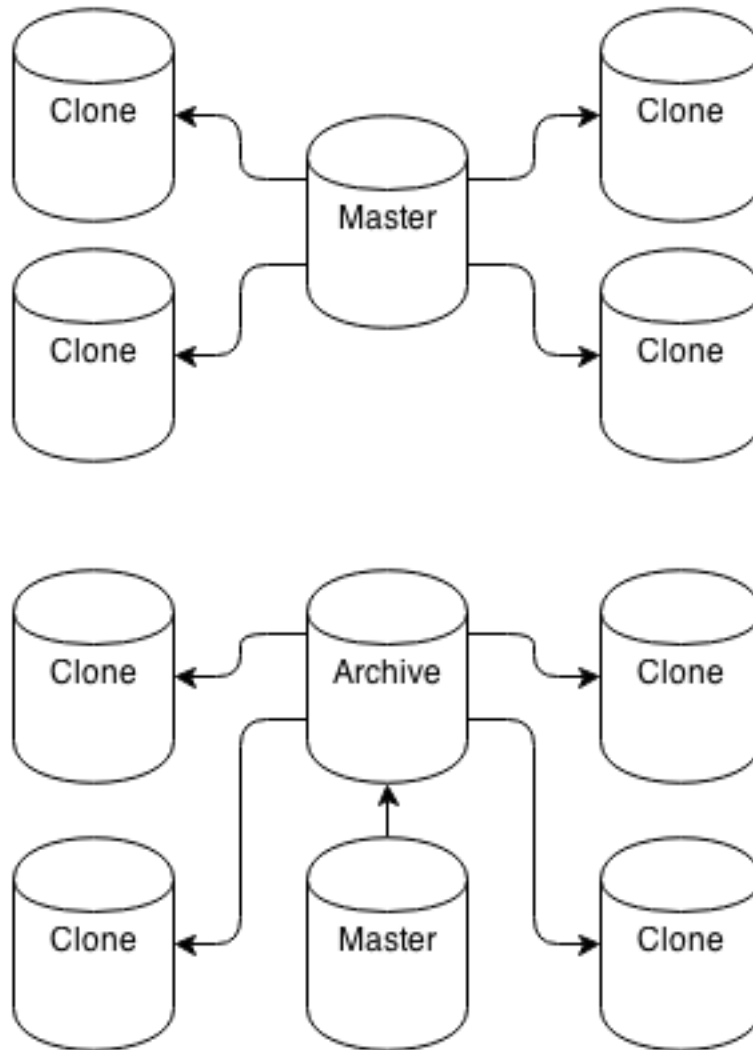
Role	Name	Upstream	Connection String
FAILED	parent		host=pg-primary dbname=postgres
* master	child1		host=pg-clone dbname=postgres

```
postgres@pg1:~$ repmgr -f /etc/repmgr.conf cluster show
```

Role	Name	Upstream	Connection String
standby	parent		host=pg-primary dbname=postgres
* master	child1		host=pg-clone dbname=postgres

```
[2016-10-28 20:34:56] [INFO] retrieving node list for cluster 'pgnet'
[2016-10-28 20:34:56] [ERROR] connection to database failed: could not connect to server:
Connection refused
Is the server running on host "pg-primary" (192.168.56.10) and accepting
TCP/IP connections on port 5432?

[2016-10-28 20:34:56] [INFO] checking cluster configuration with schema 'repmgr_pgnet'
[2016-10-28 20:34:56] [INFO] checking node 3 in cluster 'pgnet'
[2016-10-28 20:34:56] [INFO] reloading configuration file and updating repmgr tables
[2016-10-28 20:34:56] [INFO] starting continuous standby node monitoring
```



Modifying postgresql.conf for WAL management...

- * Checking wal_level: changed to replica.
- * Checking max_wal_senders: changed to 5. (Minimum value)
- * Checking archive_mode: ok. (on)
- * Checking archive_command: changed to '/usr/bin/walctl_push %p'.
- * Checking hot_standby: ok. (on)

Done modifying config.

Reloading PostgreSQL configuration files... done.

NOTICE: Some config values changed require PostgreSQL restart.
 Restart PostgreSQL with this command to enable these:
 /usr/lib/postgresql/9.6/bin/pg_ctl -D /db/pgdata restart

Create a Bucket - Select a Bucket Name and Region Cancel

A bucket is a container for objects stored in Amazon S3. When creating a bucket, you can choose a Region to optimize for latency, minimize costs, or address regulatory requirements. For more information regarding bucket naming conventions, please visit the [Amazon S3 documentation](#).

