

Upgrading and Testing Rollback of a Multi-TB Amazon Aurora PostgreSQL Database Using Blue-Green Deployment

PRESENTED BY

Shailesh Rangani

Director and Global Practice Lead – PostgreSQL
Datavail



datAvail

Shailesh Rangani

Director & Global
Practice Lead -
PostgreSQL,
Datavail



- ✔ Shailesh Rangani is Practice Lead of PostgreSQL Services with 18+ years' experience in the database domain.
- ✔ He holds certifications on cloud platforms like AWS, Azure, and OCI, along with database platforms like Oracle, PostgreSQL, MongoDB, and Db2 LUW.
- ✔ He is an expert in the design, deployment, administration, and management of data-intensive applications that enable organizations to effectively analyze and process large volumes of structured and unstructured data.


Connect or follow me on LinkedIn:



<https://www.linkedin.com/in/shaileshrangani/>

Datavail at a Glance

Delivering a superior approach to leverage data through application of a tech-enabled global delivery model & deep specialization in databases, data management, and application services.



16⁺

Years

building and operating mission critical data and application systems



\$25^M

Invested

in IP that improves the service experience and drives efficiency



1,000⁺

Employees

staffed 24x7, resolving over 2,000,000 incidents per year



Fill out the session
evaluation and enter to

**WIN FREE
BEER FOR A
YEAR!**



What Are the Major Upgrade Options?

