I) Oracle Database

- Single-instance
- Non cluster Database

Although this architecture does not have the node or database redundancy, there are a numerous high availability features that can be used in this architecture and any subsequent database architectures.
These features make the standalone database on a single computer attractive and available for certain failures and planned maintenance activities.

1) **Fast-Start Fault Recovery**
   Bounds and optimizes instance and database recovery times (Backup and Recovery).

2) **Automatic Storage Management (ASM)**
   Tolerates storage failures and optimizes storage performance and usage.

3) **Oracle Flashback Technology**
   Optimizes logical failure repair. Oracle recommends that you use automatic undo management with sufficient space to attain your desired undo retention guarantee, enable Flashback Database and allocate sufficient space and I/O bandwidth in the flash recovery area.

4) **Recovery Manager (RMAN)**
   RMAN optimizes local repair of data failures. Oracle recommends that you create and store the local backups in the flash recovery area.

5) **Flash Recovery Area**
   Manages local recovery related files.

6) **Online Reorganization and Redefinition**
   Allows for dynamic data changes.

7) **Oracle Security Features**
   Prevent unauthorized access and changes.

8) **Data Recovery Advisor**
   Provides intelligent advice and repair of different data failures.

9) **Data Block Corruption Prevention and Detection Parameters**
   Detects and prevents some corruption and lost writes.

10) **Dynamic Resource Provisioning**
    Allow for dynamic system changes

11) **Online Patching**
    Allows for dynamic system changes

12) **Oracle Secure Backup**
    Provide a centralized tape backup management solution

Shows a basic, single-node Oracle Database that includes an ASM instance. This architecture takes advantage of several high availability features, including Flashback Database, Online Redefinition, Recovery Manager, and Oracle Secure backup.

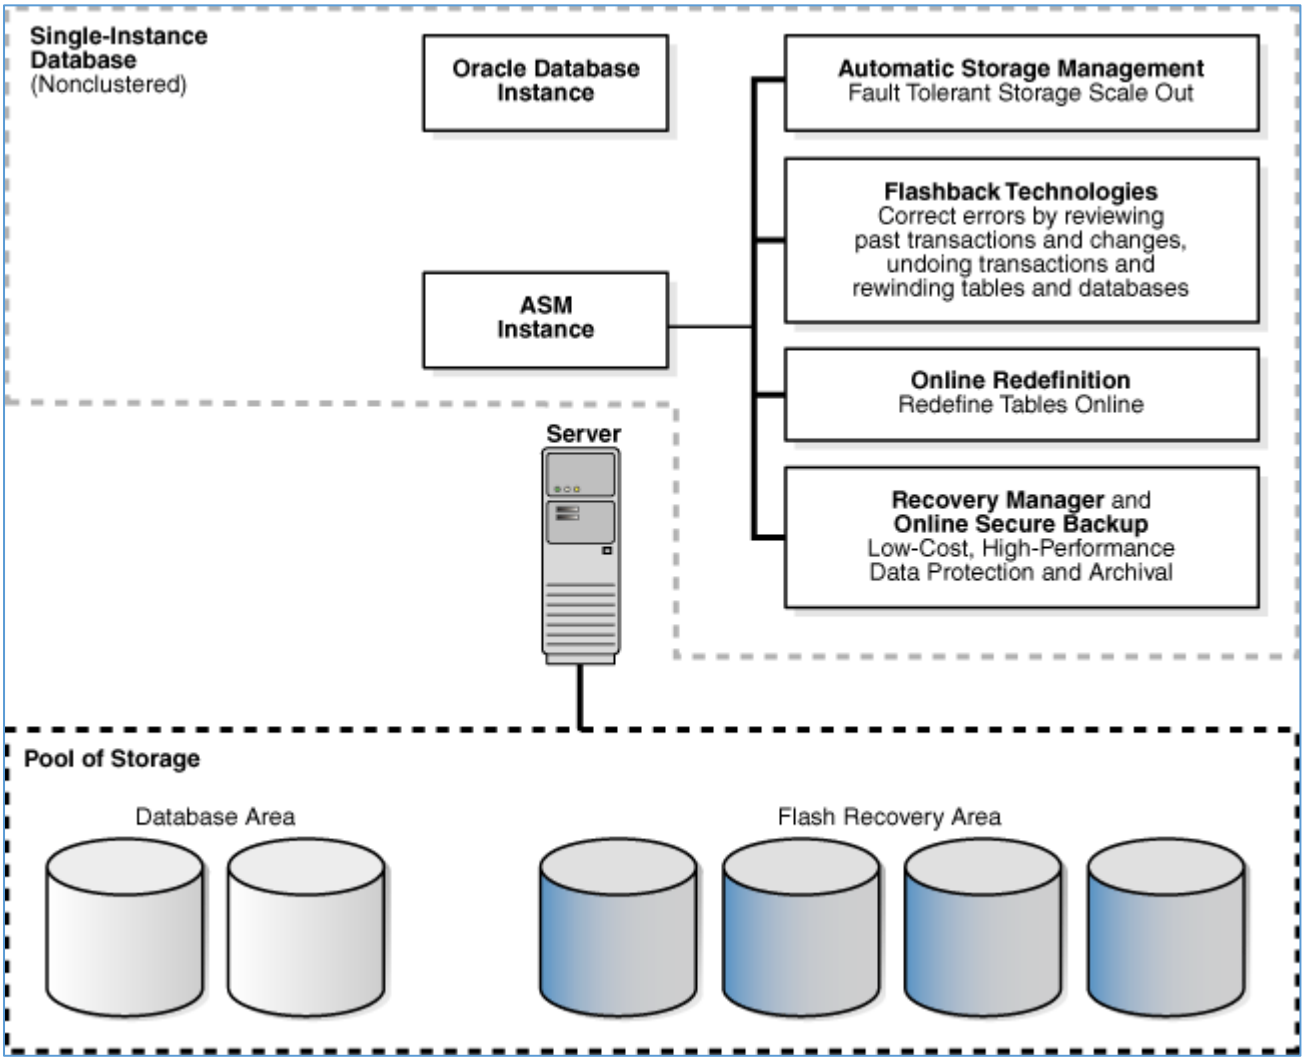

II) Oracle Database with Oracle Clusterware (Cold Failover Cluster)

With Oracle Clusterware you can provide a *cold failover cluster* to protect an Oracle instance from a system or server failure.
The basic function of a cold failover cluster is to monitor a database instance running on a server, and if a failure is detected, to restart the instance on a spare server in the cluster.