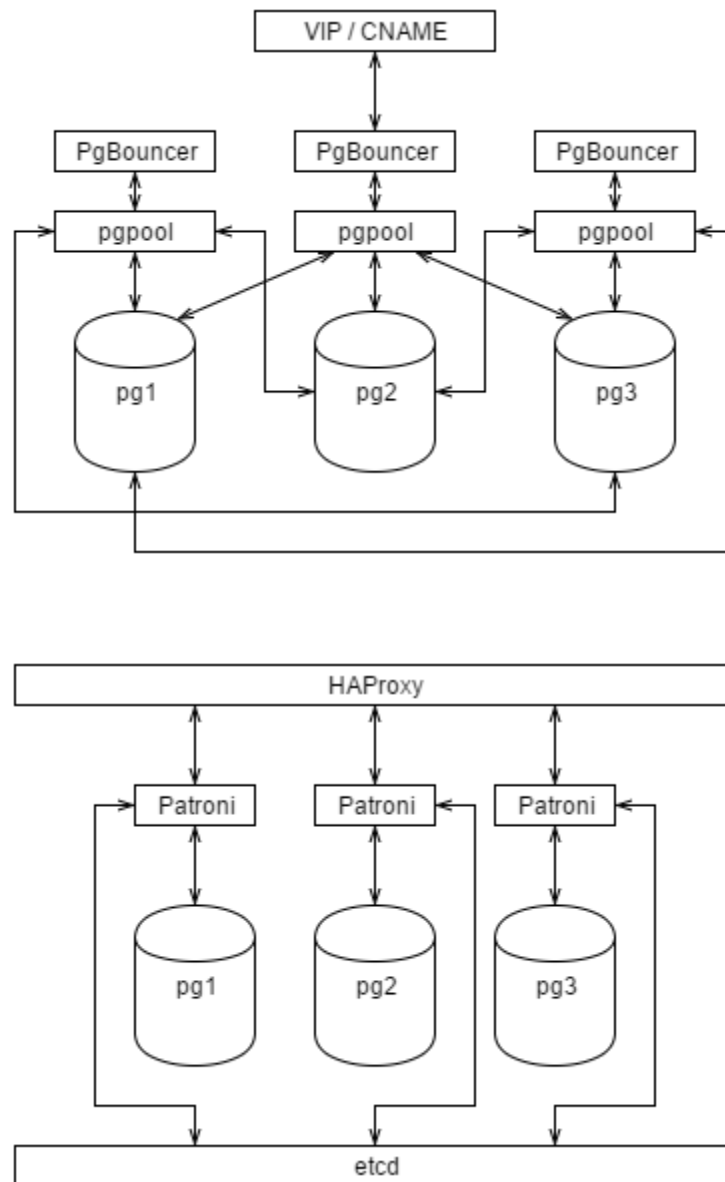
Bucket Name:

Region:

```
postgres@pg1:~$ envdir /etc/wal-e/env wal-e wal-push test.txt
wal_e.main ERROR MSG: AWS support requires module "boto"
HINT: Try running "pip install boto".
STRUCTURED: time=2016-10-29T21:20:09.243400-00 pid=7136
```

```
postgres@pg1:~$ tail -f /var/log/postgresql/postgresql-9.6-main.log | grep "DETAIL"
DETAIL: Uploading "pg_xlog/00000003000000000000002F" to
"s3://postgres-ha-cookbook/wal/wal_005/00000003000000000000002F.lzo".
DETAIL: Archiving to "s3://postgres-ha-cookbook/wal/wal_005/00000003000000000000002F.lzo"
complete at 204.568KiB/s.
```

Chapter 8: Simple Stack



```
postgres@pg1:/db/pgdata$ ETCDCCTL_API=3 etcdctl get ha-cookbook-1 ha-cookbook-9
ha-cookbook-3
Hello World!
ha-cookbook-2
Hello World!
```

```
postgres@pg3:~$ psql -h pg3 -p 5000 -c "select inet_server_addr();"
inet_server_addr
-----
192.168.56.10
```

```
2016-11-06 14:04:11,601 INFO: establishing a new patroni connection to the postgres cluster
2016-11-06 14:04:11,730 INFO: Lock owner: pg1; I am pg2
2016-11-06 14:04:11,730 INFO: does not have lock
2016-11-06 14:04:11,734 INFO: no action. i am a secondary and i am following a leader
```

```
postgres@pg1:~$ patronictl list -c /etc/patroni/stampede.yml stampede
```