Aurora PostgreSQL – Major Upgrade

In-Place



Aurora PostgreSQL – v11
Aurora PostgreSQL – v15


Import/Export



Aurora PostgreSQL – v11
EXPORT

Aurora PostgreSQL – v15
IMPORT

Logical Replication/DMS/Pub-Sub

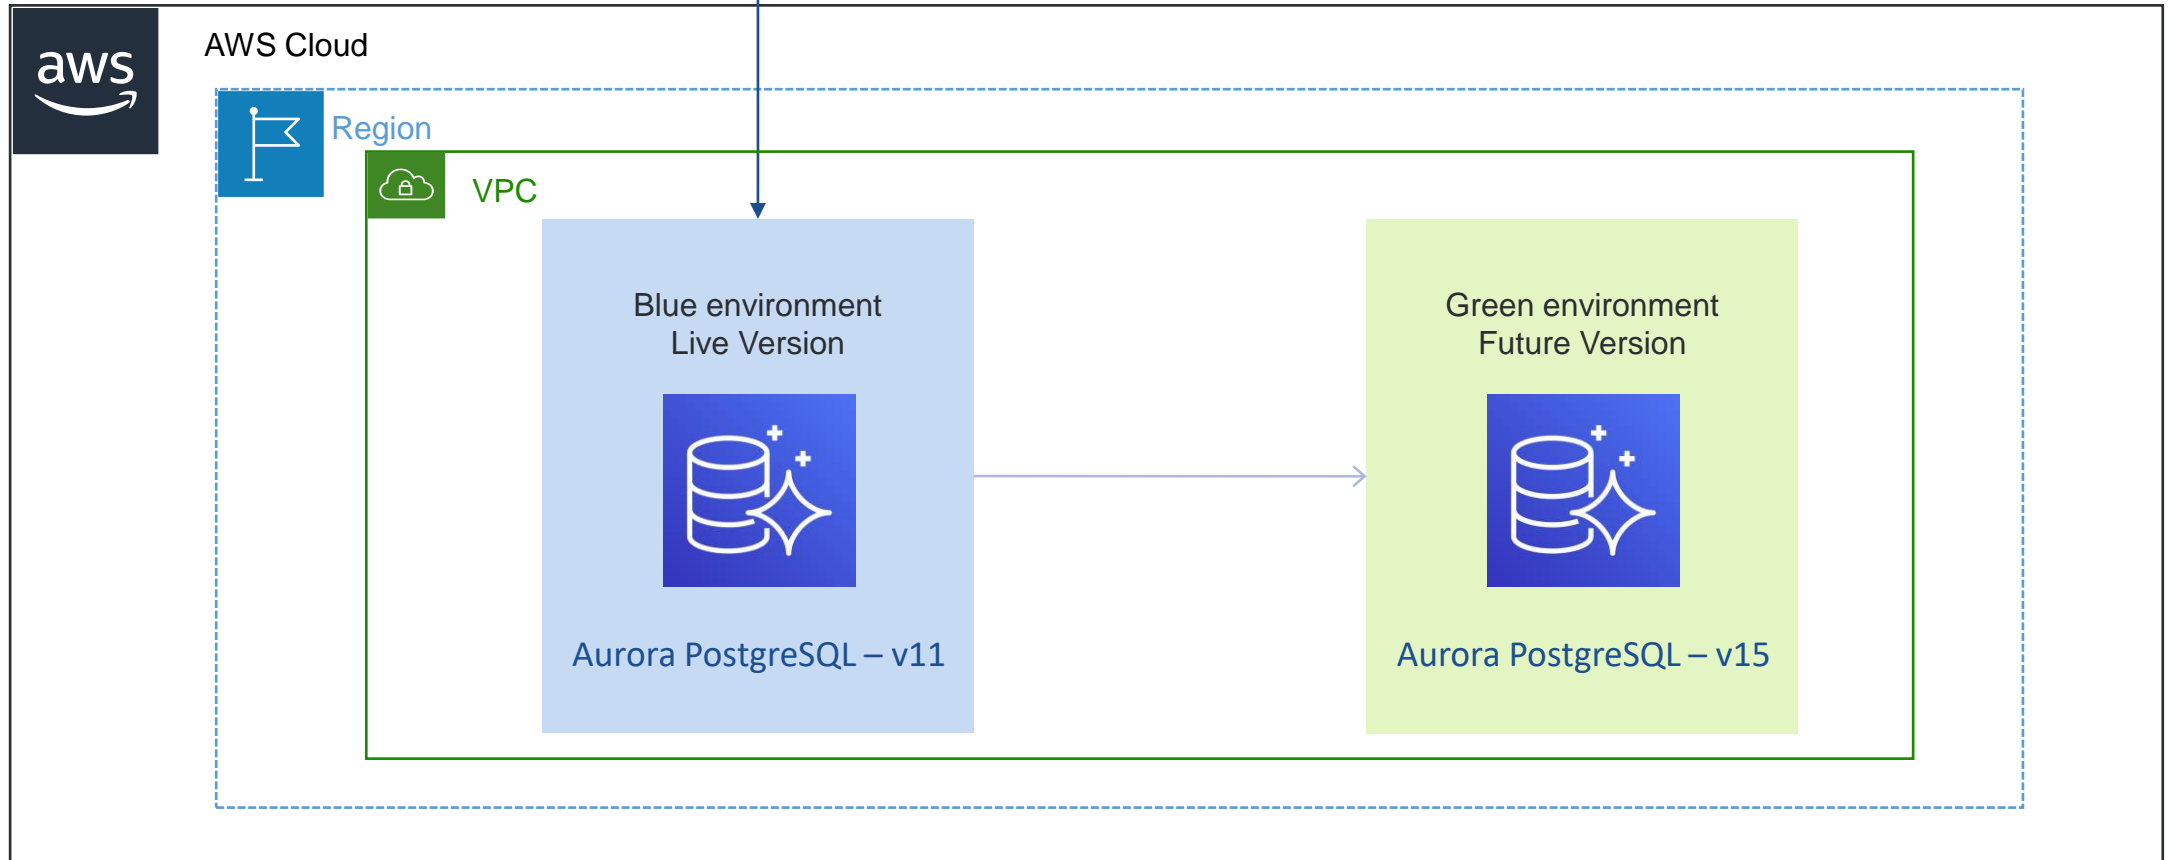



Aurora PostgreSQL – v11

AWS DMS

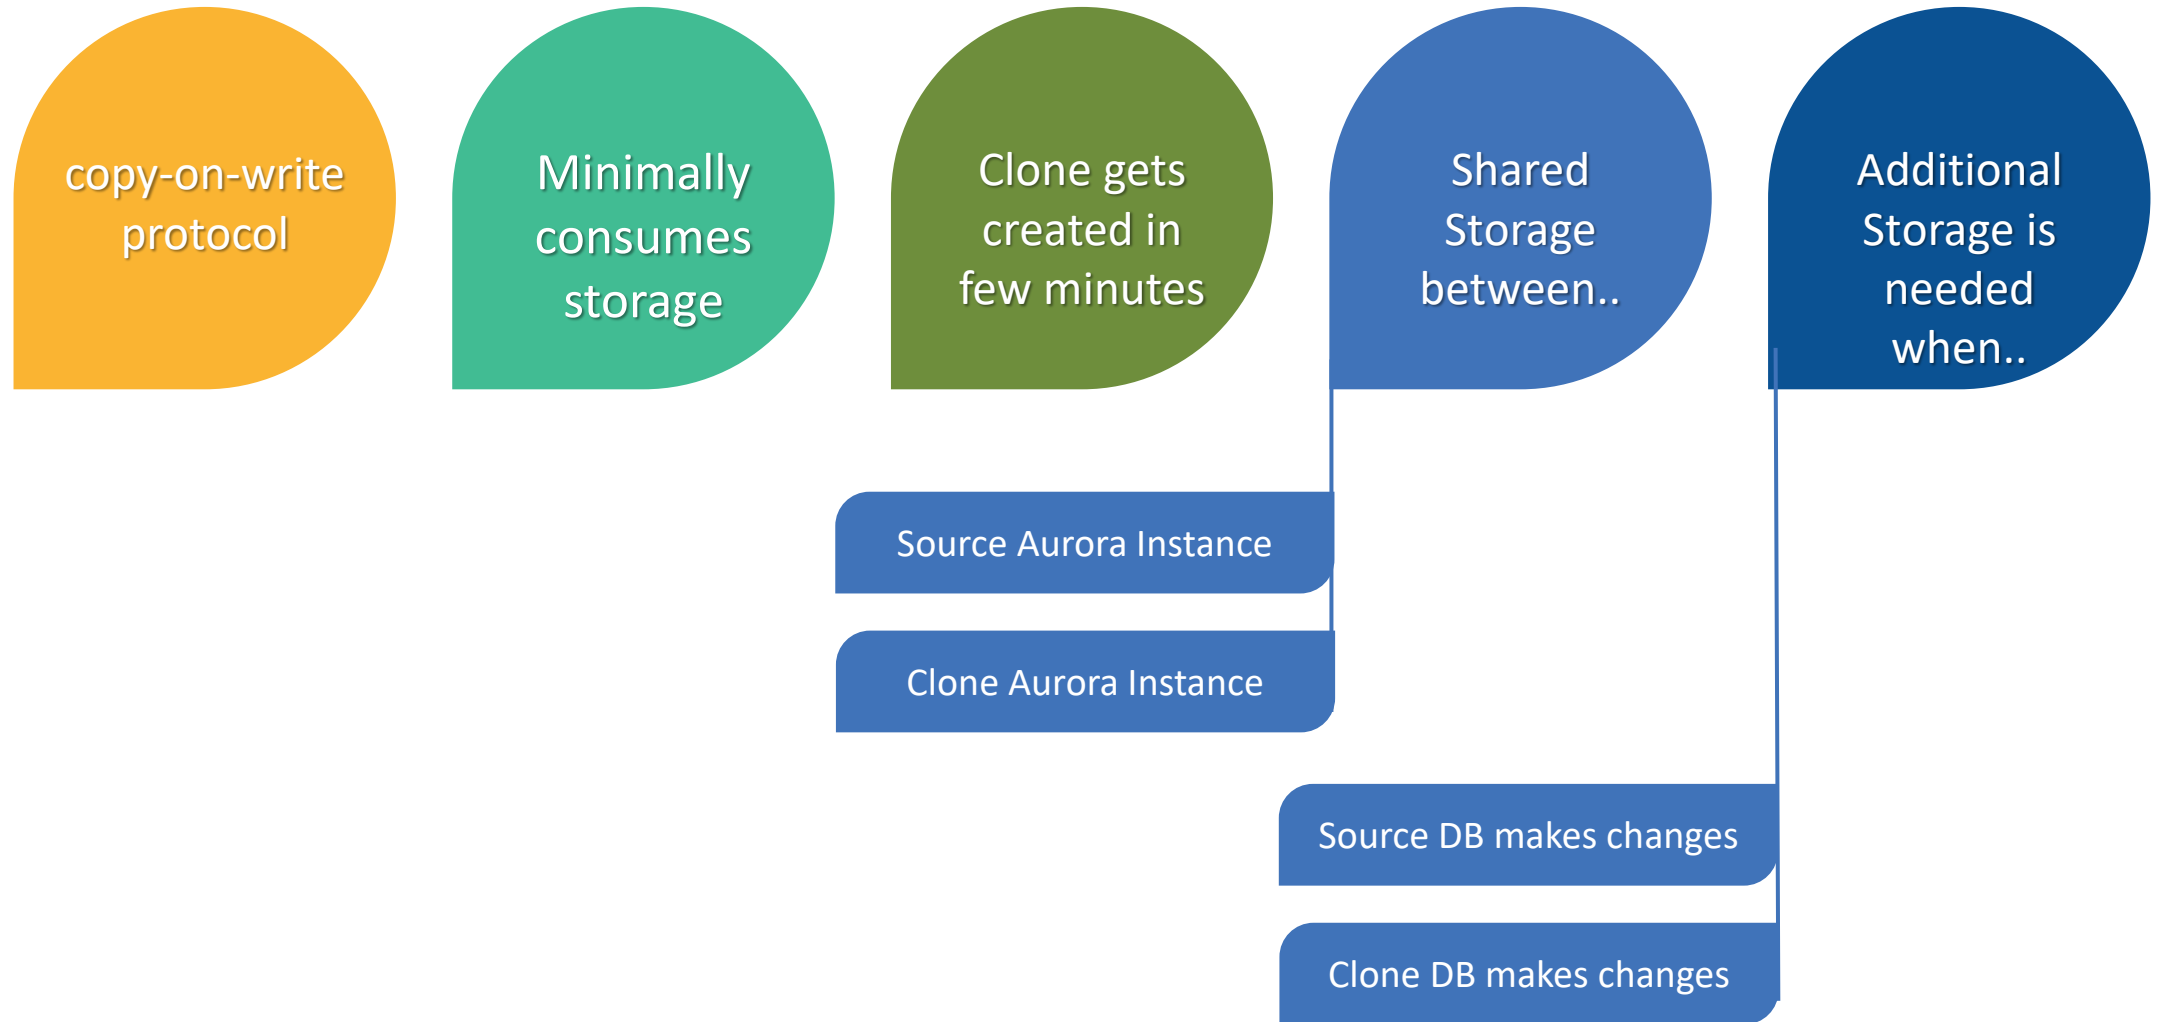
Aurora PostgreSQL – v15

Upgrade using Blue/Green deployment




Aurora Fast Cloning

Aurora Fast Cloning




Aurora PostgreSQL – Source



Aurora PostgreSQL – Source


Aurora PostgreSQL – Clone



Actions ▲	Restore from S3
Stop temporarily	
Start database activity stream	
Delete	
Set up EC2 connection	
Set up Lambda connection	
Upgrade now	
Upgrade at next window	
Add AWS Region	
Add reader	
Create cross-Region read replica	
Create clone	


