19c On Oracle SE Upgrade

HA, DR, and Backup Compared

Introduction

Backups, High Availability and Disaster Recovery (Standby Databases) are complementary technologies that may be implemented as part of your business continuity plan based on the importance of a particular database.

For Oracle databases these technologies are key components within Oracle's Bronze, Silver and Gold Maximum Availability architectures. The recommended technologies and architecture depend on the role and importance of the database within your organisation, where a development database would have significantly lower requirements than a business-critical database.

Backups, High Availability and Disaster Recovery can mainly be evaluated based on their RPO and RTO performance, resilience to different types of failures, and cost. First, we will cover these classifications, then look at each of the technologies, followed by a summary of the technologies.

Evaluation Criteria

Recovery Point Objective (RPO)

RPO is the point in time which your systems will be restored from when brought back online. For example, daily backup has a recovery point at worst that is 24hrs behind the production database (RPO = 24hrs).

Your organisation's RPO requirement should be defined as the maximum tolerable period in which data might be lost from a disaster. Learn more in Dbvisit's RTO and RPO article.

Recovery Time Objective (RTO)

RTO is the time required for IT to bring your systems back online into production. If you have an RTO of 2 days, and RPO of 1 day and you fail at Tuesday noon, your systems will be back on Thursday noon using Monday noon's data.

Your RTO requirement can be defined as the time in which a business process must be restored after a disaster to avoid unacceptable consequences associated with a break in business continuity. Learn more in Dbvisit's RTO and RPO article.

Coverage & Resilience

Your disaster recovery plan should cover all types of disaster, including operational failures, natural disaster, and human-caused events. Forrester Research in 2014 published an excellent document showing how unplanned downtime causes were spread over a variety of causes.

