

TECHNICAL FEATURES

Integrity

100% Physical replication

The standby database is an exact replica of the primary at the lowest (binary) level to ensure data integrity. All indexes, pointers, and tables are transferred to ensure database consistency.

Graceful switchover

A simple automated process for performing a planned switchover (role change) between primary and standby database environments with zero-data-loss.

Pre-flight checks

Automated pre-install checks ensure the standby server meets technical prerequisites before standby creation. This saves the admin time by avoiding rework.

Real-time monitoring

Continual monitoring of the time gap, tasks, and issues, combined with smart notifications.

Configurable lag time

A “lag” (delay) can be configured to keep the standby “X” amount of time behind the primary database. This can help guard against human error, as changes to the standby database can be stopped during the lag time.

Cascading standby database support

StandbyMP supports the creation of multiple cascading standby databases from a single primary.

Network encryption

All data sent over the network is encrypted by default.

Speed

Low overhead architecture

StandbyMP requires few server resources, has low latency requirements, and is data efficient.

Network compression

Logs are compressed during transfer providing significant savings on bandwidth requirements.

Script post processing

Shell or batch scripts can be configured to perform other tasks before or after Standby activation. Configure these post-processing tasks for advanced options like Graceful Switchover or activation.

Backups from standby

Utilize your standby for the creation of backups.

Hot or warm standby database for fast recovery

The standby database is always on and can take over in just a few minutes by a simple command or the automated Observer.

Automatic failover (or manual)

The Observer monitors the status of both the primary and standby databases. If any errors are detected, a notification is sent, and the system can perform an automatic failover based on pre-defined rules.

Continuously updating (2 min RPO)

Depending on the Replication Mode used, the standby is either kept up-to-date in real-time (data streaming), or archive logs are continuously transferred and applied to the standby, ensuring typical maximum data loss of less than 2 minutes, but can also be configured to user needs.

Network compression

Logs are compressed during transfer providing significant savings on bandwidth requirements.

Reporting from standby (read only)

Use a Hot Standby database for read-only queries to boost ROI and reduce load on the primary system. Recovery can still proceed while the Hot Standby is queried.

Clarity

One intuitive UI

A consistent browser-based UI to create, view and manage all your Disaster Recovery configurations.

Guided user experience on GUI

Powerful yet simple, the GUI enables administrators to perform tasks quickly, easily, and with confidence. Lower barriers to use reduce dependency on key employees.

Effortless standby creation

Creation of one or more standby clusters in a fast, streamlined process.

Smart notifications (log gap, heartbeat, status)

Use email and Slack to notify admins of status and issues.

Advanced task tracking

Tasks in the Standby Central Console offer extensive information about all initiated processes and database events.

User management and access controls

Track and limit user access to the product based on a set of predefined Roles.

Standardized DR

PostgreSQL has different workflows & utilities depending on version, replication method, and underlying operating system. StandbyMP adapts to any environment, to standardize DR processes.

Best-in-class UI

Building on Dbvisit's 15 years of experience in DR, StandbyMP offers the first truly easy to use UI for best-practice PostgreSQL DR management.

Automation

Failover assistant (automated or guided)

Following near instantaneous issue detection by the Observer, failover begins automatically, or manually after notification to the administrator.

Archive log management

Automatic management of the archive log files on both the primary and standby systems.

Automated cluster & service management

During standby creation & switchover, StandbyMP will reconfigure PostgreSQL clusters and services as necessary to ensure smooth operation.

Support package creation

The support package collects detailed data automatically for fast issue resolution together with Dbvisit support.

Compatibility

Version support

StandbyMP supports PostgreSQL versions 12+, across Ubuntu, CentOS and Windows.

Cloud Ready

Supports cloud-based solutions or hybrid solutions for Oracle Cloud, Microsoft Azure, and AWS, where the primary database is run on-site, and the standby database is running in a hosted environment, or the complete configuration is in the Cloud.