



Migrating PostgreSQL Databases to VMware Data Services Manager

Best Practices Document

Table of contents

Executive Summary	3
Business Benefits	3
Target Audience	3
Technology Overview	4
VMware Cloud Foundation	4
VMware Data Services Manager	4
PostgreSQL Migration Methodology	5
Prepare Source PostgreSQL database	6
Prepare Target PostgreSQL database in Data Services Manager	7
Installation and Deployment	7
Configure Infrastructure Policy	7
Configure object storage	7
Provision PostgreSQL databases in Data Services Manager	7
Use Case 1 – Migration with Logical Replication	8
Migration Prerequisites	8
Migration Steps	8
Tested Configuration	12
Use Case 2 – Migration with Backup and Restore Methodology	13
Migration Prerequisites	13
Migration Steps (pg_dump to migrate single database)	13
Migration Steps (pg_dumpall to migrate multiple databases at the source)	15
Tested Configuration	16
Use Case 3 – Migration with Bucardo	18
Best Practices and Key Considerations	19

Executive Summary

VMware Data Services Manager™ (DSM) offers a comprehensive solution for modern database and data services management within VMware Cloud Foundation environments. It functions as a powerful database-as-a-service (DBaaS) toolkit, enabling organizations to achieve on-demand provisioning and automated management of PostgreSQL and other databases, all within a highly virtualized infrastructure.

This document specifically addresses best practices for a "brownfield migration", which involves migrating existing PostgreSQL databases to new PostgreSQL databases that are managed by VMware Data Services Manager. Furthermore, the document outlines a step-by-step approach for executing the migration, ensuring that the desired business requirements for performance, availability, and scalability are fully met in the new DSM-managed environment.

Business Benefits

Migration to VMware Data Services Manager will bring the following business benefits:

- **Faster provisioning for rich data services:** Data services templates that are pre-packaged and certified by VMware and partners enabling organizations to quickly spin up desired data platforms.
- **Automated operations:** Common routine tasks such as lifecycle management, resource handling, scaling, data protection and availability can all be automated with Data Services Manager.
- **Extensive monitoring:** Gain visibility into database and underlying resources, including metrics such as the utilization of CPU, memory, network, local, and cloud storage, as well as curated alarms and notifications for improved visibility into the health of both the database and the underlying infrastructure.
- **Ready business for Private AI:** Enable enterprise to customize models and run generative AI and RAG workloads with vector database capability to enable fast querying of data and enhanced outputs of different LLMs.
- **Seamless integration with the private cloud platform:** Empower IT organizations with enterprise-hardened data services that are deeply integrated with VMware Cloud Foundation.

Target Audience

This document assumes a knowledge and understanding of VMware Cloud Foundation, VMware Data Services Manager and PostgreSQL databases. The document is designed for architectural and management staff, infrastructure and database administrators to plan and migrate their existing PostgreSQL database infrastructure to VMware Data Services Manager. The goal is to help model the business requirements and take advantage of management simplicities and operational efficiencies achieved with various components of VMware Cloud Foundation platform.

Technology Overview

VMware Cloud Foundation

VMware Cloud Foundation™ is strategically positioned to empower organizations in their pursuit of IT modernization, by addressing three pivotal areas: infrastructure modernization, providing a cloud experience for developers, and enhancing security and resilience. This comprehensive approach ensures that enterprises are well-equipped to face the digital challenges of today and tomorrow. Empowering agility within your organization to drive innovation.

VMware Cloud Foundation is built on trusted VMware technologies for compute, storage, networking, and cloud management. These technologies are the building blocks of a full software-defined data center, integrating all these capabilities into a single platform that is easy to consume.