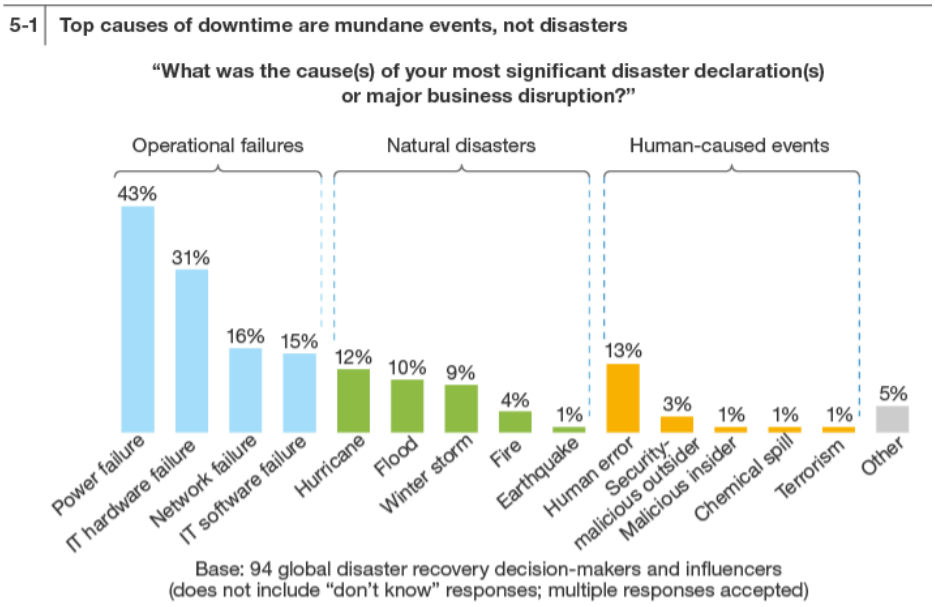

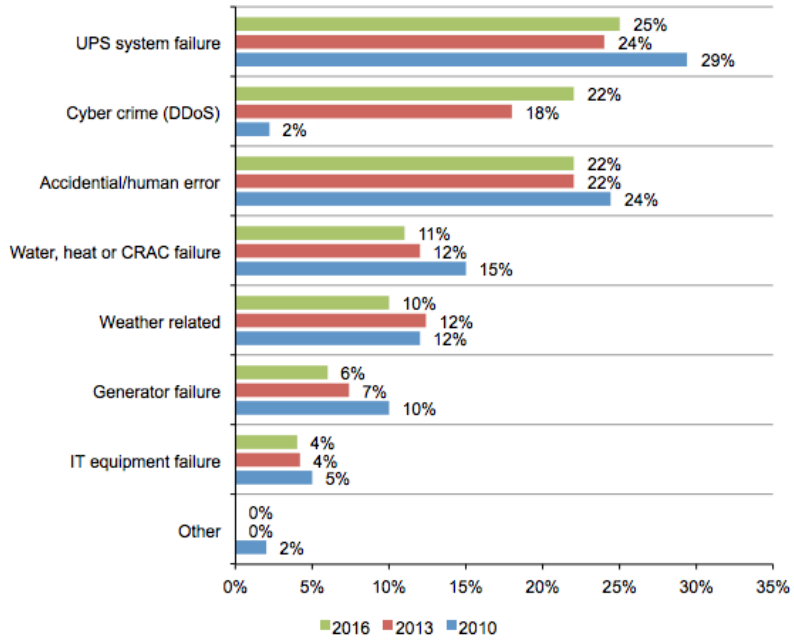

The above diagram by Ponemon Institute shows a significant spread of causes for unplanned downtime, with many of the failures affecting sitewide capabilities. Average cost per outage in 2016 was $740,357 (up from $690,204) in 2013.

Review of key Terminologies

Backup (RMAN on Oracle SE):

Backups performed by RMAN are a copy of your organisation's systems and data that can be used to bring a failed system back online to a particular point in time.

- **Data Completeness (RPO) - Poor:**
  - They are to the time when the backup was completed, such as every 24hrs, and likely will not include latest data.
  - Incremental backups backup the datafile blocks that have changed since the last backup. If a failure occurs the incremental backup is applied to a restored (full) backup reducing the RPO (data lost).
- **Restoration (RTO) - Slow:**
  - Backups can take significant time (many days) to utilise, restore and put into production (example RPO = 24hrs, RTO = 3 days).
  - Furthermore, applying incremental backups and archive log files to the restored backup to reduce lost data (RPO) adds further to the RTO.
  - Do not usually include infrastructure. The data may be available, but the production hardware may be damaged. Provisioning hardware can take additional time, adding further to the RTPO.
  - Testing of the database to confirm it is error free is required.
- **Testing (Integrity) - Difficult:**
  - It's difficult to test backups in a production environment with production volumes. Typically, a “test” machine is therefore provisioned for the purpose of doing a recovery test. The work required makes this an infrequent activity, increasing risk.
- **Resilience - Good**
  - Multiple copies of the data can be made, with more than 3 recommended.
  - Different storage types can be used to minimise risk.
  - Geographically remote sites can be used to protect against disaster.

High Availability (SE2HA for 19c on Oracle SE)

High Availability (HA) on Oracle has multiple servers each running Oracle RDBMS software while accessing a single shared database to create a clustered database. Standard Edition High Availability (SE2HA) available on Oracle SE has two nodes in an Active/Passive arrangement with shared storage. If the active node has a minor hardware failure the database is relocated to the passive node and starts up. Applications are then redirected to the new node keeping the database available to the application and users with only a few minutes downtime. As the storage is shared this presents a single-point-of-failure risk and SE2HA requires a backup or DR solution to mitigate these risks.

- **Data Completeness (RPO) - Excellent:**
  - Shared storage enables failover to the passive node without any data loss when a node fails.
- **Restoration (RTO) - Excellent**:
  - The passive node requires only a few minutes to start up. Once online, applications are automatically redirected to the new active node.
- **Testing (Integrity) - Excellent:**
  - The nodes use shared storage so there is no requirement for testing.
- **Resilience - Poor:**
  - Shared storage creates a single point of failure (SPOF). This can be partially mitigated using redundant storage options, such as a SAN.
  - Storage is vulnerable to user error and corruption of the shared storage.
  - Both nodes and storage are located within the same facility increasing risk related to data center failure
- **Summary:**
  - SE2HA provides zero data loss and high database availability to applications and users in the event of minor hard failures. However, as the platform has single points of failure and the nodes are not geographically dispersed it is not a DR solution and can only be part of a balanced DR plan.

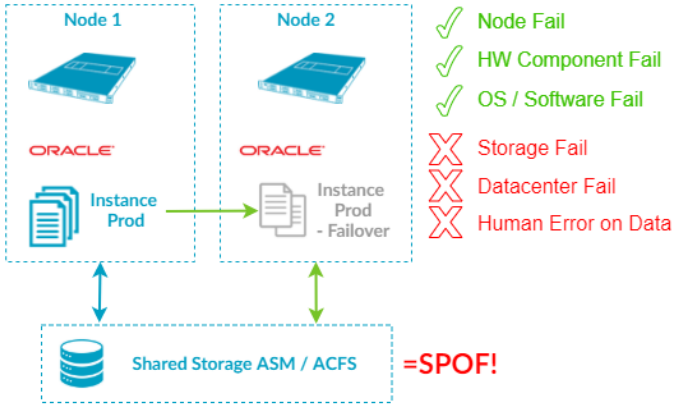

Disaster Recovery Solution (Dbvisit Standby for Oracle SE)

Disaster Recovery (DR) is a more advanced form of system copies that includes processing capabilities and has minimal data gap between the primary and the copy (standby). Oracle Enterprise Edition includes the physical replication software, Oracle Data Guard, that creates and manages a standby database for Disaster Recovery. On Oracle SE a standby database can be created and managed by the proven Dbvisit Standby software using physical replication technology.