Shows a configuration that uses Oracle Clusterware to extend the basic Oracle Database architecture and provide cold failover cluster. In the figure, the configuration is operating in normal mode in which Node 1 is the active instance connected to the Oracle Database that is servicing applications and users. Node 2 is connected to Node 1 and to the Oracle Database, but it is currently standby mode.

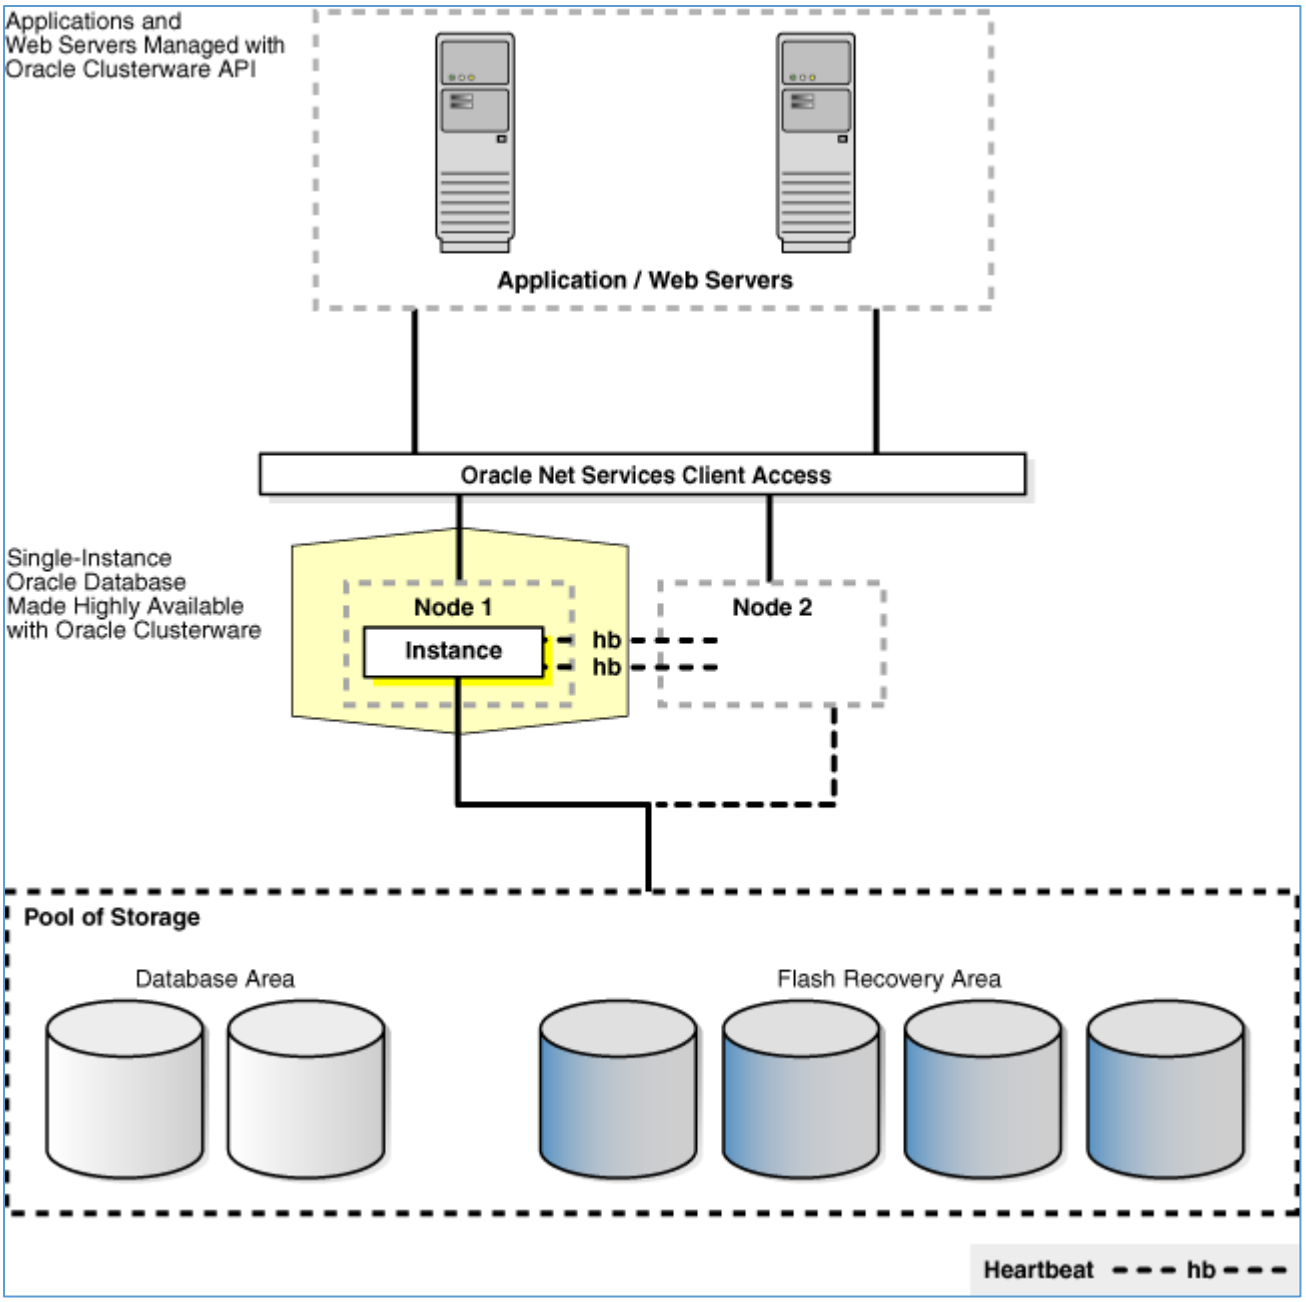

III) Oracle Database with Oracle Real Application Clusters (Oracle RAC)

An architecture that combines the Oracle Database with Oracle RAC. Oracle RAC combines the processing power of multiple interconnected computers to provide system redundancy, scalability, and high availability.