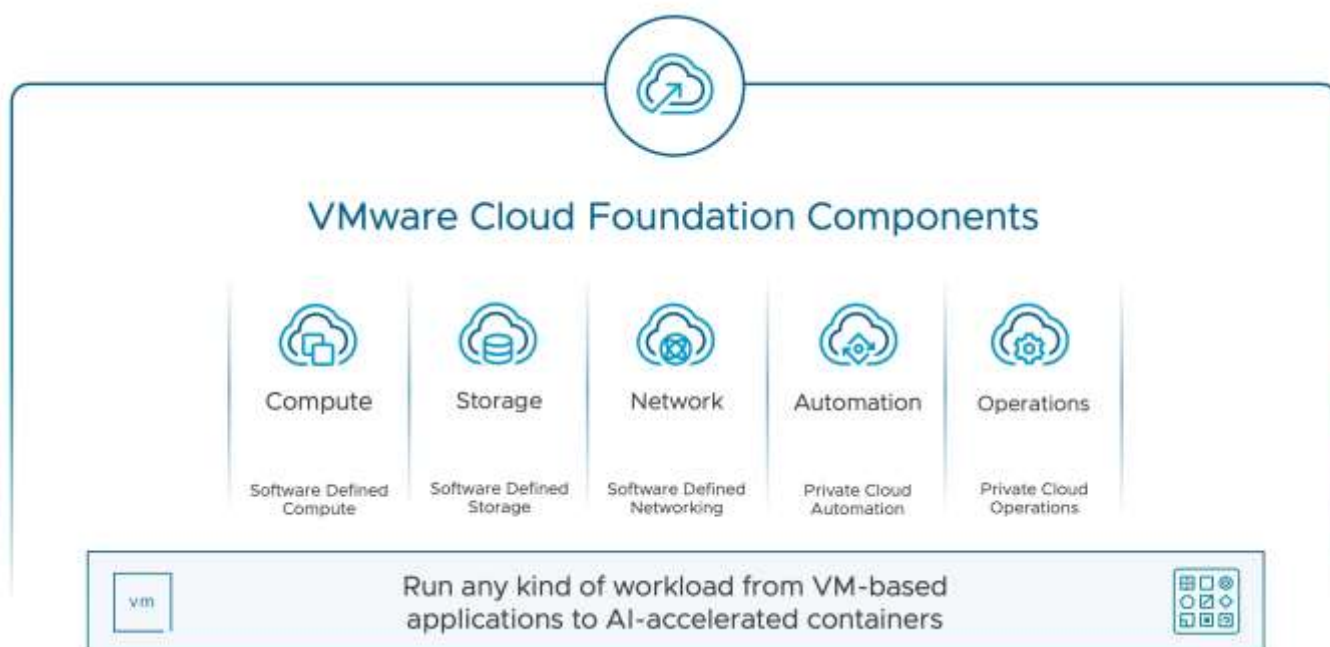


Figure 1 VMware Cloud Foundation Components

VMware Data Services Manager

VMware Data Services Manager allows a user to provision and manage at scale PostgreSQL, MySQL and Microsoft SQL Server (Tech Preview) relational databases. For Generative-AI applications, customers can also use Data Services Manager to deploy a pgvector-enabled PostgreSQL database, providing intelligence about storing and querying vector embeddings for Retrieval Augmented Generation.



PostgreSQL Migration Methodology

To migrate externally managed PostgreSQL databases and clusters to Data Services Manager, you can use the following methodology:

- Migration with PostgreSQL logical replication - replicating data objects and their changes, based upon their replication identity with a publish and subscribe model. This methodology can help minimize the business downtime during the migration process
- Migration with *pg_dump*, *pg_dumpall* and *pg_restore* utilities – migrate databases using traditional backup and restore. This methodology is for general purpose migration but requires certain business downtime during the migration process.
- Migrating legacy PostgreSQL versions into a supported version by using a tool such as Bucardo. This methodology can be considered when the source PostgreSQL is with older versions that do not natively support logical replication and the user wants to minimize the business downtime.

Before finalizing a migration methodology, it is recommended to try out the chosen methodology in a test-dev environment similar to the production setup prior to using it for production database migration..

Prepare Source PostgreSQL database

Before you start migration, prepare the source PostgreSQL database with the following considerations:

- Prepare a list of source databases under migration. Choose the appropriate migration methodology based on the PostgreSQL version, database size, and business SLAs etc.
- Compute resource consideration – collect the source PostgreSQL database sizing in terms of CPU and memory and create VM class in Data Services Manager accordingly.
- Network consideration – verify the network of source PostgreSQL database is reachable to the target data services network configured in the Data Services Manager.
- Storage consideration – collect the source PostgreSQL database storage sizing and performance requirements. If you have storage HA solutions implemented in the source cluster, make sure you have the same or equivalent solution in the target Data Services Manager environment. For example, if you run the target Data Services Manager on vSAN Express Storage Architecture (ESA), it provides FTT=1 by default with RAID-5 and FTT=2 with RAID-6 in storage policy settings.
- PostgreSQL configuration – collect the configuration parameters on the source PostgreSQL databases and you may apply those configurations on the target PostgreSQL in Data Services Manager with the advanced database options.
- PostgreSQL extensions - collect the list of extensions on the source database and make sure all of them are included in the [supported PostgreSQL extension of Data Services Manager](#).

Prepare Target PostgreSQL database in Data Services Manager

You must prepare the target Data Services Manager before migrating the source PostgreSQL databases.

Installation and Deployment

You deploy Data Services Manager as a plug-in with an .ova file available in the Broadcom Software download page. Refer to [Deploying the VMware Data Services Manager plug-in in vSphere Client](#) documentation for detailed steps.