Aurora PostgreSQL – Source

Aurora PostgreSQL – Clone



Aurora PostgreSQL – Source

Aurora PostgreSQL – Clone



Steps for Aurora PostgreSQL Major Upgrade - Blue/Green



Preparing Source DB


Preparing Source DB (PG11)

- ✓ Enable logical replication

```
rds.logical_replication = 1
```

- ✓ Setup WAL cache

```
rds.logical_wal_cache=1GB
```



(PG Version >= 14.5,
13.8, 12.12, and 11.17)

- ✓ Validate that all the tables have primary keys (including Partition tables)
- ✓ Analyze/Group tables based on number of transactions/loads
- ✓ Validate if DB has sequences

- ✓ Validate replication slot configuration

```
max_logical_replication_workers;  
max_replication_slots;
```


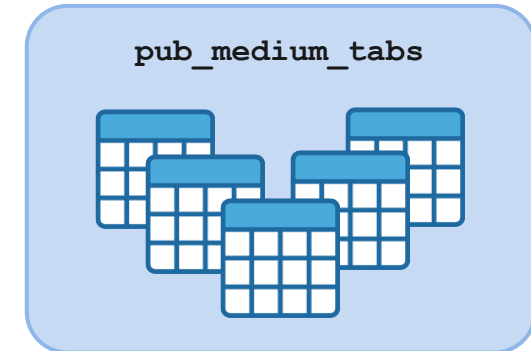
Preparing Source DB (PG11)

✓ Create multiple publications

```
CREATE PUBLICATION pub_large_tab1;  
CREATE PUBLICATION pub_large_tab2;  
CREATE PUBLICATION pub_large_tab3;  
CREATE PUBLICATION pub_medium_tabs;  
CREATE PUBLICATION pub_small_tabs;
```

✓ Add Tables to Publications

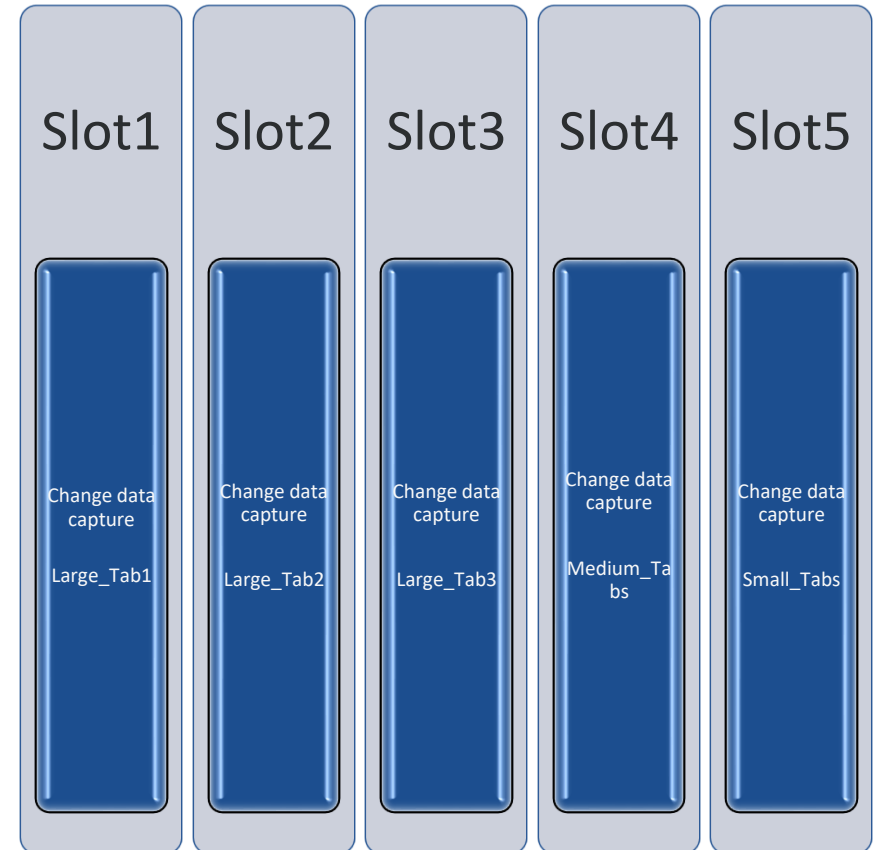
```
ALTER PUBLICATION pub_large_tab1 add table large_table1;  
ALTER PUBLICATION pub_large_tab2 add table large_table2;  
ALTER PUBLICATION pub_large_tab3 add table large_table3;  
ALTER PUBLICATION pub_medium_tabs add table table_mid_1, table_mid_2,...;  
ALTER PUBLICATION pub_small_tabs add table table_sml_1, table_sml_2,...;
```



Preparing Source DB (PG11)

✔ Create replication slots

```
SELECT pg_create_logical_replication_slot('pub_large_tab1', 'pgoutput');  
SELECT pg_create_logical_replication_slot('pub_large_tab2', 'pgoutput');  
SELECT pg_create_logical_replication_slot('pub_large_tab3', 'pgoutput');  
SELECT pg_create_logical_replication_slot('pub_medium_tabs', 'pgoutput');  
SELECT pg_create_logical_replication_slot('pub_small_tabs', 'pgoutput');
```





Create Clone and Prepare Target

Preparing Target DB (PG11)

✔ Create clone

datavail-cluster

Modify Actions ▲

Stop temporarily
Start database activity stream
Delete
Set up EC2 connection
Set up Lambda connection
Add AWS Region
Add reader
Create cross-Region read replica
Create clone
Promote

Related

Filter by databases

DB identifier	Status	Role	Engine	Engine version	Region & AZ	Size	CPU	Current activity
datavail-cluster	Available	Regional cluster	Aurora PostgreSQL	11.20	us-east-1	2 instances	-	-
datavail-cluster-instance-0	Available	Writer instance	Aurora PostgreSQL	11.20	us-east-1a	db.t4g.medium	12.59%	0.00 sessions
datavail-cluster-instance-1	Available	Reader instance	Aurora PostgreSQL	11.20	us-east-1c	db.t4g.medium	13.19%	0.00 sessions

