DBVISIT

Oracle Standard Edition

TECHNICAL FEATURES

Database Continuity

Continually verified database
The standby database is continuously self-verified, guaranteeing database integrity and successful failover at any time.

100% Physical replication
The standby database is an exact replica of the primary at the lowest (binary) level which ensures data integrity.

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

DR testing
One-click automated DR testing.

Pre-flight checks
Automated pre-install checks ensure the standby server meets technical prerequisites before standby creation.

Continuously updating
Archive logs are continuously transferred and applied to the standby, ensuring typical maximum data loss of 5 minutes, but can also be configured to suit user needs.

The Zero Data Loss option, creates an additional streaming replication layer for sub-second data loss in most environments.

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.

Point-in-time recovery
Recover to a precise SCN or timestamp by utilising data automatically synchronised to the standby. The 'Intelligent Activate' feature guarantees successful activation.

'Intelligent Activate' provides recoverable SCN or timestamps to recover to, using data automatically synchronised to the standby.

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

Network encryption
Network encryption is enabled by default for all data sent over the network.

Ease of Use

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

Guided user experience in the 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 databases in a fast, streamlined process. Standby requires few changes to your existing environment and supports ASM and non-ASM file systems.

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.

Advanced Features

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

Warm standby database for fast recovery
The standby database is warm, operational and can take control in just a few minutes with a simple command or via 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.

Proven archive log replication (5 min RPO)
Archive logs are continuously transferred and applied to the standby, ensuring typical maximum data loss of 5 minutes, but can also be configured to suit user needs.

Optional Zero Data Loss (<1s RPO)
ZDL adds a second layer of protection, continuously monitoring the online redo logs and streaming changes to the standby as they happen for sub-second data loss in most environments.

Low-overhead architecture
StandbyMP requires few server resources, has low latency requirements, and is data efficient.

Reporting and test/dev with Snapshots
StandbyMP's Snapshot feature quickly creates point-in-time or regularly updated environments for reporting, DR testing, and test/dev. Available on Linux only.

Reporting from standby (read-only)
Use the standby for read-only queries to boost ROI and reduce load on the primary system. The archive logs are still sent to the standby database, and recovery can proceed once it is placed back into Recovery Mode.

Backups from standby
Utilize your standby for the creation of backups.

Transportable media option
Create a standby database using a local storage device when the database size is prohibitively large or the network is unstable.

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

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

Automation

Failover assistant (automated or guided)
Following near instantaneous issue detection by the Observer, failover begins automatically, or manually after notification to the administrator.

MultiPlatform API & CLI
A powerful new command-line syntax, unified across all database types, together with a fully-featured API, allow for standardised, streamlined integrations and automation across complex workflows.

Support package creation
The support package automatically collects detailed data to help Dbvisit Support resolve issues quickly.

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

Compatibility

Full Oracle technology support
Full support of Oracle features such as ODA, ODA KVM, RAC, SE2HA, ASM, ACFS, and Oracle Managed Files (OMF).

Oracle Multitenant support
Automatically transfer and synchronise up to three Pluggable Databases (PDBs) to your standby environment on one StandbyMP license.

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.