Configure Infrastructure Policy

The infrastructure Policy allows VI admins to define which vSphere infrastructure resources should be used for the placement of data services provisioned by Data Services Manager. It includes the following components.

- Compute resources – a vSphere cluster or resource pool for provisioning the data services.
- Network port groups – networks that a data services virtual machine can be attached to.
- IP Pools – a collection of available IP addresses used for dynamic IP assignment for data services. Based on the number of virtual machines and deployment type (single instance or cluster), assign enough IP addresses for the IP Pools.
- VM classes – specify the vCPU and Memory for a data services virtual machine. Based on the source PostgreSQL instance sizing, create corresponding VM classes in Data Services Manager.
- Storage Policies – specify the datastore placement of a data services virtual machine. For vSAN ESA as an example for target storage, an ESA default storage policy with RAID 5 or RAID 6 is recommended.

Configure object storage

Data Services Manager requires an S3-compatible object storage solution for the following buckets.

- A bucket for storing support bundles
- A bucket for Provider backups
- A bucket for database backups

It is required to configure the desired object storage with network encryption (TLS). Follow the [Data Services Manager Certificate Requirements](#) when generating a certificate for the object storage configured above. Use IP address for the Subject Alternative Name (SAN) and do not leave blank.

Provision PostgreSQL databases in Data Services Manager

Data Services Manager supports both single server and cluster deployment for PostgreSQL databases. You may select the desired deployment type based on your business SLAs.

Data Services Manager supports customized PostgreSQL parameters in Database Options. The default settings of parameters such as *shared_buffers* and *max_connections* for the target database provisioned by Data Services Manager can be overridden with ease according to the source database settings for different use cases. For more details, refer to [PostgreSQL on Data Services Manager](#).

Based on the source PostgreSQL version, it is recommended to provision the same or newer version of PostgreSQL on the Data Services Manager deployment. Check the [VMware Data Services Manager release notes](#) for more details about supported PostgreSQL database versions.

Use Case 1 – Migration with Logical Replication

Logical Replication uses a publish and subscribe model, copying and replicating data objects and their changes based on their replication identity. It allows fine-grained control over data replication between different PostgreSQL databases. Logical Replication replicates individual tables and their changes including INSERTs, UPDATEs, DELETEs. The subscribers pull data from the publications and apply the data in the same order as the publisher so that transactional consistency is guaranteed. Logical Replication supports the following migration use cases.

- The source and target PostgreSQL with different major versions
- The source and target PostgreSQL with different platforms (Linux, Windows, etc.)

Migration Prerequisites

The logical replication replicates the Data Modification Language (DML) changes like INSERT, UPDATE, DELETE and TRUNCATE from the source PostgreSQL database to the target Data Services Manager PostgreSQL database. Before you migrate with logical replication, validate the following checklist

- Verify that you have the necessary CLI tools installed in your source environment including the *pg_dump* and *psql* utilities.
- Verify that you have configured the *pg_hba.conf* file to allow connection between the source environment and the target database environment provisioned by Data Services Manager.
- Verify that you have granted the corresponding replication privilege to the user account that will be used for the logical replication.
- Obtain a list of all the source databases needed to migrate. Logical replication is configured at the database level. For multiple databases, you need to configure logical replication for each of them.
- Understand the restrictions of logical replication before you start the migration process. Refer to [Logical Replication Restrictions](#) for more details.

Migration Steps

- Step 1 – Configure the source PostgreSQL database for logical replication. On the source PostgreSQL database and its replicas, set **wal_level** to **logical**.

Example:

```
sourcedb1=# show wal_level;
wal_level
-----
replica
(1 row)

sourcedb1=# ALTER SYSTEM SET wal_level = logical;
ALTER SYSTEM
```

Restart PostgreSQL service and check **wal_level** and make sure it's changed to **logical**.

Example:

```
sourcedb1=# show wal_level;
wal_level
-----
logical
(1 row)
```


Migrating PostgreSQL Databases to VMware Data Services Manager

- Step 2 – Extract Schema from source database. To perform this step, you must have access to the objects in the database.