DB identifier	Status	Role	Engine	Engine version	Region & AZ	Size	CPU
datavail-cluster	Available	Regional cluster	Aurora PostgreSQL	11.20	us-east-1	2 instances	-
datavail-cluster-instance-0	Available	Writer instance	Aurora PostgreSQL	11.20	us-east-1c	db.t4g.medium	14.57%
datavail-cluster-instance-1	Available	Reader instance	Aurora PostgreSQL	11.20	us-east-1a	db.t4g.medium	14.82%
datavail-cluster-pg15	Available	Regional cluster	Aurora PostgreSQL	11.20	us-east-1	1 instance	-
datavail-cluster-pg15-instance-0	Available	Writer instance	Aurora PostgreSQL	11.20	us-east-1c	db.t4g.medium	13.48%

Preparing Target DB (PG11)

- ✓ Connect to the Clone DB
- ✓ Note down start LSN from the clone

```
datavail=> SELECT * FROM aurora_volume_logical_start_lsn();
 aurora_volume_logical_start_lsn
-----
1CC/324A7EA8
```

- ✓ Drop replication slots

```
select pg_drop_replication_slot('pub_large_tab1');
select pg_drop_replication_slot('pub_large_tab2');
select pg_drop_replication_slot('pub_large_tab3');
select pg_drop_replication_slot('pub_medium_tabs');
select pg_drop_replication_slot('pub_small_tabs');
```



Upgrade Cluster from PG11 to PG15

Upgrade Cluster to PG15 (PG11)

✓ Upgrade to PG15

RDS > Databases > datavail-cluster-pg15

datavail-cluster-pg15

Refresh Modify Actions

Related

Filter by databases

DB identifier	Status	Role	Engine	Engine version	Region & AZ	Size	CPU	Current
<input checked="" type="radio"/> datavail-cluster-pg15	✓ Available	Regional cluster	Aurora PostgreSQL	11.20	us-east-1	1 instance	-	-
<input type="radio"/> datavail-cluster-pg15-instance-0	✓ Available	Writer instance	Aurora PostgreSQL	11.20	us-east-1c	db.t4g.medium	12.12%	

Summary of modifications

You are about to submit the following modifications. Only values that will change are displayed. Carefully verify your changes and click Modify cluster.

Attribute	Current value	New value
Engine version	11.20	15.3
DB cluster parameter group	datavail-cluster-pg-11	default.aurora-postgresql15
DB parameter group		default.aurora-postgresql15

Perform Post upgrade tasks(PG15 - CLONE)

✓ Custom Parameter group validation

✓ Analyze database



Perform Maintenance Tasks

Maintenance Tasks on (PG15 - CLONE)

✓ Vacuum full on bloated tables

```
datavail=> vacuum full large_table1;  
VACUUM  
datavail=> vacuum full large_table2;  
VACUUM  
:  
:
```

✓ REINDEX on bloated indexes

```
datavail=> REINDEX DATABASE datavail;  
REINDEX
```

```
reindexdb -h <aurora-end-point> -p 5432 -d datavail -U postgres -e -j 15 -v
```

✓ Analyze database

```
datavail=> analyze;  
ANALYZE
```

```
vacuumdb -h <aurora-end-point> -p 5432 -d prod -U postgres -e -j 15 -v -Z
```



Setup Replication Between PG11 and PG15

Setup Replication (PG15)

✓ Create subscriptions

```
datavail=> CREATE SUBSCRIPTION sub_large_tab1
  CONNECTION 'host=datavail-cluster.cluster-czfcfxk38byn.us-east-1.rds.amazonaws.com dbname=datavail
port=5432 user=sysdba password=*****'
  PUBLICATION pub_large_tab1
  WITH (
    copy_data = false,
    create_slot = false,
    enabled = false,
    connect = true,
    slot_name = 'pub_large_tab1'
  );
CREATE SUBSCRIPTION
:
:
```

✓ Get replication origin names

```
datavail=> SELECT * FROM pg_replication_origin;
 roident |  roname
-----+-----
       1 | pg_53563397
       2 | pg_53563398
       3 | pg_53563399
       4 | pg_53563400
       5 | pg_53563401
```

Setup Replication (PG15)

- ✔ Start replication from the clone time (LSN)

```
datavail=> SELECT pg_replication_origin_advance('pg_53563397', '1CC/324A7EA8');
datavail=> SELECT pg_replication_origin_advance('pg_53563398', '1CC/324A7EA8');
datavail=> SELECT pg_replication_origin_advance('pg_53563399', '1CC/324A7EA8');
datavail=> SELECT pg_replication_origin_advance('pg_53563400', '1CC/324A7EA8');
datavail=> SELECT pg_replication_origin_advance('pg_53563401', '1CC/324A7EA8');
```

- ✔ Enable Replication

```
datavail=> ALTER SUBSCRIPTION sub_large_tab1 ENABLE;
datavail=> ALTER SUBSCRIPTION sub_large_tab2 ENABLE;
datavail=> ALTER SUBSCRIPTION sub_large_tab3 ENABLE;
datavail=> ALTER SUBSCRIPTION pub_medium_tabs ENABLE;
datavail=> ALTER SUBSCRIPTION pub_small_tabs ENABLE;
```

Validate Replication (PG11)

✓ Validate replication status

```
datavail=> select slot_name, active, restart_lsn, confirmed_flush_lsn from  
pg_catalog.pg_replication_slots;
```

slot_name	active	restart_lsn	confirmed_flush_lsn
pub_large_tab1	t	1CC/3554DD50	1CC/355BBA58
pub_large_tab2	t	1CC/3554DD50	1CC/355BBA58
pub_large_tab3	t	1CC/3554DD50	1CC/355BBA58
pub_medium_tabs	t	1CC/3554DD50	1CC/355BBA58
pub_small_tabs	t	1CC/3554DD50	1CC/355BBA58