All single instance HA features, such as the Flashback technologies and online reorganization, also apply to Oracle RAC.

Unlike the cold cluster model where one node is completely idle, all instances and nodes can be active to scale your application.

- Scalability across database instances
- Flexibility to increase processing capacity using commodity hardware without downtime or changes to the application
- Ability to tolerate and quickly recover from computer and instance failures (measured in seconds)
- Rolling upgrades for system and hardware changes
- Rolling patch upgrades for some interim patches
- Fast, automatic, and intelligent connection and service relocation and failover
- Load balancing advisory and runtime connection load balancing
- Comprehensive manageability integrating database and cluster features

The Oracle Database with Oracle RAC architecture is designed primary as a scalability solution that resides in a single data center.

Shows the Oracle Database with Oracle RAC architecture

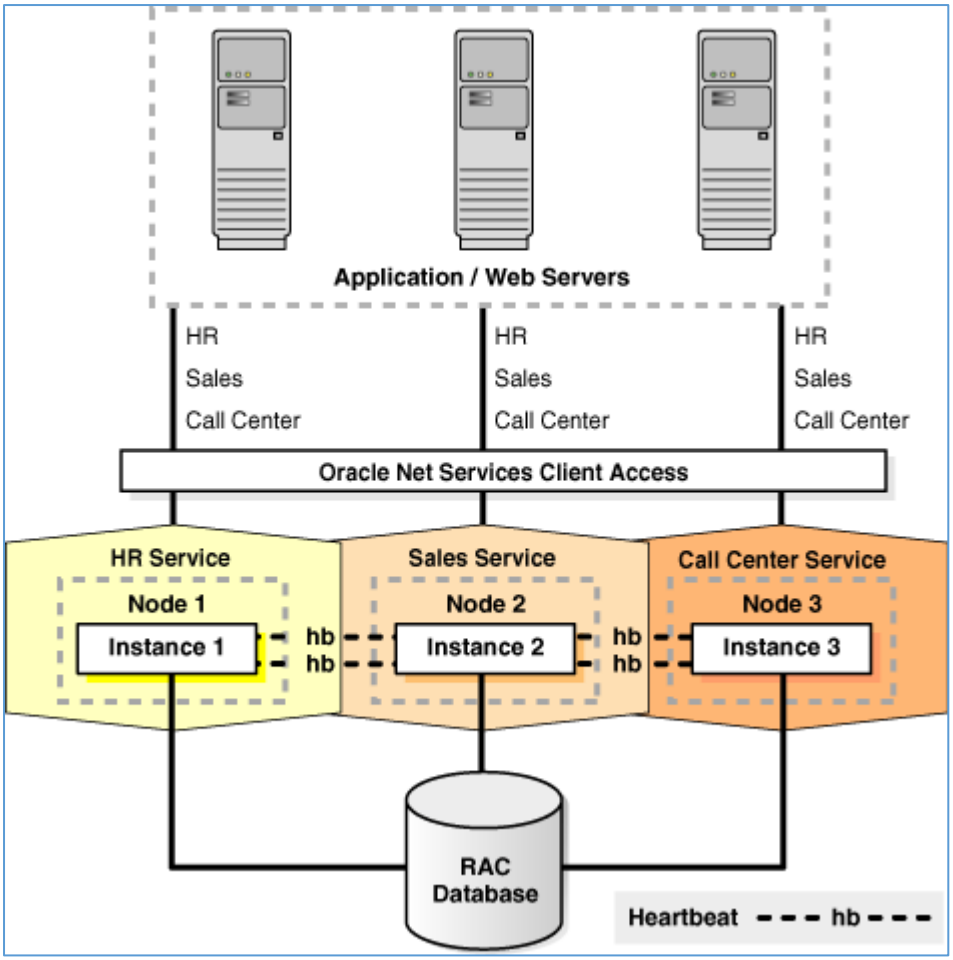

IV) Oracle Database with Oracle RAC on Extended Clusters

The Oracle Database with Oracle RAC architecture is designed primary as a scalability solution that resides in a single data center.

It is possible, under certain circumstances, to build and deploy an Oracle RAC system where the nodes in the cluster are separated by greater distances. This architecture is referred to as an *extended cluster*.