Example:

```
pg_dump "host=source-cluster-host user=source-db-user password=source-db-passsword dbname=source-db-name sslmode=prefer"
--schema-only --clean --if-exists --no-owner --no-privilege --no-subscriptions --no-publications > /tmp/output.log
```

In this example, we populated the sample source database named “sourcedb1” using *pgbench* utility with scale factor set to 10000. The sample database size was about 146GB.

```
~$ pgbench -i -s 10000 sourcedb1
dropping old tables...
NOTICE: table "pgbench_accounts" does not exist, skipping
NOTICE: table "pgbench_branches" does not exist, skipping
NOTICE: table "pgbench_history" does not exist, skipping
NOTICE: table "pgbench_tellers" does not exist, skipping
creating tables...
generating data (client-side)...
1000000000 of 1000000000 tuples (100%) done (elapsed 941.42 s, remaining 0.00 s)
vacuuming...
creating primary keys...
done in 1514.07 s (drop tables 0.00 s, create tables 0.00 s, client-side generate 946.15 s, vacuum 1.14 s, primary keys 566.78 s).
```

The sample output.log file may look like this:

```
~$ more /tmp/output.log
--
-- PostgreSQL database dump
--

-- Dumped from database version 16.8 (Ubuntu 16.8-0ubuntu0.24.04.1)
-- Dumped by pg_dump version 16.8 (Ubuntu 16.8-0ubuntu0.24.04.1)

SET statement_timeout = 0;
SET lock_timeout = 0;
SET idle_in_transaction_session_timeout = 0;
SET client_encoding = 'UTF8';
SET standard_conforming_strings = on;
SELECT pg_catalog.set_config('search_path', '', false);
SET check_function_bodies = false;
SET xmloption = content;
SET client_min_messages = warning;
SET row_security = off;

ALTER TABLE IF EXISTS ONLY public.pgbench_tellers DROP CONSTRAINT IF EXISTS pgbench_tellers_pkey;
ALTER TABLE IF EXISTS ONLY public.pgbench_branches DROP CONSTRAINT IF EXISTS pgbench_branches_pkey;
ALTER TABLE IF EXISTS ONLY public.pgbench_accounts DROP CONSTRAINT IF EXISTS pgbench_accounts_pkey;
DROP TABLE IF EXISTS public.pgbench_tellers;
DROP TABLE IF EXISTS public.pgbench_history;
DROP TABLE IF EXISTS public.pgbench_branches;
DROP TABLE IF EXISTS public.pgbench_accounts;
SET default_tablespace = '';
```

- Step 3 – Create the schema objects in the target database

Example:

```
psql connection_string_to_target_dsm_database < /tmp/output.log
```

Sample output:

```
~$ psql postgresql://pgadmin:Password@10.211.198.233:5432/dsmg1 < /tmp/output.log
SET
SET
SET
SET
SET
set_config
-----
(1 row)

SET
SET
SET
SET
ALTER TABLE
ALTER TABLE
ALTER TABLE
DROP TABLE
DROP TABLE
DROP TABLE
DROP TABLE
SET
SET
CREATE TABLE
ALTER TABLE
CREATE TABLE
ALTER TABLE
CREATE TABLE
ALTER TABLE
CREATE TABLE
ALTER TABLE
ALTER TABLE
ALTER TABLE
ALTER TABLE
```

- Step 4 – Create publication on the source database

Example:

```
create publication onboard_to_dsm for all tables;
```

Sample output:

```
sourcedb1=# create publication onboard_to_dsm for all tables;
CREATE PUBLICATION
```

- Step 5 – Create subscription on the target database (DSM-managed)

```
CREATE SUBSCRIPTION onboard_to_db_name CONNECTION 'dbname=source-db host=source-db-host port=port user=source-db-user
password=source-db-password sslmode=prefer' PUBLICATION onboard_to_dsm WITH(create_slot = true, slot_name =
onboard_to_dbname, enabled=true);
```

Sample output:

```
dsmPg1=# CREATE SUBSCRIPTION onboard_to_dsm_target CONNECTION 'dbname=sourcedb1 host=10.211.198.200 port=5432
user=postgres password=Password sslmode=prefer' PUBLICATION onboard_to_dsm WITH(create_slot=true, slot_name=slot1,
enabled=true);
NOTICE: created replication slot "slot1" on publisher
CREATE SUBSCRIPTION
```

- Step 6 – Monitor the logical replication status

Use the `pg_subscription_rel` catalog to view the logical replication status

```
select srsubid, class.relname, srsubstate from pg_subscription_rel join pg_class class on oid = srrelid;
```

Check the `srsubstate`. The sample output below with all “r” indicates the replication has been completed.