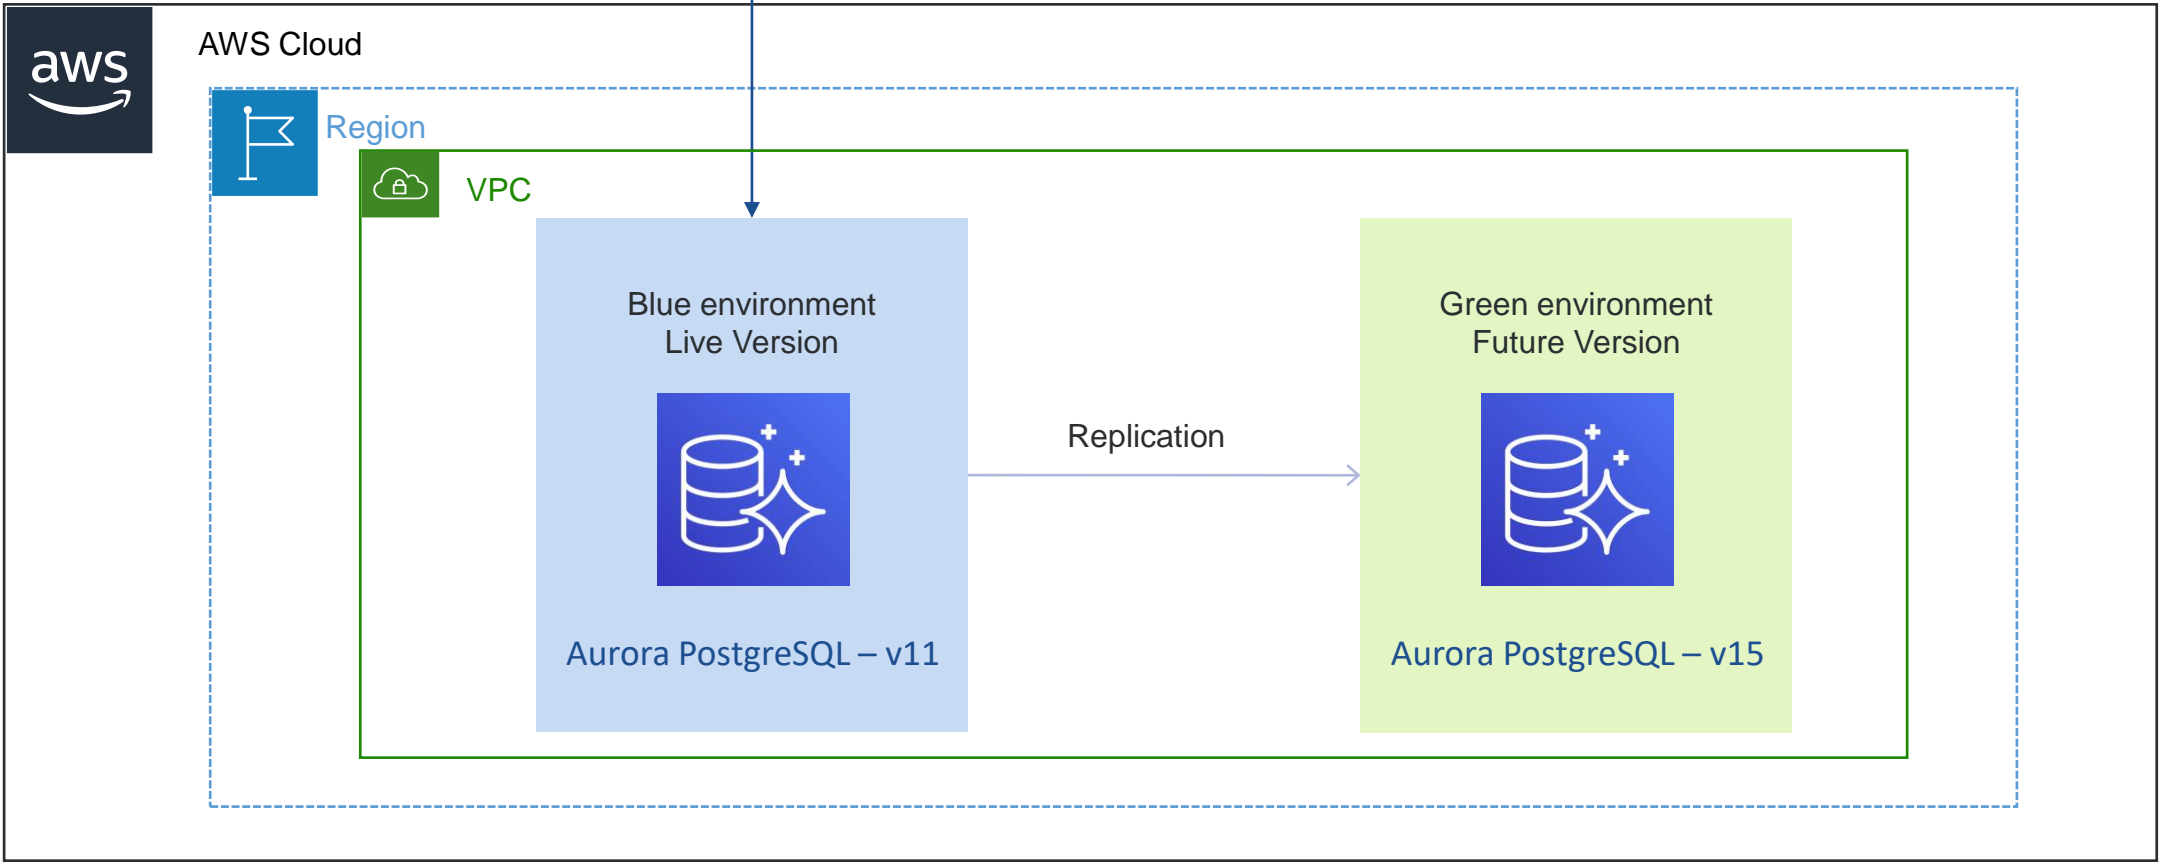
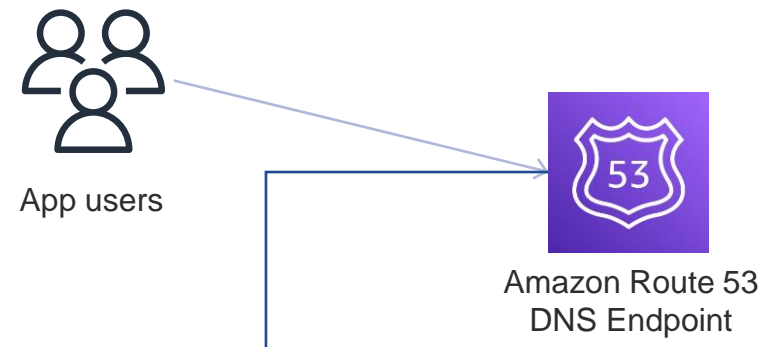
(5 rows)

✓ Validate replication status - By size

```
datavail=> SELECT slot_name, active, active_pid,  
pg_size_pretty(pg_wal_lsn_diff(pg_current_wal_lsn(), confirmed_flush_lsn)) AS diff_size,  
pg_wal_lsn_diff(pg_current_wal_lsn(), confirmed_flush_lsn) AS diff_bytes  
FROM pg_replication_slots  
WHERE slot_type = 'logical';
```

slot_name	active	active_pid	diff_size	diff_bytes
pub_large_tab1	t	1791	136 bytes	136
pub_large_tab2	t	1966	136 bytes	136
pub_large_tab3	t	2117	136 bytes	136
pub_medium_tabs	t	2117	136 bytes	136
pub_small_tabs	t	2117	136 bytes	136

(5 rows)