Cluster	Member	Host	Leader	State	Lag in MB
stampede	pg1	pg1	*	running	0.0
stampede	pg2	pg2		running	0.0
stampede	pg3	pg3		running	0.0

```
2016-11-11 19:55:08.88969 Successfully failed over to "pg2"
```

```
postgres@pg1:~$ patronictl list -c /etc/patroni/stampede.yml stampede
```

Cluster	Member	Host	Leader	State	Lag in MB
stampede	pg1	pg1		stopped	7824.0
stampede	pg2	pg2	*	running	
stampede	pg3	pg3		running	

```
postgres@pg1:/db/pgdata$ patronictl list -d pg1:2379 stampede
```

Cluster	Member	Host	Leader	State	Lag in MB
stampede	pg1	pg1		running	0.0
stampede	pg2	pg2	*	running	0.0
stampede	pg3	pg3		running	0.0

```
2016-11-06 13:53:39,666 INFO: cleared rewind flag after becoming the leader
2016-11-06 13:53:39,788 INFO: promoted self to leader by acquiring session lock
2016-11-12 16:25:50,982 INFO: trying to bootstrap from leader 'pg1'
2016-11-12 16:25:52,484 INFO: replica has been created using basebackup
2016-11-12 16:25:52,484 INFO: bootstrapped from leader 'pg1'
2016-11-12 16:25:52,484 INFO: Starting new HTTP connection (2): 192.168.56.20
2016-11-12 16:25:55,128 INFO: Lock owner: pg1; I am pg2
2016-11-12 16:25:55,128 INFO: bootstrap from leader 'pg1' in progress
2016-11-12 16:26:03,505 INFO: establishing a new patroni connection to the postgres cluster
2016-11-12 16:26:03,524 INFO: Lock owner: pg1; I am pg2
```

```
postgres@pg1:/db/pgdata$ patronictl list -d pg1:2379 stampede
```

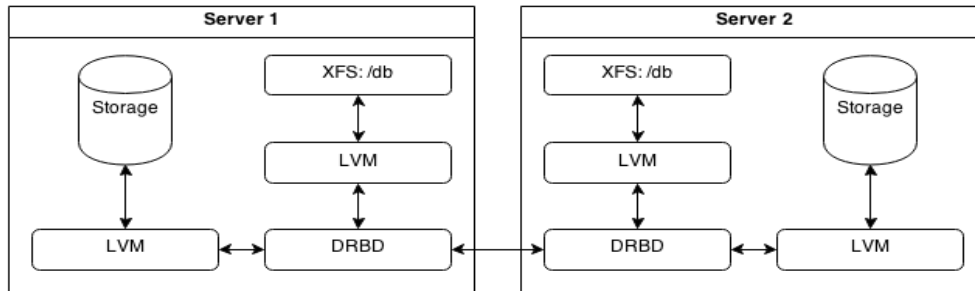

Cluster	Member	Host	Leader	State	Lag in MB
stampede	pg1	pg1	*	running	0.0
stampede	pg2	pg2		running	0.0
stampede	pg3	pg3		running	0.0
stampede	pg4	pg4		running	0.0

```
2016-11-13 14:57:11,164 INFO: Connecting to pg3:2181
2016-11-13 14:57:11,337 INFO: Zookeeper connection established, state: CONNECTED
2016-11-13 14:57:11,589 WARNING: Postgresql is not running.
2016-11-13 14:57:11,610 INFO: Lock owner: None; I am pg1
2016-11-13 14:57:11,611 INFO: starting as a secondary
2016-11-13 14:57:21,414 INFO: establishing a new patroni connection to the postgres cluster
2016-11-13 14:57:21,445 INFO: cleared rewind flag after becoming the leader
2016-11-13 14:57:21,860 INFO: promoted self to leader by acquiring session lock
```

```
2016-11-13 16:36:02,893 INFO: Starting new HTTP connection (1): pg1
2016-11-13 16:36:03,295 WARNING: Postgresql is not running.
2016-11-13 16:36:03,297 INFO: Lock owner: None; I am pg1
2016-11-13 16:36:03,297 INFO: starting as a secondary
2016-11-13 16:36:13,111 INFO: establishing a new patroni connection to the postgres cluster
2016-11-13 16:36:13,330 INFO: cleared rewind flag after becoming the leader
2016-11-13 16:36:13,945 INFO: promoted self to leader by acquiring session lock
```