- **Data Completeness (RPO) - Excellent:**
  - A continuous stream of archive logs are sent from the source database to the standby site creating a minimal gap between the primary and standby.
  - RPO of 10 minutes
- **Restoration (RTO) - Good**:
  - DR includes standby infrastructure where the standby database is always running ready for a failover or graceful switchover event (switchover with zero data loss).
  - Failover to the standby site can be performed in a few minutes. The only work is to point the application to the standby site.
- **Testing (Integrity) - Excellent:**
  - Dbvisit Standby continually checks the integrity of the standby database.
  - Testing of the standby database can be easily carried out.
- **Resilience - Excellent:**
  - No single-point-of failures
  - Creating geographic distance between the primary and standby is easy. The DR could also be done to the cloud, such as AWS, Azure or Oracle Cloud.
  - As a bonus, the standby server can also be used to run your backups and reduce overhead on your primary server by providing read-only query access.
- **Summary:**
  - Dbvisit Standby provides a full-featured DR platform enabling rapid failover in the event of any type of disaster or operational failure. Because the standby system can be geographically separate from the primary environment and can be continually updated, the RTO and RPO time frames provide for rapid recovery with minimal data loss.

Comparing Backup, HA and DR on Oracle SE

Simplified Performance Comparison

Backup, High Availability and Disaster Recovery for Oracle SE each were developed with a different focus and deliver different capabilities. A high-level overview of these capabilities is provided in the table below.

	Backup	HA (SE2HA)	DR (Dbvisit Standby)
Continuous Data Flow	No	Yes	Yes
Zero data loss	No	Yes	No (~10 min)
Fast switchover	No (full restore)	Yes	Yes (few min.)
Highly resilient	Yes	No (shared storage)	Yes
Geographically Dispersed	Yes	No (same facility)	Yes
Continuously tested	No	N/A	Yes
DR testing environment	No (difficult)	N/A	Yes
RPO	~24hrs	0 minutes	10 minutes
RTO	~Days	Few minutes	Few minutes plus application layer

Table Explanation:

Backup and Restore delivers a lower cost solution that when well configured can be resilient and cover most disasters. However, because organisations are increasingly data-driven and reliant on data integrity, the periodic nature of backups means they no-longer meet many organisations RPO requirements. Furthermore, as backups can require lengthy restoration, they may also not meet RTO requirements for business continuity.

High Availability using Standard Edition High Availability (SE2HA) prioritises application availability and switchover speed over data integrity and resilience. The focus of SE2HA is to provide minimal disruptions to business from more common server outages at a reasonable price to the organisation. However, due to shared storage requirements they do not provide protection against storage issues including human error, serious hardware failures, data center outages, or natural disasters. For this reason, SE2HA (or RAC) cannot be considered complete Disaster Recovery. HA should be supplemented by Backup and Restore where RTO and RPO requirements are lower, or Disaster Recovery using physical replication where RTO and RPO requirements are higher.

Disaster Recovery using physical replication such as Dbvisit Standby is designed to protect and maintain database integrity through minimal data loss, and continual testing and monitoring of the database and standby environment, and remote geographic location. In this regard Disaster Recovery prioritises database integrity and solution resilience over application availability. For this reason, RPO and RTO are slightly higher than High Availability but can be depended on in nearly any event.

Comparison of RPO and RTO

To provide a very simplified visual comparison of the Backup, High Availability and Disaster Recovery we have plotted them on a 2-axis graph. Please do not be misled by the graph - it does NOT show medium-end DR, and high-end DR. A Standby DR solution is not usually used for brief outages and normal maintenance. Similarly, HA on Oracle SE is not a complete Disaster Recovery solution and would not help in recovering from a fire or event where both nodes are unavailable.