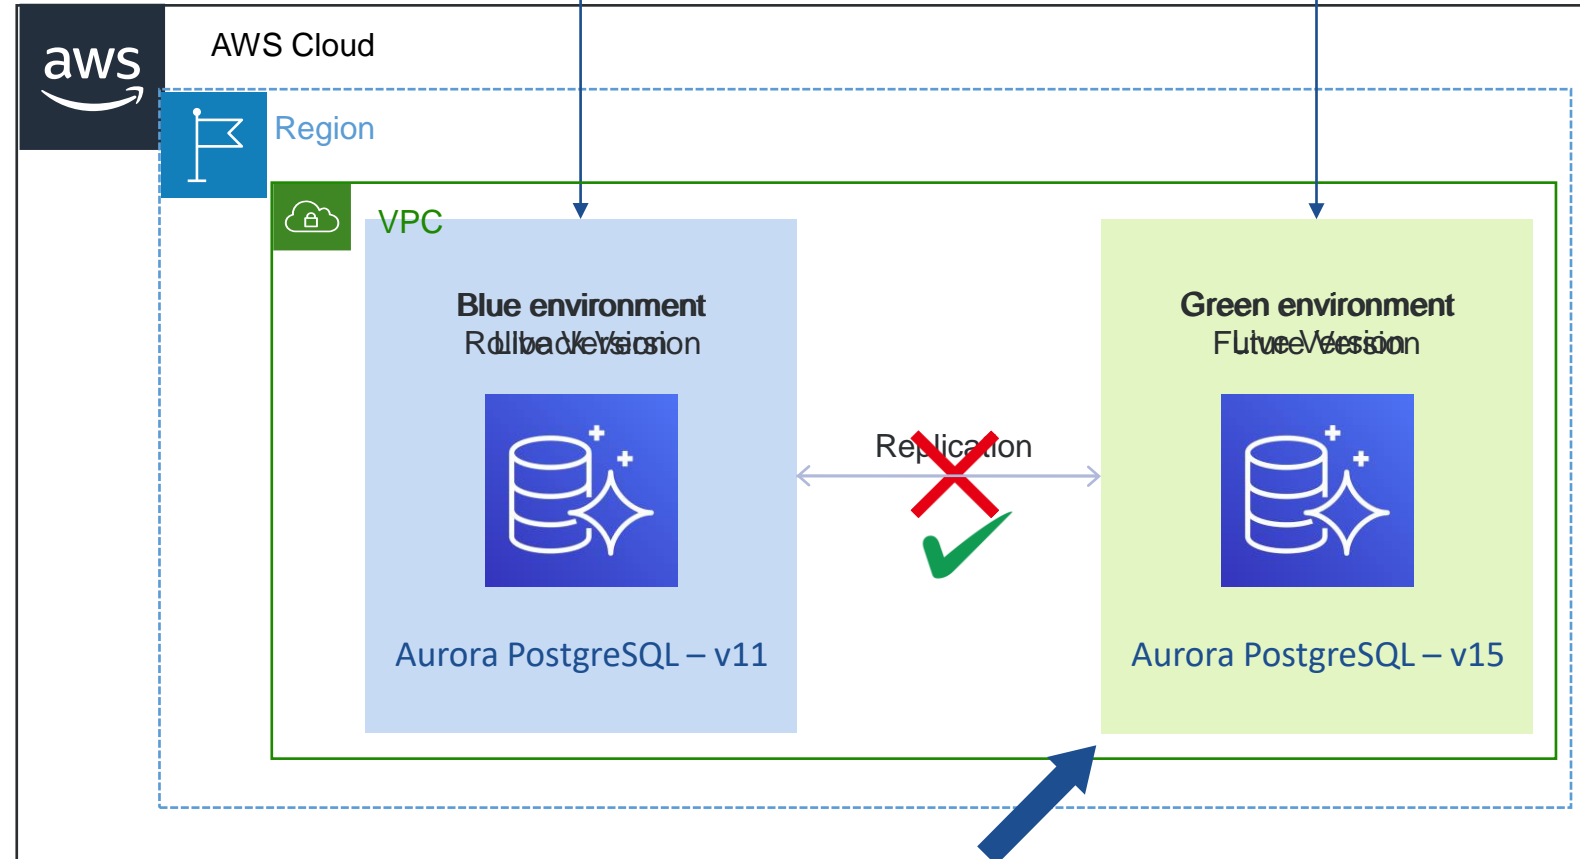
Cutover to PG15



App users



Amazon Route 53
DNS Endpoint



- ✔ Setup maintenance page OR Stop applications/update security groups
- ✔ Validate replication lag (Should be 0)
- ✔ Reset sequences on Green (PG15)
- ✔ Break replication between PG11 and PG15
 - Drop Subscription on Green (PG15)
 - Drop Publication on Blue (PG11)
- ✔ Switch the roles of the Blue/Green
- ✔ PG15
 - Create publications
 - Create replication slots
- ✔ PG11
 - Create subscriptions
- ✔ Validate replication lag (Should be 0)
- ✔ Update CNAME
- ✔ Remove maintenance page/Start applications