- i = initialize
- d = data is being copied
- f = finished table copy
- s = synchronized
- r = ready (normal replication)

```
dsmPg1=# select srsubid, class.relname, srsubstate from pg_subscription_rel join pg_class class on oid = srrelid;
srsubid | relname      | srsubstate
-----+-----+-----
 16457 | pgbench_accounts | r
 16457 | pgbench_branches | r
 16457 | pgbench_history  | r
 16457 | pgbench_tellers  | r
(4 rows)
```

You may also check the log sequence number (LSN) to make sure the logical replication is completed. On the source PostgreSQL database.

```
sourcedb1=# select pg_current_wal_lsn();
pg_current_wal_lsn
-----
25/3AF7E5E0
(1 row)
```

On the target Data Services Manager, check if the LSN matches the source PostgreSQL database.

```
dsmPg1=# select received_lsn, latest_end_lsn from pg_catalog.pg_stat_subscription;
received_lsn | latest_end_lsn
-----+-----
25/3AF7E5E0 | 25/3AF7E5E0
(1 row)
```

- Step 7 – Confirm the necessary data is replicated on the target DSM database
- Step 8 – Switch your application to start using the PostgreSQL database provisioned by Data Services Manager and drop the publication and subscription

On the target DSM database:

```
dsmPg1=# drop subscription if exists onboard_to_dsm_target;
```

```
NOTICE: dropped replication slot "slot1" on publisher
DROP SUBSCRIPTION
```

On the source database:

```
sourcedb1=# drop publication if exists onboard_to_dsm;
DROP PUBLICATION
```

Tested Configuration

Once you have completed all the migration steps, you may configure your applications to start using the PostgreSQL databases provisioned by VMware Data Services Manager.

We measured the migration result of the following test scenario. Both source and target PostgreSQL databases are configured with 2 vCPU and 8GB memory. Note the estimated migration duration is subject to change due to the hardware configuration, network bandwidth, resource allocation and other factors and should be only used for reference purposes.

- In the first scenario, we applied no workload during the migration process. This is a recommended way to migrate with PostgreSQL Logical Replication unless your production workload can not stop due to business requirements. The estimated time for 146GB data migration is about 50 minutes.
- In the second scenario, we applied workload on the source database after the first scenario was completed. The Logical Replication will continue to replicate the transaction workloads on the source database to the target database in Data Services Manager. The estimated migration duration is 50 minutes plus additional 15 minutes to complete the new transactions from the source to target.
- In the third scenario, we applied transaction workload during the migration process at the source side. As both Logical Replication and transaction workload will consume the resource at source and target side at the same time, the estimated total migration duration is 73 minutes. Note this is not a recommended way to migrate your databases to Data Services Manager, but only a simulation of a use case where the source workload can not stop due to certain business reasons.

Logical Replication migration result

Scenario	Description	Estimated Migration Duration
Migration with no workload on the source database	Migrate sample source PostgreSQL with about 146GB data	~50 minutes
Workload applied on the source database after migration completed	Migrate sample source PostgreSQL with about 146GB data, then apply pgbench workload with 100000 transactions	50 + 15 minutes
Workload applied on the source database during migration	Migrate sample source PostgreSQL with about 146GB data, and apply pgbench workload with 100000 transactions during the migration process	73 minutes

Use Case 2 – Migration with Backup and Restore Methodology

The backup and restore methodology refers to using *pg_dump*, *pg_restore* and *pg_dumpall* utilities to export the desired source databases and metadata, to the target databases and objects managed by Data Services Manager. The *pg_dump* utility dumps in plain-text files containing the SQL commands required to reconstruct the database to the state it was in at the time it was saved. The restore can be completed by feeding the dump files to *psql* utility or *pg_restore* utility to rebuild the database at the target side.

Migrating PostgreSQL Databases to VMware Data Services Manager

Here is the difference between *pg_dump* and *pg_dumpall*.

- *pg_dump* only dumps a single database.
- *pg_dump* supports flexible output files with a custom format which is compressed by default.
- *pg_dumpall* backup the entire database cluster including all databases.
- *pg_dumpall* also backs up the global objects that are common to a cluster, such as database roles and tablespaces.

Migration Prerequisites

The *pg_dump* command requires a staging location for the output files, in a format that is flexible to restore the target database into a Data Services Manager deployed PostgreSQL database. Read the following prerequisites before migrating with the backup and restore methodology

- Verify that you have the necessary CLI tools installed in your source environment including the *pg_dump*, *pg_restore*, *psql* and *pg_dumpall* utilities.
- Obtain a list of databases under migration. The *pg_dump* command can export with custom output with compression enabled by default, but can only migrate one database at a time, while *pg_dumpall* is more useful to migrate multiple databases or the entire database cluster within a batch.
- Use *pg_dumpall* with the “*-schema-only*” option to migrate global database objects including roles, tablespaces etc.
- *pg_dumpall* needs to connect to the database with **superuser privileges** to read tables from all databases on the source side and execute the saved script on the target side. For more details, read [pg_dumpall requirements](#).