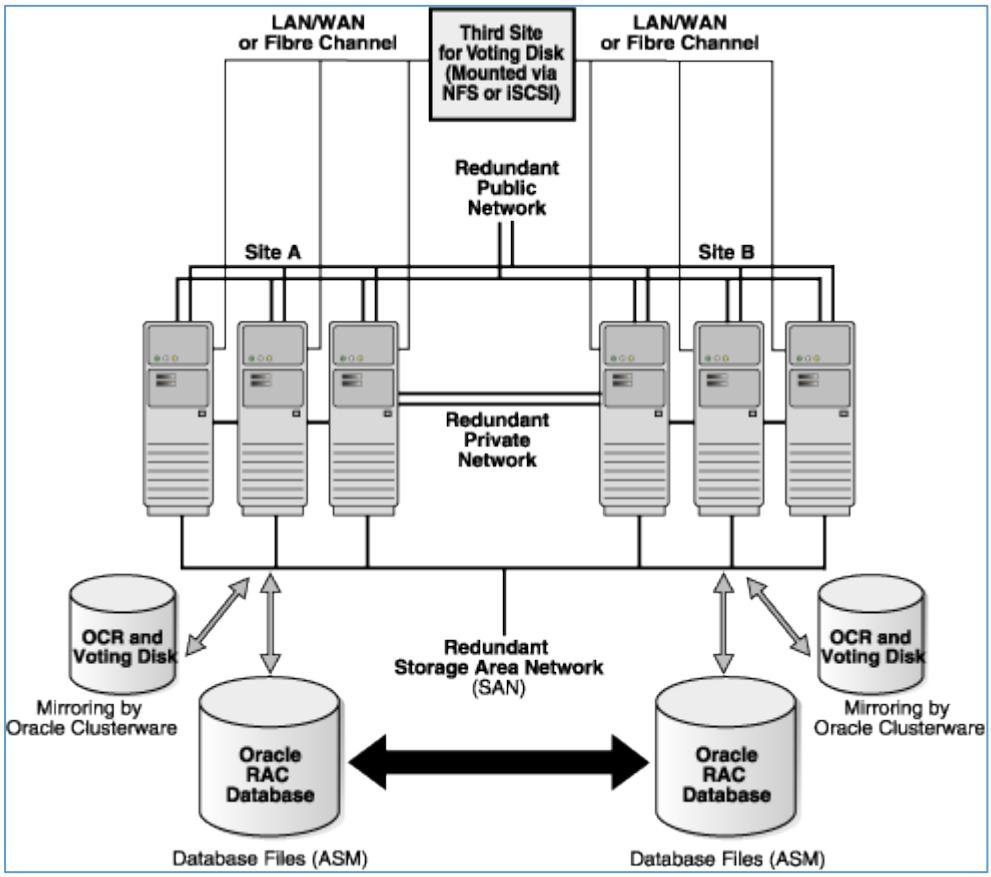

V) Oracle Database with Data Guard

Oracle Data Guard is a high availability and disaster-recovery solution that provides very fast automatic failover (referred to as fast-start failover) in the case of database failures, node failures, corruption, and media failures.
Furthermore, the standby databases can be used for ready-only access and subsequently for reader farms, for reporting purposes, and for testing and development purposes.

Data Guard provides a number of advantages over traditional solutions, including the following:

- Fast, automatic or automated failover for data corruptions, lost writes, and database and site failures
- Protection against data corruptions and lost writes on the primary database
- Reduced downtime with Data Guard rolling upgrade capabilities
- Ability to offload primary database activities, such as backups, queries or reporting without sacrificing RTTO and RPO
- Site failures do not require instance restart, storage remastering, or application reconnections
- Transparent to applications
- Effective network utilization

While traditional solutions (such as backup and recovery from tape, storage based remote mirroring, and database log shipping) can deliver some level of high availability, Data Guard provides the most comprehensive HA and disaster recovery solution for Oracle databases.

- **Better Network Efficiency**
  Only the redo data needs to be sent to the remote site. However, if a remote mirroring solution is used For data protection, typically you must mirror the database files, the online redo logs, the archived redo logs and the control file. If the flash recovery area is on the source volume that is remotely mirrored, then you must also remotely mirror the flashback logs. Thus, compared to Data Guard, a remote mirroring solution must transmit each change many more times to the remote site.

- **Better Performance**
  Data Guard is designed so that it does not affect the Oracle database writer (DBWR) process that writes to data files, because anything that slows down DBWR process affects database performance.
  Data Guard enables you to use the standby database for updates while it continues to protect the primary database.

- **Better suited for WANs**
  Remote mirroring solutions based on storage systems often have a distance limitation due to the underlying communication technology (Fibre Channel, ESCON) used by the storage systems. Maximum distance between these two boxes connected in a point-to-point fashion and running synchronously can be only 10 km. Using specialized devices this distance can be extended to 66 Km.
  You must use a series of repeaters and converters from third-party vendors > 66 Km.
  These devices convert ESCON/Fibre Channel to the appropriate IP, ATM or SONET networks

- **Better resilience and data protection**
  Corruptions are eliminated by Data Guard. If the Host Bus Adaptor corrupts a block as it is written to disk, then a remote mirroring solution may propagate this corruption to the DR site.

- **Higher Flexibility**
  Data Guard only requires a standard TCP/IP-based network link between the two computers.
  It also allows the storage to be laid out in a different fashion from the primary. For example, you can put the files on different disks, volumes, file systems, and so on.

- **Better Functionality**
  Data Guard, with its full suite of data protection features (Redo Apply for **physical standby databases** and SQL Apply for **logical standby databases**, multiple protection modes, push-button automated switchover and failover capabilities, automatic gap detection and resolution, GUI-driven management and monitoring framework, cascaded redo log destinations), is a much more comprehensive and effective solution optimized for data protection and disaster recovery than remote mirroring solutions.

- **Higher ROI**

1) Overview of Single Standby Database Architectures

A single standby database architecture consists of the following key traits and recommendations.

- Primary database resides in Site A.
- Standby database resides in Site B.
  Synchronous redo transport does not impose any physical distance limitation.

- Fast-start failover is recommended to provide automatic failover
- Use a physical standby database if read-only access is sufficient
- Evaluate logical standby databases if additional indexes are required for reporting purposes and if application only uses data types supported by logical standby database and SQL Apply

Shows the relationship of Primary and Standby Databases and the Observer during Fast-Start Failover