Cluster	Member	Host	Leader	State	Lag in MB
stampede	pg1	pg1		running	0.0
stampede	pg2	pg2	*	running	0.0
stampede	pg3	pg3		running	0.0

Chapter 9: Advanced Stack



```
Model: ATA VBOX HARDDISK (scsi)
Disk /dev/sdb: 4295MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:
```

Number	Start	End	Size	File system	Name	Flags
1	1049kB	4294MB	4293MB		primary	lvm

```
root@pg1:~# lvdisplay VG_DRBD/LV_DATA | grep LV
LV Path                /dev/VG_DRBD/LV_DATA
LV Name                 LV_DATA
LV UUID                jG3938-TtR9-UL8D-Xybw-kD7h-wRkd-geF79g
LV Write Access        read/write
LV Creation host, time dev1, 2016-11-19 16:31:23 -0600
LV Status               available
LV Size                 4.00 GiB
```

```
initializing activity log
NOT initializing bitmap
Writing meta data...
New drbd meta data block successfully created.
```

```
root@pg1:~# cat /proc/drbd
version: 8.4.5 (api:1/proto:86-101)
srcversion: D496E56BBEBA8B1339BB34A
0: cs:SyncSource ro:Secondary/Secondary ds:UpToDate/Inconsistent C r-----
ns:193056 nr:0 dw:0 dr:193056 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:3996988
[>.....] sync'ed: 4.7% (3996988/4190044)K
finish: 0:06:21 speed: 10,488 (9,192) K/sec
```



```
root@pg1:~# vgdisplay VG_POSTGRES | grep Size
VG Size          3.99 GiB
PE Size          4.00 MiB
Alloc PE / Size  970 / 3.79 GiB
Free PE / Size   52 / 208.00 MiB
```

```
version: 8.4.5 (api:1/proto:86-101)
srcversion: D496E568BEBA8B1339BB34A
0: cs:VerifyS ro:Primary/Secondary ds:UpToDate/UpToDate C r-----
ns:4190088 nr:0 dw:44 dr:5181812 al:1 bm:0 lo:0 pe:577 ua:0 ap:0 ep:1 wo:f oos:0
[==>.....] verified: 23.6% (3204676/4190044)K
finish: 0:02:15 speed: 23,676 (13,876) want: 26,160 K/sec
```

```
version: 8.4.5 (api:1/proto:86-101)
srcversion: D496E568BEBA8B1339BB34A
0: cs:StandAlone ro:Primary/Unknown ds:UpToDate/DUnknown r-----
ns:0 nr:4190088 dw:4190088 dr:4190180 al:0 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:0
```

```
meta-data=/dev/VG_POSTGRES/LV_DATA isize=512    agcount=128, agsize=7760 blks
          =                               sectsz=512   attr=2, projid32bit=1
          =                               crc=1        finobt=1, sparse=0
data      =                               bsize=4096   blocks=993280, imaxpct=25
          =                               sunit=0      swidth=0 blks
naming    =version 2                       bsize=4096   ascii-ci=0 ftype=1
log       =internal log                    bsize=4096   blocks=2560, version=2
          =                               sectsz=512   sunit=0 blks, lazy-count=1
realtime  =none                             extsz=4096   blocks=0, rtextents=0
```

```
root@pg1:~# df /dev/mapper/VG_POSTGRES-LV_DATA
Filesystem          1K-blocks  Used Available Use% Mounted on
/dev/mapper/VG_POSTGRES-LV_DATA 3962880 37424 3925456 1% /db
```

```
root@pg1:~# xfs_db -f -c frag /dev/VG_POSTGRES/LV_DATA
actual 1224, ideal 1024, fragmentation factor 16.34%
```