Migration Steps (pg_dump to migrate single database)

- Step 1 – On the source database, use **pg_dump** to create the dump files. You may specify “-Fc” to allow custom format for the dump file.

```
pg_dump -h source-host -p port-number -U user-name -Fc -b -v -f dump-file-pwd -d database-name
```

Example:

```
~$ pg_dump -h localhost -p 5432 -U postgres -Fc -b -v -f dumpsourcedb.sql -d sourcedb1
pg_dump: last built-in OID is 16383
pg_dump: reading extensions
pg_dump: identifying extension members
pg_dump: reading schemas
...
...
pg_dump: dumping contents of table "public.pgbench_accounts"
pg_dump: dumping contents of table "public.pgbench_branches"
pg_dump: dumping contents of table "public.pgbench_history"
pg_dump: dumping contents of table "public.pgbench_tellers"
```

- Step 2 – Restore the dump files to the target database in Data Services Manager

```
pg_restore -v -h target-host -U user-name -d database-name -j number-of-jobs dump-file-pwd
```

Example:

```
~$ pg_restore -v -h 10.211.198.233 -U pgadmin -d dsmpg1 -j 2 dumpsourcedb.sql
pg_restore: connecting to database for restore
```

```
Password:
pg_restore: processing item 3416 ENCODING ENCODING
pg_restore: processing item 3417 STDSTRINGS STDSTRINGS
pg_restore: processing item 3418 SEARCHPATH SEARCHPATH
pg_restore: processing item 3419 DATABASE sourcedb1
pg_restore: processing item 217 TABLE pgbench_accounts
pg_restore: creating TABLE "public.pgbench_accounts"
...
...
pg_restore: finished item 3412 TABLE DATA pgbench_accounts
pg_restore: launching item 3263 CONSTRAINT pgbench_accounts pgbench_accounts_pkey
pg_restore: creating CONSTRAINT "public.pgbench_accounts pgbench_accounts_pkey"
pg_restore: finished item 3263 CONSTRAINT pgbench_accounts pgbench_accounts_pkey
pg_restore: finished main parallel loop
```

- Step 3 – On the source database, export the global objects using *pg_dumpall*

```
pg_dumpall -U user-name -h source-host -f dump-global-pwd --no-role-passwords -g
```

Example:

```
~$ pg_dumpall -U postgres -h localhost -f dumpglobal.sql --no-role-passwords -g
```

- Step 4 – Import the global object into the database provisioned by Data Services Manager

```
psql -h target-host -U user-name -f dump-global-pwd
```

Example:

```
~$ psql -h 10.211.198.233 -U pgadmin -f dumpglobal.sql dsmpg1
Password for user pgadmin:
SET
SET
SET
ALTER ROLE
ALTER ROLE
ALTER ROLE
ALTER ROLE
```

Migration Steps (*pg_dumpall* to migrate multiple databases at the source)

- Step 1 – Get a list of databases on the source side to migrate to Data Services Manager.

In this example, we populated the two sample source databases named “sourcedb1” and “sourcedb2” using the *pgbench* utility with scale factor set to 5000. Each sample database is 73GB.

```
postgres-# \l+
                                List of databases
 Name | Owner  | Encoding | ... | Size | Tablespace | Description
-----+-----+-----+---+-----+-----+-----
 bucardo | postgres | UTF8   | ... | 8636 kB | pg_default |
 postgres | postgres | UTF8   | ... | 7524 kB | pg_default | default administrative connection database
 sourcedb1 | postgres | UTF8   | ... | 73 GB  | pg_default |
```



```
sourcedb2 | postgres | UTF8 | ... | 73 GB | pg_default |  
template0 | postgres | UTF8 | ... | 7345 kB | pg_default | unmodifiable empty database  
template1 | postgres | UTF8 | ... | 7564 kB | pg_default | default template for new databases  
  
(6 rows)
```

- Step 2 – Use `pg_dumpall` on the source host to create the dump files of multiple databases under migration

```
pg_dumpall -U user-name -f dump-file-pwd --exclude-database=PATTERN
```

Example:

```
~$ pg_dumpall -U postgres -f dumpall.sql --exclude-database=bucardo --exclude-  
database=postgres --exclude-database=template0 --exclude-database=template1
```

- Step 3 – Reload the multiple databases into a Data Services Manager-managed database at the target side

```
psql -h target-host -U user-name -p port-number -f dump-file-pwd target-connection-  
database
```

Example:

The target database instance has been created by Data Services Manager. In this example, it is called “dsmpg1”. The command recreates the two databases from the source side (named “sourcedb1” and “sourcedb2”) on the target side.

```
psql -h 10.211.198.233 -U pgadmin -p 5432 -f dumpall.sql dsmpg1  
  
...  
  
CREATE DATABASE  
ALTER DATABASE  
psql (16.9 (Ubuntu 16.9-0ubuntu0.24.04.1), server 16.8 (VMware Postgres 16.8.1))  
SSL connection (protocol: TLSv1.3, cipher: TLS_AES_256_GCM_SHA384, compression: off)  
You are now connected to database "sourcedb1" as user "pgadmin".  
  
...  
  
CREATE DATABASE  
ALTER DATABASE  
psql (16.9 (Ubuntu 16.9-0ubuntu0.24.04.1), server 16.8 (VMware Postgres 16.8.1))  
SSL connection (protocol: TLSv1.3, cipher: TLS_AES_256_GCM_SHA384, compression: off)  
You are now connected to database "sourcedb2" as user "pgadmin".  
SET  
SET  
SET  
...
```

- Step 4 – Verify the multiple databases in the target database instance provisioned by Data Services Manager. You can see the multiple databases “sourcedb1” and “sourcedb2” have been successfully migrated to the target database instance.

```
dsmpg1=# \l+
```

List of databases					
Name	Owner	...	Size	Tablespace	Description
-----+-----+-----+-----+-----					
dsmgp1	pgautofailover_replicator	...	7684 kB	pg_default	
postgres	postgres	...	7724 kB	pg_default	default administrative connection database
sourcedb1	postgres	...	73 GB	pg_default	
sourcedb2	postgres	...	73 GB	pg_default	
template0	postgres	...	7505 kB	pg_default	unmodifiable empty database
template1	postgres	...	7772 kB	pg_default	default template for new databases

(6 rows)

Tested Configuration

Once the migration steps are completed, configure your applications to start using the PostgreSQL databases provisioned by VMware Data Services Manager.

The backup and restore methodology using *pg_dump* or *pg_dumpall* requires certain business downtime during the migration process, compared to migration using Logical Replication. Here's the tested configuration running both source and target databases on vSAN and we measured the test result. Both source and target PostgreSQL databases are configured with 2 vCPU and 8GB memory. Note the estimated migration duration is subject to change due to the hardware configuration, network bandwidth, resource allocation and other factors and should be only used for reference purposes.

From the test result we can see the same compute/network/storage resource allocation, backup and restore methodology may take more time than migration using Logical Replication. However *pg_dump* and *pg_dumpall* can help migrate both database and global objects, and support database DDL migration (e.g. schema/tablespace creation etc.) to the target side, which can not be achieved by Logical Replication only.

Migration with backup and restore test result (pg_dump)

Scenario	Description	Estimated Migration Duration
Backup using <i>pg_dump</i>	Backup sample source PostgreSQL with about 146GB data, with “-Fc” option (compression) enabled	~15 minutes
Restore using <i>pg_restore</i>	Restore the sample database to target database in Data Services Manager	~50 minutes
Backup and Restore global objects	Export and import global objects using <i>pg_dumpall</i>	Completed at once

Migration with backup and restore test result (pg_dumpall)

Scenario	Description	Estimated Migration Duration
Backup using <i>pg_dumpall</i>	Backup two source PostgreSQL database with total 146GB data. Note this is with no compression for dump files	~12 minutes
Restore databases in Data Services Manager	Restore the two sample databases to target database in Data Services Manager	~50 minutes

Use Case 3 – Migration with Bucardo

Bucardo is an open-source asynchronous PostgreSQL replication system that allows for multi-source, multi-target database migrations. Bucardo is a Perl daemon that listens for notify requests and acts on them, by connecting to target databases and copying data back and forth. All the specific information that the daemon needs is stored in the main bucardo database, including a list of all the databases involved in the replication and how to reach them, all the tables that are to be replicated, and how each is to be replicated. The replication operation occurs at the table level which is achieved by database triggers.

We recommend using Bucardo for PostgreSQL migrations to Data Services Manager, if the source database is a legacy version (prior to 10.x) of PostgreSQL and there is a requirement for minimal downtime.

Here are the limitations of migration using Bucardo.

- Bucardo migrates PostgreSQL databases with the granularity of table, which does not support migrating a table without a primary key.
- Bucardo does not migrate the DDL operations. Use *pg_dump -schema-only* to migrate the schema first from the source to the target database.
- Bucardo requires a central database which can either be deployed at source PostgreSQL or on a PostgreSQL on a third location. You can not deploy the central database on the PostgreSQL provisioned by the Data Services Manager.
- The initial table migration from the source database to the target (Data Services Manager-managed) database can be achieved via [onetimecopy](#) option available in Bucardo. This setting overrides the default pushdelta mode to use fullcopy mode.