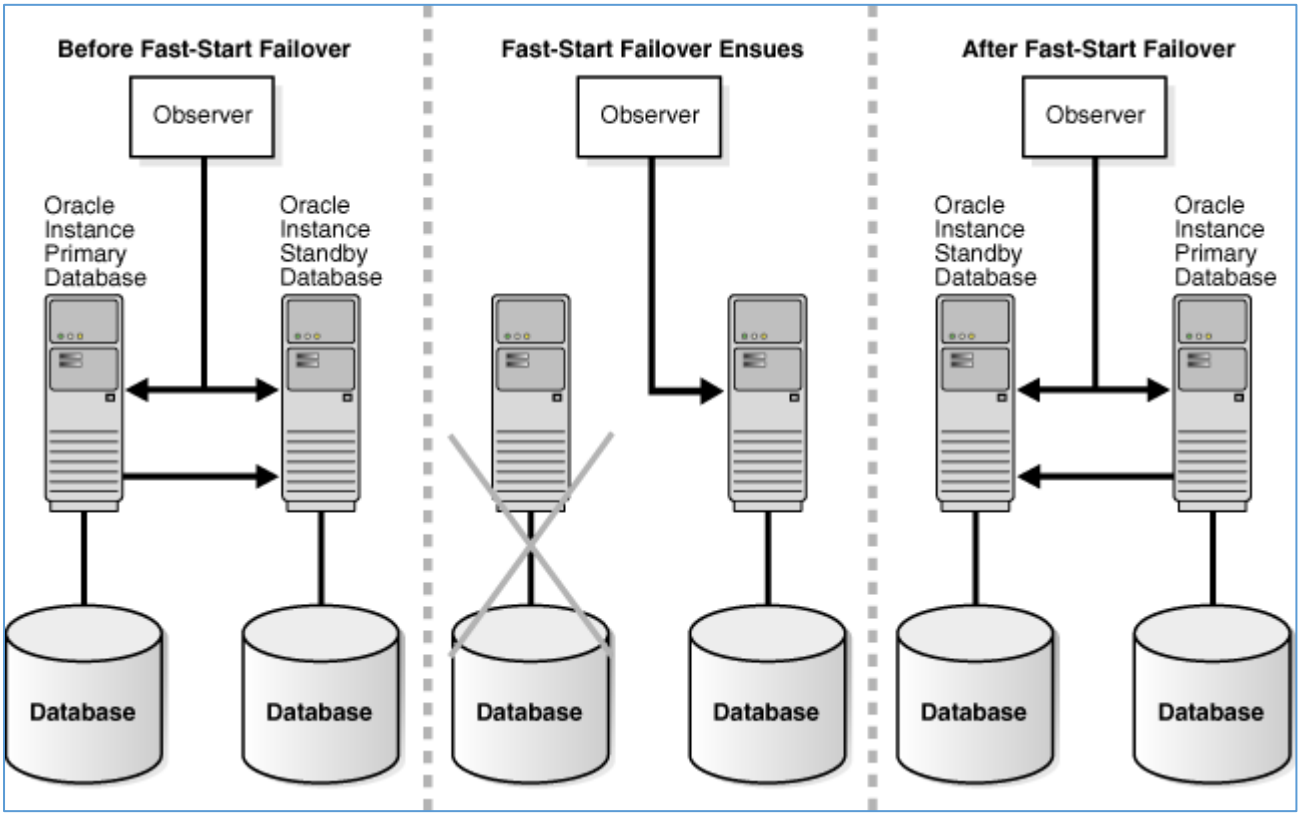

2) Overview of Multiple Standby Databases Architectures

There are multiple standby databases in the same Data Guard configuration.

- Continuous and transparent disaster or HA protection in case of an outage at the primary database or the targeted standby database
- Reader farms or look up databases
- Reporting databases
- Regional reporting or reader databases for better response time
- Synchronous transport transmits to a more local standby database, and asynchronous transport to a more remote standby database to provide optimum levels of performance and data protection
- Testing and development clones using snapshot standby databases.
- Rolling upgrades

It is possible to convert a physical standby database to a logical standby database or to a snapshot standby database, or you can create additional logical standby databases or snapshot standby databases.

Transient logical standby databases:

Can be used to minimize downtime for database upgrades. It is helpful in Data Guard architectures where there are no logical standby databases.

In a multiple standby database environment, you can create a transient logical standby database temporarily (for planned maintenance) and then convert it back to the physical standby database role. For example, you can use transient logical standby databases to minimize downtime for database upgrades, when required. There is no need to create a separate logical standby database to perform upgrades. The high-level steps for rolling upgrades with a transient logical standby database are as follows:

1. Start performing a rolling database upgrade with the physical standby database.
2. Temporarily convert the physical standby database to a logical standby database to perform the upgrade. (Note that data type restrictions are limited for the short window of time required to perform an upgrade.)
3. Revert the logical standby database back to the physical standby database role.

Snapshot standby databases:

Can be used as a clone or a test database to test new functionality and new releases. The snapshot standby database continues to receive and queue redo data so data protection and RPO are not sacrificed. Snapshot standby databases diverge from the primary database over time because redo data from the primary database is not applied when it is received. Redo Apply does not apply the redo data until you convert the snapshot standby database back into a physical standby database, and all local updates that were made to the snapshot standby database are discarded. Although the local updates to the snapshot standby database cause additional divergence, the data in the primary database is fully protected by means of the redo logs that are located at the standby site.

Shows the production database at the primary site and multiple standby databases at secondary sites.