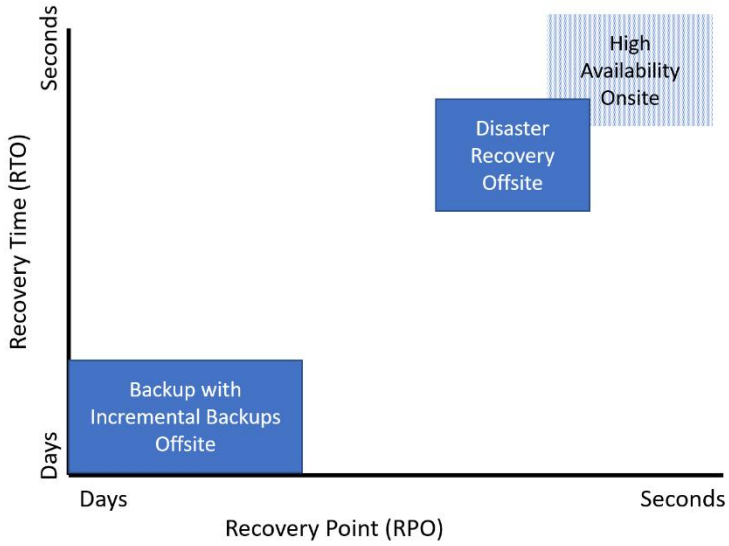

At the bottom left, backups have many hours of data loss if a failure occurs, even when using incremental backups. Backups also take significant time to restore. This can be further lengthened if archive logs must be applied, or the server hardware is non-operational.

High availability at the top right delivers zero data loss and fast recovery (typically a few minutes on SE2HA as the passive server starts). However, as indicated by the light blue color, the solution does not cover many failure types due to its geographic proximity to the primary server and shared storage creating a single point of failure. For this reason, it should be augmented by other technologies as part of a business continuity plan.

Disaster Recovery, such as Dbvisit Standby for Oracle SE, delivers minimal data loss (maximum of ~10min) and fast database failover (recovery) in a few minutes. Integrity of the standby database is maintained through continual exercising and testing. Standby site resilience can be enhanced through creation in a remote location, implementation of a DR test plan and even implementation of a cascaded standby database at a set time delay.

Comparison using Oracle's Maximum Availability Architecture.

The database's use can be an indicator to its importance and RPO and RTO requirements. Below is an adaption of Oracle's Maximum Availability Architectures with technologies available for Oracle SE. The goal of the table is to assist Oracle SE customers in the selection of data protection technologies based on their database use case.

In comparison to Oracle's Maximum Availability Architectures the Platinum level for 'Mission Critical' databases has been removed as this is an Enterprise Level architecture and not applicable to Oracle SE. Secondly, we have created a Silver (B) architecture that prioritizes data integrity and database protection over the application availability prioritised by Silver (A).

Adaption of Oracles Maximum Availability Architecture for Comparison Purposes

Bronze	Silver (A)	Silver (B)	Gold
Dev, test, production	Prod/departmental	Prod/departmental	Business Critical
Backup & Restore	Backup & Restore + High Availability	Backup & Restore + Disaster Recovery	Backup & Restore + High Availability + Disaster Recovery
RPO = ~24hrs RTO = days Resilience = High	RTO = 0 RPO = Few min. Resilience = Medium	RTO = 10 min. RPO = Few min. Resilience = High	RTO = 0~10 min. RPO = Few min. Resilience = High

Summary

This document has looked at and compared Backup, High Availability and Disaster Recovery technologies and their characteristics for Oracle SE.

Using RPO (tolerable data loss), RTO (tolerable recovery time), and Resilience (coverage of different disaster types) it has been shown that each technology was designed for different goals and therefore offer significantly different capabilities. Backups prioritise cost and resilience over RPO and RTO. High Availability prioritises application availability over resilience. And Disaster Recovery prioritises Data Integrity and Resilience for some reduction in RPO and RTO.

The solution appropriate for your organisation will depend on the importance the database has (and will have) to your organisation. This importance can be expressed in your RPO and RTO requirements, which should be evaluated for all disaster types including human error, hardware failure, power failure, fire, natural disaster, and so on.

As an organisation focused on Data Integrity, we believe Disaster Recovery capabilities (in the case of Oracle SE, Dbvisit Standby) is a necessity and should be prioritised over High Availability as data integrity and system resilience are of greater importance than application availability. An organisation with Disaster Recovery will likely continue relatively unaffected by a ~10min of data loss and ~10min of downtime. In comparison we believe an organisation with High Availability that is affected by a data center fire that affects both SE2HA nodes will have significant reputational damage from ~1 day of data loss (last backup) and ~2 days of downtime, even if the event is of lower probability.

Finally, for organisations wanting to deliver high levels of availability across many disaster types, solutions need a combination of high availability (HA) and disaster recovery (DR) and backups. In this architecture HA ensures that the system continues to operate after the failure of individual components, and DR ensures continued availability in the event of a “disaster”, or complete failure of the production site.