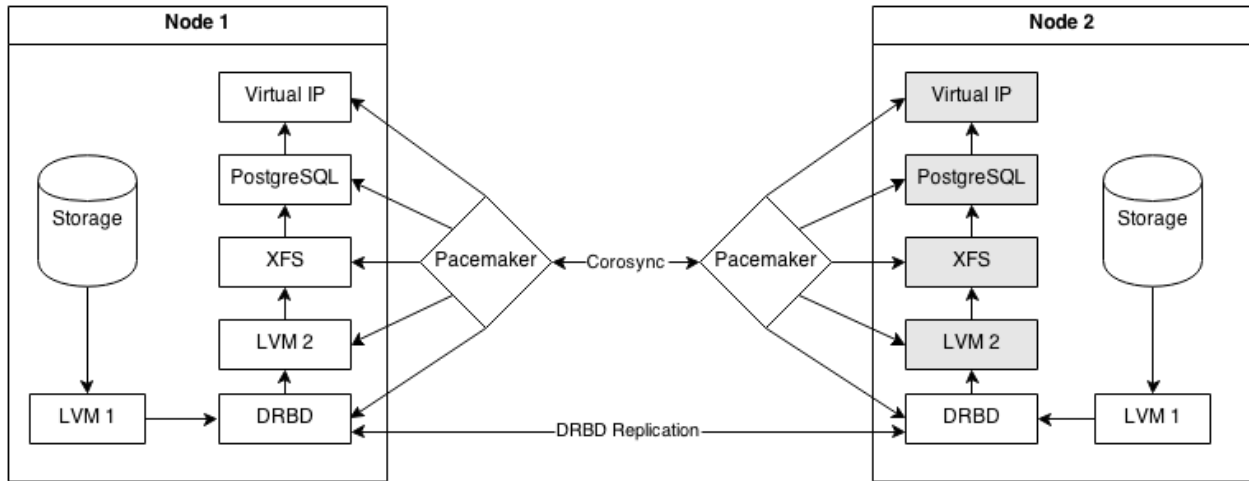
```
root@pg1:~# xfs_db -f -c frag -r /dev/VG_POSTGRES/LV_DATA
actual 1031, ideal 1024, fragmentation factor 0.68%
```

```
root@pg1:~# lvdisplay VG_POSTGRES/snap | grep snap
LV Path          /dev/VG_POSTGRES/snap
LV Name          snap
LV snapshot status active destination for LV_DATA
Allocated to snapshot 13.72%
```

```
0: cs:Connected ro:Secondary/Secondary ds:UpToDate/UpToDate C r-----
ns:13795078 nr:14943 dw:13810066 dr:16201494 al:471 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:0
```

```
0: cs:StandAlone ro:Secondary/Unknown ds:Inconsistent/DUnknown r-----
ns:86085 nr:13797935 dw:18074109 dr:4199604 al:140 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:4190044
```

Chapter 10: Cluster Control



```

Last updated: Sun Nov 20 18:47:01 2016
Last change: Sun Nov 20 18:28:16 2016 by hacluster via crmd on pg1
Stack: corosync
Current DC: pg1 (version 1.1.14-70404b0) - partition with quorum
2 nodes and 0 resources configured

Online: [ pg1 pg2 ]

```

```

root@pg1:~# crm configure show
node 1084766218: pg1
node 1084766228: pg2
property cib-bootstrap-options: \
  have-watchdog=false \
  dc-version=1.1.14-70404b0 \
  cluster-infrastructure=corosync \
  cluster-name=debian \
  stonith-enabled=false \
  no-quorum-policy=ignore \
  default-resource-stickiness=100

```

```

root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]

```

```

root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]
pg_lvm (ocf::heartbeat:LVM): Started

```

```

root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]
pg_lvm (ocf::heartbeat:LVM): Started
pg_fs (ocf::heartbeat:Filesystem): Started

```

Last updated: Mon Nov 21 19:40:12 2016
Last change: Mon Nov 21 19:37:34 2016 by hacluster via crmd on pg1
Stack: corosync
Current DC: pg2 (version 1.1.14-70404b0) - partition with quorum
2 nodes and 3 resources configured

Online: [pg1 pg2]

Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [pg1]
Slaves: [pg2]
pg_lvm (ocf::heartbeat:LVM): Started pg1

```
root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]
pg_lvm (ocf::heartbeat:LVM): Started
pg_fs (ocf::heartbeat:Filesystem): Started
pg_lsb (lsb:postgresql-ha): Started
```

```
root@pg1:~# ifconfig | grep -A3 :pgvip
eth1:pgvip Link encap:Ethernet HWaddr 08:00:27:28:9d:8f
inet addr:192.168.56.50 Bcast:192.168.56.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
```

```
root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]
pg_lvm (ocf::heartbeat:LVM): Started
pg_fs (ocf::heartbeat:Filesystem): Started
pg_lsb (lsb:postgresql-ha): Started
pg_vip (ocf::heartbeat:IPAddr2): Started
```

```
root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]
pg_lvm (ocf::heartbeat:LVM): Started
pg_fs (ocf::heartbeat:Filesystem): Started
pg_lsb (lsb:postgresql-ha): Started
pg_vip (ocf::heartbeat:IPAddr2): Started
pg_mail (ocf::heartbeat:MailTo): Started
```

```
root@pg1:~# crm resource status
Master/Slave Set: ms_drbd_pg [drbd_pg]
Masters: [ pg1 ]
Slaves: [ pg2 ]
pg_mail (ocf::heartbeat:MailTo): Started
Resource Group: PGServer
pg_lvm (ocf::heartbeat:LVM): Started
pg_fs (ocf::heartbeat:Filesystem): Started
pg_lsb (lsb:postgresql-ha): Started
pg_vip (ocf::heartbeat:IPAddr2): Started
```

```
root@pg2:~# crm configure show | egrep 'colocation|order'
colocation col_pg_drbd inf: PGServer ms_drbd_pg:Master
colocation col_pg_mail inf: pg_mail PGServer
order ord_pg inf: ms_drbd_pg:promote PGServer:start
```

```
root@pg1:~# crm resource status PGServer
resource PGServer is running on: pg2
```

Online: [pg1 pg2]

Master/Slave Set: ms_drbd_pg [drbd_pg]

Masters: [pg1]

Slaves: [pg2]

pg_mail (ocf::heartbeat:MailTo): Started pg2

Resource Group: PGServer

pg_lvm (ocf::heartbeat:LVM): Started pg1

pg_fs (ocf::heartbeat:Filesystem): Started pg1

pg_lsb (lsb:postgresql-ha): Stopped

pg_vip (ocf::heartbeat:IPAddr2): Stopped

Online: [pg2]

OFFLINE: [pg1]

Master/Slave Set: ms_drbd_pg [drbd_pg]

Masters: [pg2]

Stopped: [pg1]