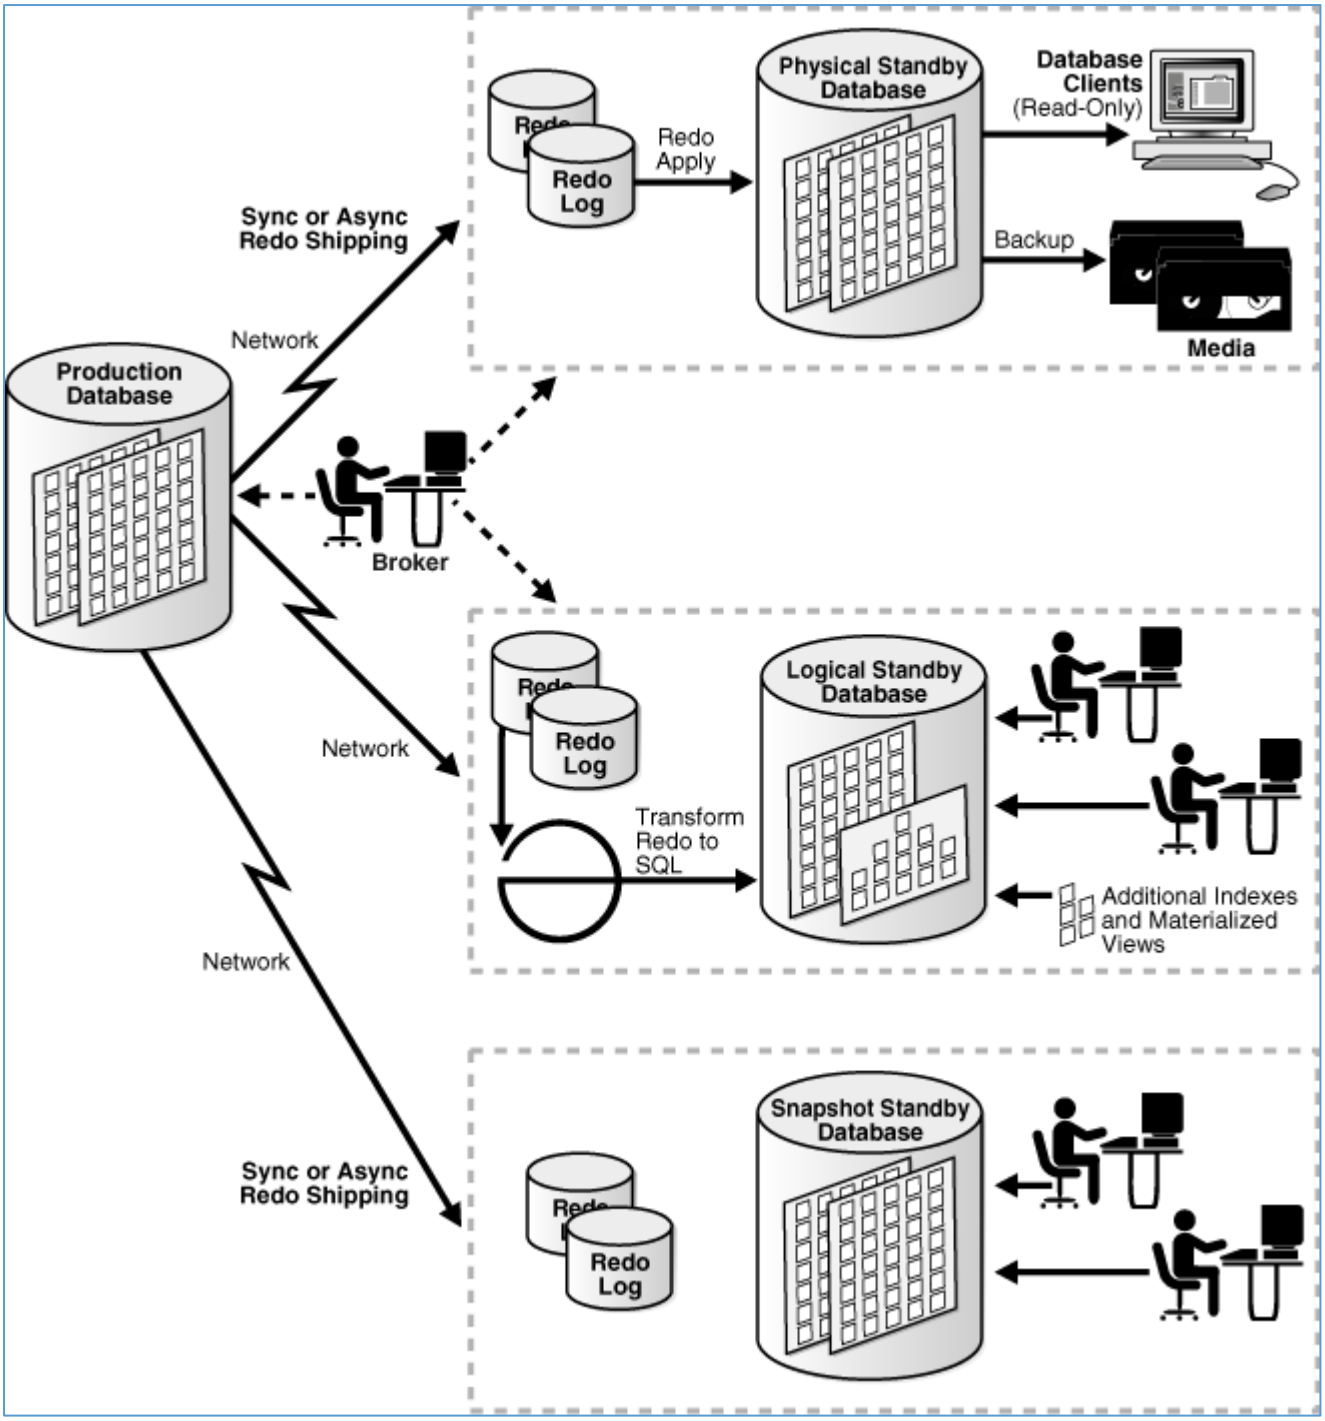

The following list describes examples of Data Guard configurations using multiple standby databases:

A world-recognized financial institution uses two remote physical standby databases for continuous data protection after failover. If the primary system should fail, the first standby database becomes the new primary. The second standby database automatically receives data from the new primary, insuring that data is protected at all times.

A nationally recognized insurance provider in the U.S. maintains two standby databases in the same Data Guard configuration, one physical and one logical standby database. Their strategy further mitigates risk by maintaining multiple standby databases, each implemented using a different architectures - Redo Apply and SQL Apply.

A world-recognized e-commerce site utilizes multiple standby databases—a mix of both physical and logical databases - both for disaster recovery purposes and to scale-out read performance by provisioning multiple logical standby databases using SQL Apply.

A global provider of information services to legal and financial institutions uses multiple standby databases in the same Data Guard configuration to minimize downtime during major database upgrades and platform migrations.

Also, for large data centers where there is a need to support many applications with Data Guard requirements, you can build **a Data Guard hub** to reduce the total cost of ownership.
With the Database Server and Storage Grid, you can build standby database and testing Hubs that leverage a pool of system resources. The system resources can be dynamically allocated and deallocated depending on various priorities. For example, if the primary database fails over to one of the standby databases in the standby hub, the new primary database acquires more system and storage resources while the testing resources may be temporarily starved. With the Oracle Grid technologies, you can enable a high level of utilization and low TCO, without sacrificing business requirements.