Here are the high-level steps to migrate a standalone PostgreSQL database to a Data Services Manager-managed PostgreSQL database using Bucardo. For more details refer to the [Bucardo documentation](#).

1. Configure the Bucardo dependencies and requirements and install Bucardo.
2. Deploy the target database using Data Services Manager
3. Migrate the schema from the source database to the target database using *pg_dump*
4. Add the source and target PostgreSQL databases to Bucardo.
5. Add the tables from the source to replicate the target with Bucardo.
6. Create Bucardo sync to start the replication.
7. Logon to PostgreSQL in Data Services Manager and monitor the migration progress.

Here are the tested configuration of migration with Bucardo

Bucardo migration result		
Scenario	Description	Estimated Migration Duration
Migrate PostgreSQL database to DSM using Bucardo	Source PostgreSQL with about 146GB data. Both source database and target database VM is configured with 2vCPU and 8GB memory	~50 minutes

Best Practices and Key Considerations

- It is highly recommended to perform the entire migration steps with a test/dev database before migration of the actual production databases. This will help identify potential issues and resolve them beforehand.
- Migrate with logical replication
 - Minimize the business downtime of the source PostgreSQL databases.
 - Replicate the subsequent changes during the migration period of your source databases.
 - Require source PostgreSQL database with version 10.0 or later.
- Migrate with backup and restore methodology (*pg_dump* and *pg_restore*)
 - Simple to configure and can be considered for small-to-medium sized source PostgreSQL database environments (requires storage for backup output files).
 - Support to migrate PostgreSQL objects including view, triggers, stored procedures etc.
 - Support to migrate PostgreSQL global objects including database roles, tablespaces, and privilege grants for configuration parameters using *pg_dumpall* utility.
 - *pg_dump* migrates single database at one time with compression for dump files by default
 - *pg_dumpall* can migrate multiple databases or the entire database instance within a batch
 - Make sure the source database has no workload / transaction during the migration process (extended down time required). *pg_dump* does not capture ongoing changes.
- Migrate with open-source tool - Bucardo - for legacy versions of PostgreSQL databases
 - Bucardo can support multi-source / multi-target migration.
 - Read the [limitations](#) of the Bucardo tool before migrating the databases.
 - If you have not used this tool previously, be sure to install it in a test-dev environment and familiarize yourself with its use and capabilities.
 - Bucardo is an open source tool maintained by the community. Check the [Bucardo page](#) for updates, support and for further open-source tool related help.
- Data Services Manager supports the following deployment options for PostgreSQL database high availability. You may choose the deployment type for the target PostgreSQL database based on your business SLAs.
 - Single server deployment - standalone option for PostgreSQL databases provided by Data Services Manager.
 - Cluster deployment - You can deploy PostgreSQL databases with primary and replica across single or multiple vSphere clusters. It provides automatic failover and resiliency to a single compute node failure.
- VM Class
 - Data Services Manager maintains a database virtual machine template with an optimized configuration so that the infrastructure administrators only need to focus on the sizing of the virtual machine class for the database.
 - Data Services Manager supports customized virtual machine classes for PostgreSQL databases. It is recommended to start with the same sizing as of the source database configuration in terms of CPU cores and memory and you can easily scale-up or scale-down for different use cases.
- IP Pools
 - Assign sufficient IP addresses to each of the IP Pools for data services provisioned by Data Services Manager. Single server deployments require one IP address for each node, as well as two additional cluster-level IP addresses for the database and control plane access. Three-node cluster deployments require 5 IP addresses.
 - Maintenance operations such as database engine upgrades or changes to cluster topology require additional IP addresses. For planning purposes, you need to calculate the size of the pool based on a higher number of servers and clusters.
- Compute Resources
 - Make sure to enable vSphere HA and vSphere DRS for the cluster used to deploy data services by Data Services Manager. For more information, refer to [How vSphere HA Works](#) in vSphere Availability and [Create a vSphere DRS Cluster](#) in vSphere Resource Management.
- Storage Policies

- For vSAN use FTT=1 or FTT=2 based on the protection level required for the data services. vSAN Express Storage Architecture provides RAID-5 for FTT=1 and RAID-6 for FTT=2 with optimal performance and capacity. For more information, refer to [Using RAID 5 or RAID 6 Erasure Coding in vSAN Cluster](#)
- For external storage, create a host-based or tag-based storage policy for Data Services Manager. See [Creating and Managing vSphere VM Storage Policies](#).



Copyright © 2024 Broadcom. All rights reserved.

The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. For more information, go to www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Item No: vmw-bc-wp-tech-temp-a4-word-2024 1/24