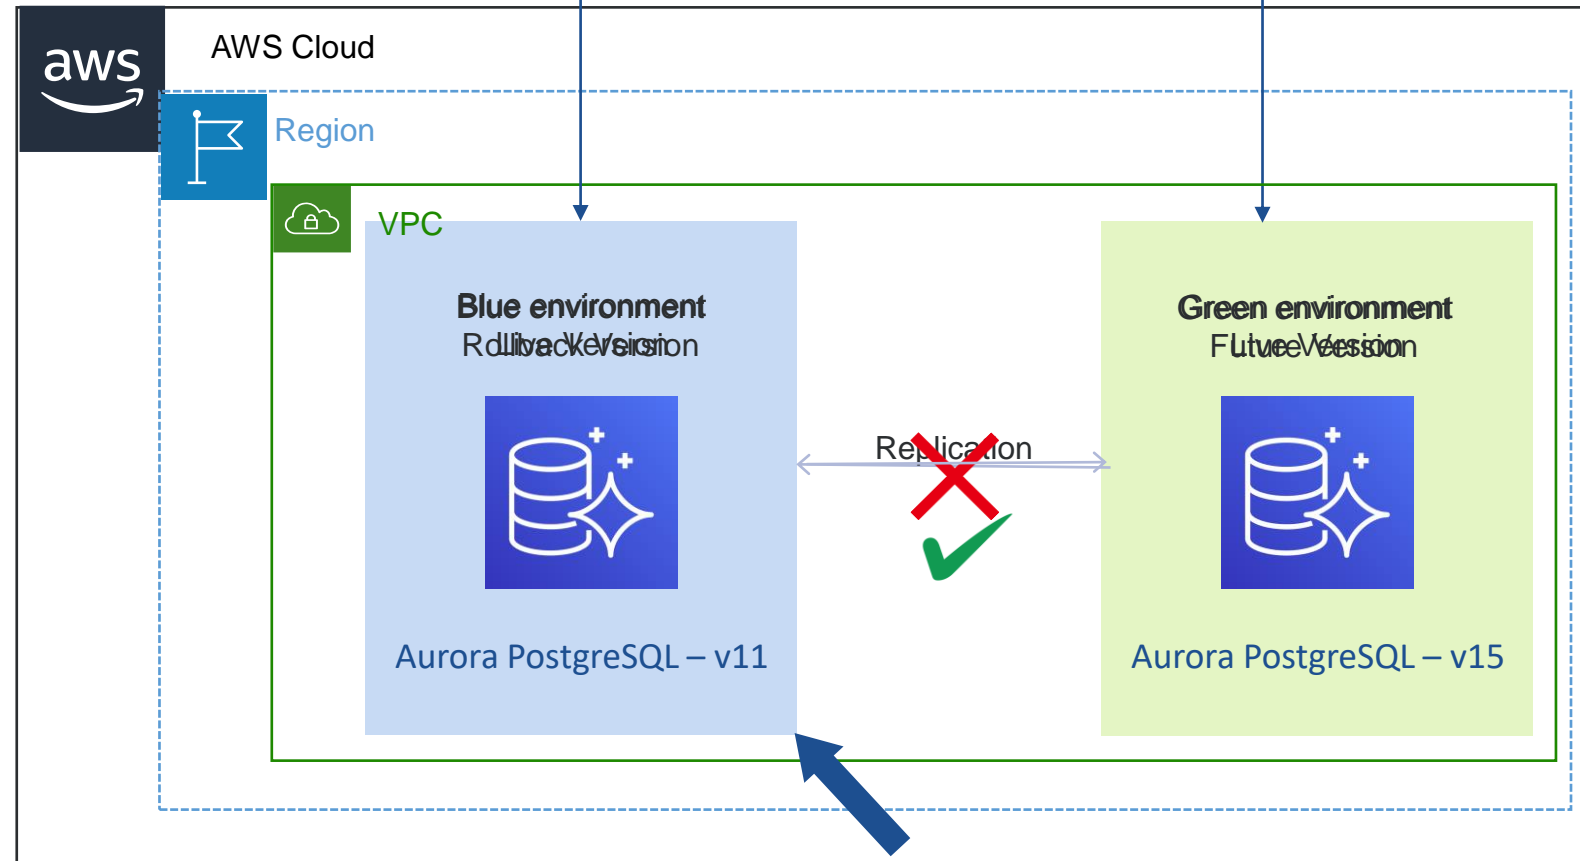
Rollback to PG11



App users



Amazon Route 53
DNS Endpoint

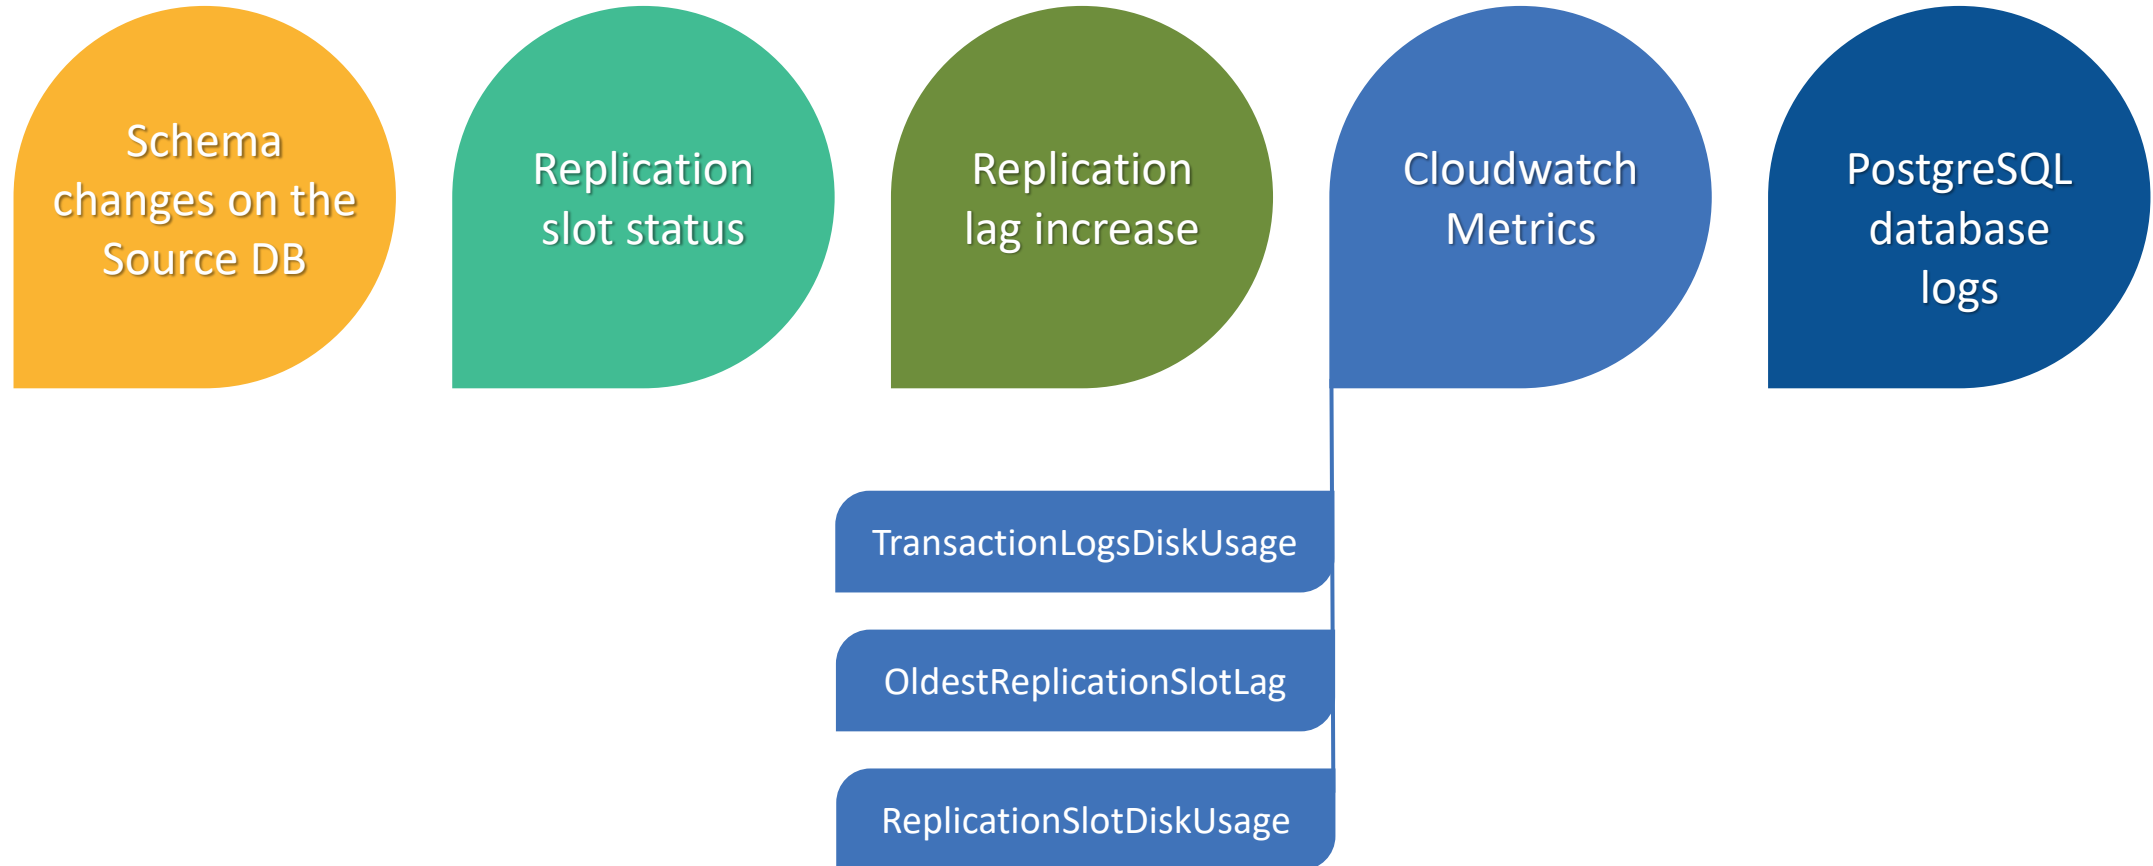


- ✔ Setup maintenance page OR Stop applications/update security groups
- ✔ Validate replication lag (Should be 0)
- ✔ Reset sequences on Blue (PG11)
- ✔ Break replication between PG15 and PG11
 - Drop Subscription on Blue (PG11)
 - Drop Publication on Green (PG15)
- ✔ Switch the roles of the Blue/Green PG11
 - Create publications
 - Create replication slots
- ✔ PG15
 - Create subscriptions
- ✔ Validate replication lag (Should be 0)
- ✔ Update CNAME
- ✔ Remove maintenance page/Start applications