A Data Guard hub can consists of:
Several standby databases in an Oracle RAC environment residing in a cluster of servers, called a grid server
Leveraging the storage grid
The premise of the standby hub is that it provides higher utilization with lower cost. The probability of failing over all the databases at the same time is unlikely. Thus, when there is a failover, you can prioritize the system resources to production activity and allocate new system resources in a grid for the standby database functions. At the time of role transition, more storage and system resources can be allocated toward that application.

For example, a Data Guard hub could include multiple databases and applications that are supported in a Grid server and storage architecture. This configuration consists of a central resource supporting 10 applications and databases in the grid compared to managing 10 separate system or storage units in a nongrid infrastructure.

Another possible configuration might be a testing hub consisting of snapshot standby databases. With the snapshot standby database hub, you can leverage the combined storage and server resources of a Grid instead of building and managing individual servers for each application.

VI) Oracle Database with Oracle Clusterware and Data Guard

If your business does not require the scalability and additional high availability benefits provided by Oracle RAC, but you still need all the benefits of Oracle Data Guard and cold failover cluster, then this architecture is a good compromise. With Oracle Database 11g, Oracle Clusterware cold failover cluster combined with Oracle Data Guard makes a tightly integrated solution in which failover to the secondary node in the cold failover cluster is transparent and does not require you to reconfigure the Data Guard environment or perform additional steps.

Shows an Oracle Clusterware and Oracle Data Guard architecture that consists of a primary and a secondary site. Both the primary and secondary sites contain Oracle application servers, two database instances, and an Oracle Database.

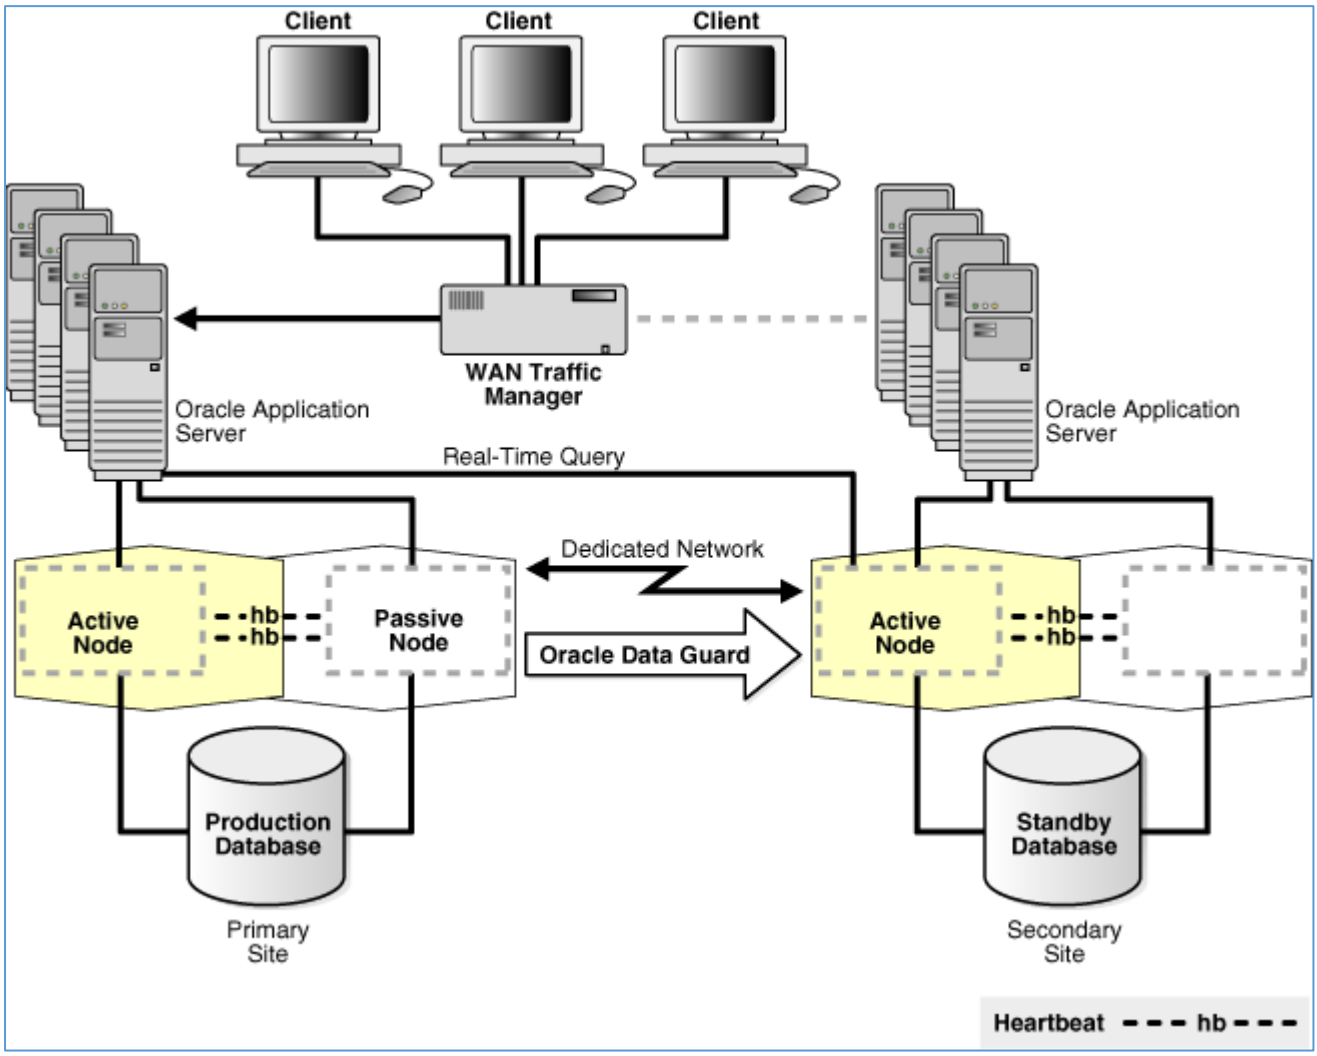

- The application servers on the secondary site are connected to the WAN traffic manager by a dotted line to indicate that they are not actively processing client requests at this time. The application server on the secondary site can be active and processing client requests such as queries if the standby database is a physical standby database with the Active Data Guard option enabled, or if it is a logical standby database.
- Oracle Data Guard transmits redo data from the primary database to the secondary site to keep the databases synchronized.
- Oracle Clusterware manages the availability of both the user applications and Oracle databases.
- Oracle Clusterware provides tolerance of node failures, while Data Guard provides additional protection against data corruptions, lost writes, and database and site failures. (See Oracle Database with Data Guard for a complete description.)