pg_mail (ocf::heartbeat:MailTo): Started pg2

Resource Group: PGServer

pg_lvm (ocf::heartbeat:LVM): Started pg2

pg_fs (ocf::heartbeat:Filesystem): Started pg2

pg_lsb (lsb:postgresql-ha): Started pg2

pg_vip (ocf::heartbeat:IPAddr2): Started pg2

Chapter 11: Data Distribution

```
database_name | transactions | writes | queries
-----+-----+-----+-----
pgbench      |      183786 | 20104396 | 336705
(1 row)
```

```
table_name | num_rows | size_mb | writes
-----+-----+-----+-----
pgbench_accounts | 20000000 | 2993 | 20025503
(1 row)
```

```
pgbench=# SELECT srvname, srvoptions
pgbench=# FROM pg_foreign_server;
 srvname | srvoptions
-----+-----
primary_db | {host=pg-primary,dbname=pgbench}
(1 row)
```

```
user_name | server_name | map_options
-----+-----+-----
bench_user | primary_db | {user=bench_user,password=testing}
(1 row)
```

```
postgres@pg-report:~$ psql pgbench -c '\d pgbench_accounts'
Foreign table "public.pgbench_accounts"
 Column | Type | Modifiers | FDW Options
-----+-----+-----+-----
aid     | integer | not null |
bid     | integer |          |
abalance | integer |          |
filler  | character(84) |          |
Server: primary_db
FDW Options: (table_name 'pgbench_accounts')
```

```
pgbench=# DROP TABLE pgbench_accounts;
ERROR: "pgbench_accounts" is not a table
HINT: Use DROP FOREIGN TABLE to remove a foreign table.
```

```
Foreign Scan on public.pgbench_accounts
(cost=100.00..628372.08 rows=4 width=12)
Output: aid, bid, abalance
Remote SQL: SELECT aid, bid, abalance
            FROM public.pgbench_accounts
            WHERE ((aid >= 500000)) AND ((aid <= 500004))
```

```
Aggregate (cost=628372.09..628372.10 rows=1 width=8)
Output: sum(abalance)
-> Foreign Scan on public.pgbench_accounts
(cost=100.00..628372.08 rows=4 width=4)
Output: aid, bid, abalance, filler
Remote SQL: SELECT abalance
            FROM public.pgbench_accounts
            WHERE ((aid >= 500000)) AND ((aid <= 500004))
```

```
Hash Join (cost=628472.13..1631744.17 rows=4 width=12)
Output: a2.aid, a2.bid, a2.abalance
Hash Cond: (a2.aid = a1.aid)
-> Foreign Scan on public.pgbench_accounts a2
(cost=100.00..928372.00 rows=20000000 width=12)
Output: a2.aid, a2.bid, a2.abalance, a2.filler
Remote SQL: SELECT aid, bid, abalance
            FROM public.pgbench_accounts
-> Hash (cost=628372.08..628372.08 rows=4 width=4)
Output: a1.aid
-> Foreign Scan on public.pgbench_accounts a1
(cost=100.00..628372.08 rows=4 width=4)
Output: a1.aid
Remote SQL: SELECT aid
            FROM public.pgbench_accounts
            WHERE ((aid >= 500000)) AND ((aid <= 500004))
```

```
Foreign Scan on public.pgbench_accounts_self
(cost=100.00..300100.08 rows=4 width=12)
Output: aid, bid, abalance
Remote SQL: SELECT aid, bid, abalance
            FROM public.v_pgbench_accounts_self_join
            WHERE ((aid >= 500000)) AND ((aid <= 500004))
```

```
Index Scan using idx_pgbench_accounts_aid on pgbench_accounts
(cost=0.29..10.41 rows=46 width=97)
(actual time=0.007..0.017 rows=50 loops=1)
Index Cond: ((aid >= 400001) AND (aid <= 400050))
```

```
Planning time: 0.099 ms
Execution time: 0.037 ms
```

```

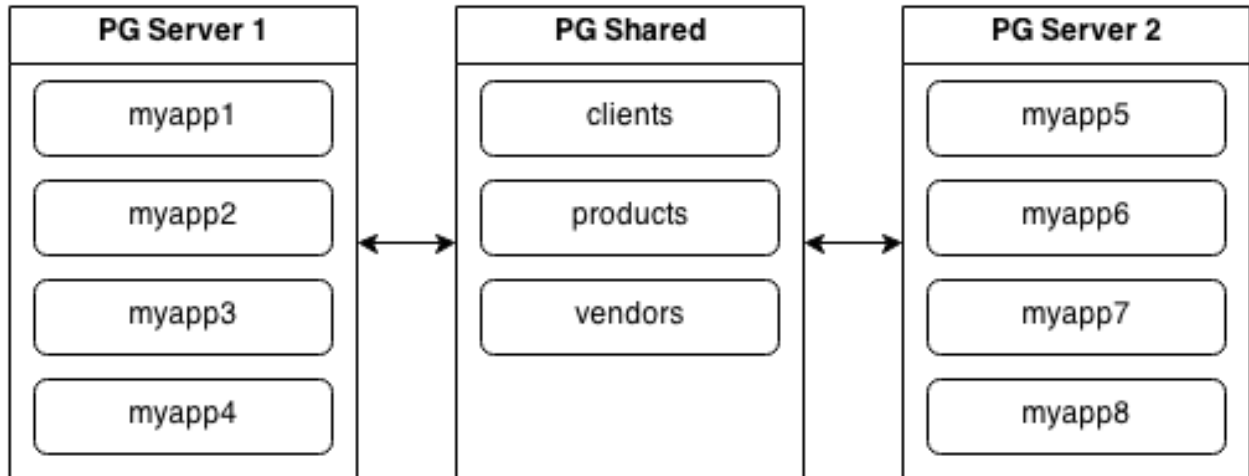
shard=# SELECT (newval & 2047) AS id_value,
shard=#         (newval >> 11) & 2047 AS shard_id,
shard=#         (newval >> 22) / 1000 / 3600 / 24 AS days
shard=# FROM (SELECT shard.next_unique_id(15)
shard=#         AS newval) nv;

```

```

id_value | shard_id | days
-----+-----+-----
      14 |        15 | 328

```



```

shard=# SELECT * FROM shard.shard_map;
map_id | shard_id | source_schema | shard_schema | server_name
-----+-----+-----+-----+-----
      1 |        1 | myapp         | myapp1       | pg-primary
      2 |        2 | myapp         | myapp2       | pg-primary
(2 rows)

```