Observability & Alerting

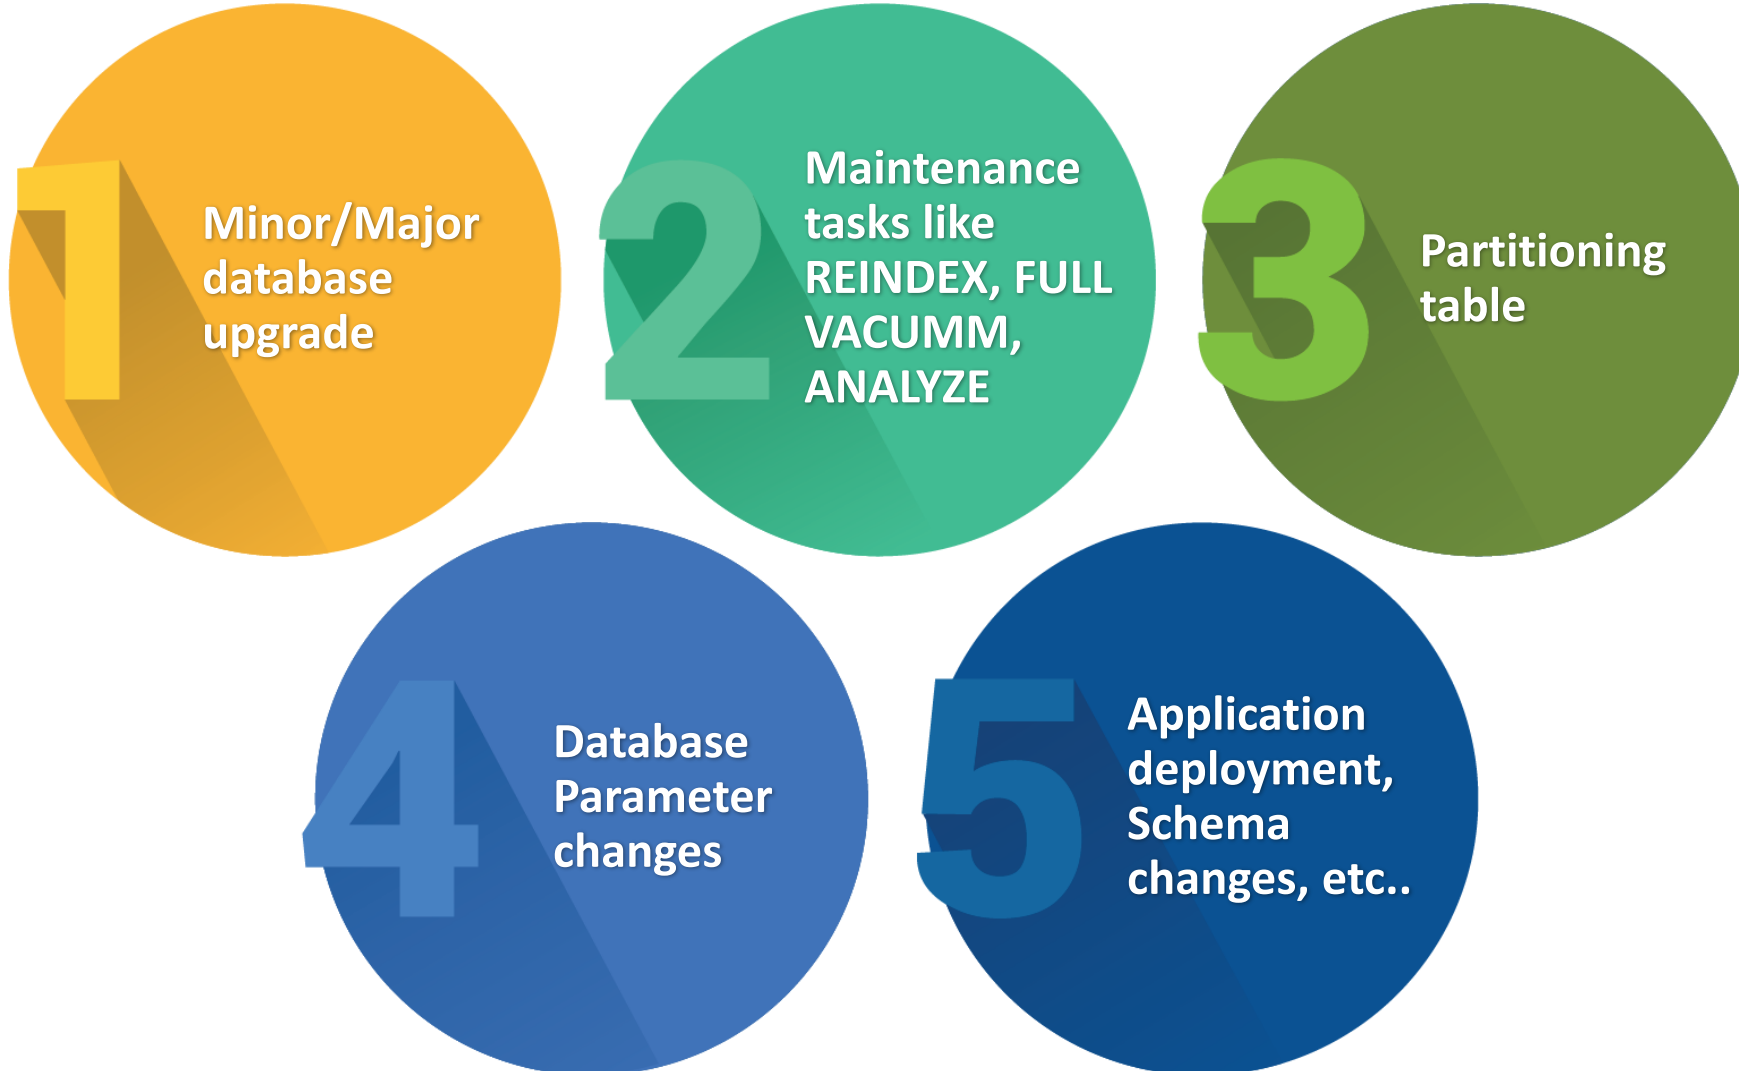
Observability/Alerting





Use Cases

Blue/Green Deployment Use Cases



Fill out the session
evaluation and enter to

**WIN FREE
BEER FOR A
YEAR!**





Q

&

A

Thank You



Shailesh Rangani

Director & Global Practice Lead : PostgreSQL Services

shailesh.rangani@datavail.com

+1 866-623-4956

datAvail