- Although cold failover cluster is not shown in Figure 4-8, you can configure it by adding a passive node on the secondary site.

VII) Oracle Database with Oracle RAC and Data Guard

This architecture combines the benefits of both Oracle RAC and Data Guard and it is the recommended architecture for Maximum Availability Architecture (MAA).
The MAA recommends Oracle RAC and Data Guard reside on separate systems (clusters) and data centers. Configuring symmetric sites is recommended to ensure that each site can accommodate the performance and scalability requirements of the application after any role transition. Furthermore, operational practices across role transitions is simplified when the sites are symmetric

Shows Oracle Database with Oracle RAC and Data Guard => MAA

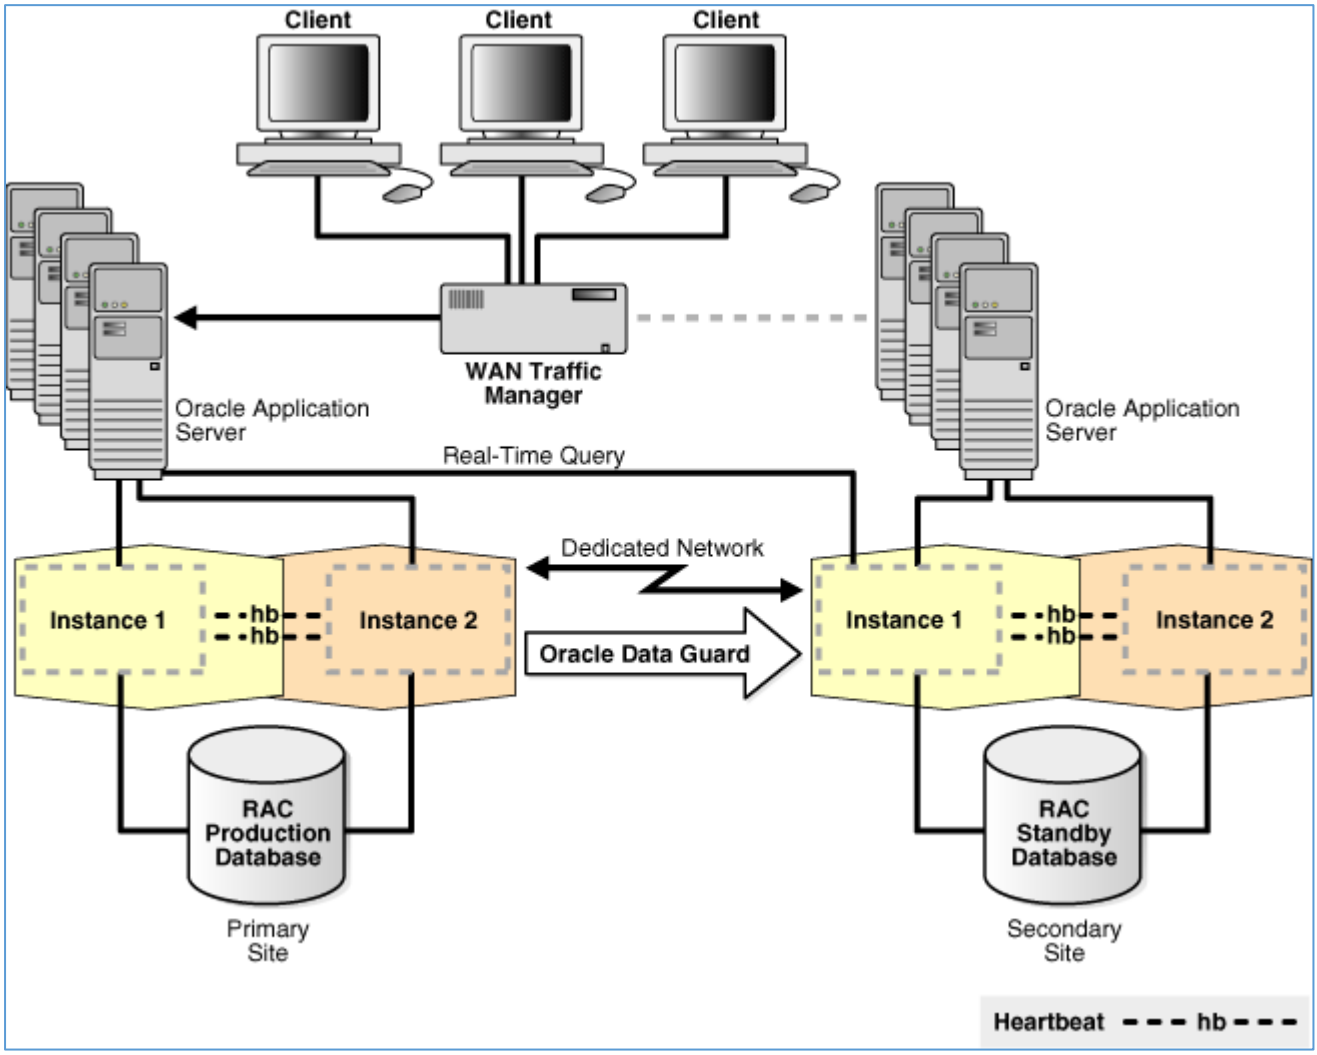

Although cold failover cluster is not shown in Figure 4-8, you can configure it by adding a passive node on the secondary site.

VIII) Oracle Database with Streams