



High Availability Options for Oracle Database

An IT Convergence presentation by Dan Norris

Agenda

Database High Availability Spectrum

Oracle Export/Import

Oracle 10g Data Pump

Oracle RMAN

Oracle Data Guard

Oracle Advanced Replication

Oracle Streams

Failover Clusters

Oracle Real Application Clusters

Summary

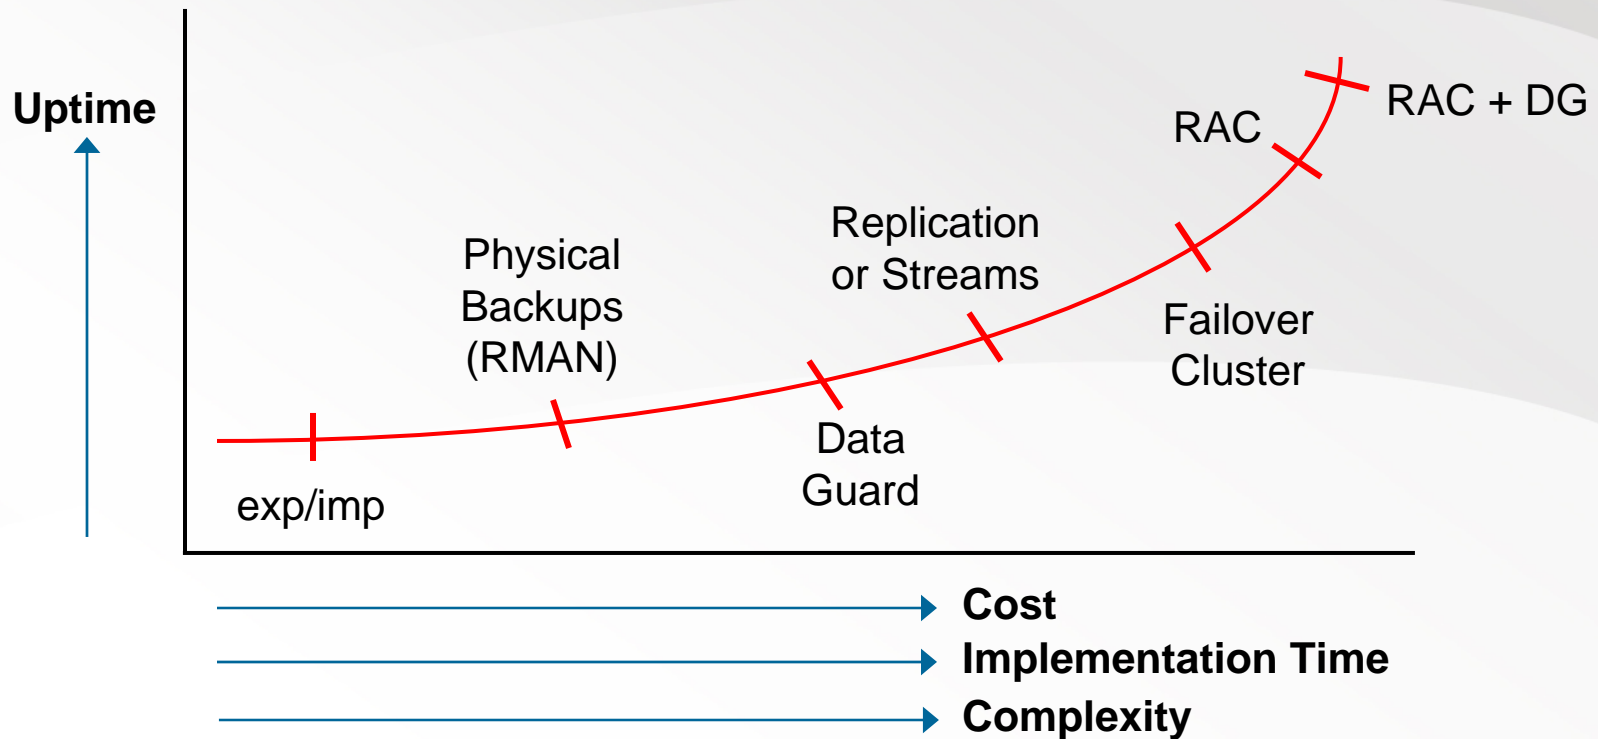
Oracle Database High Availability Spectrum

Many options for availability

Some options are very basic, but necessary

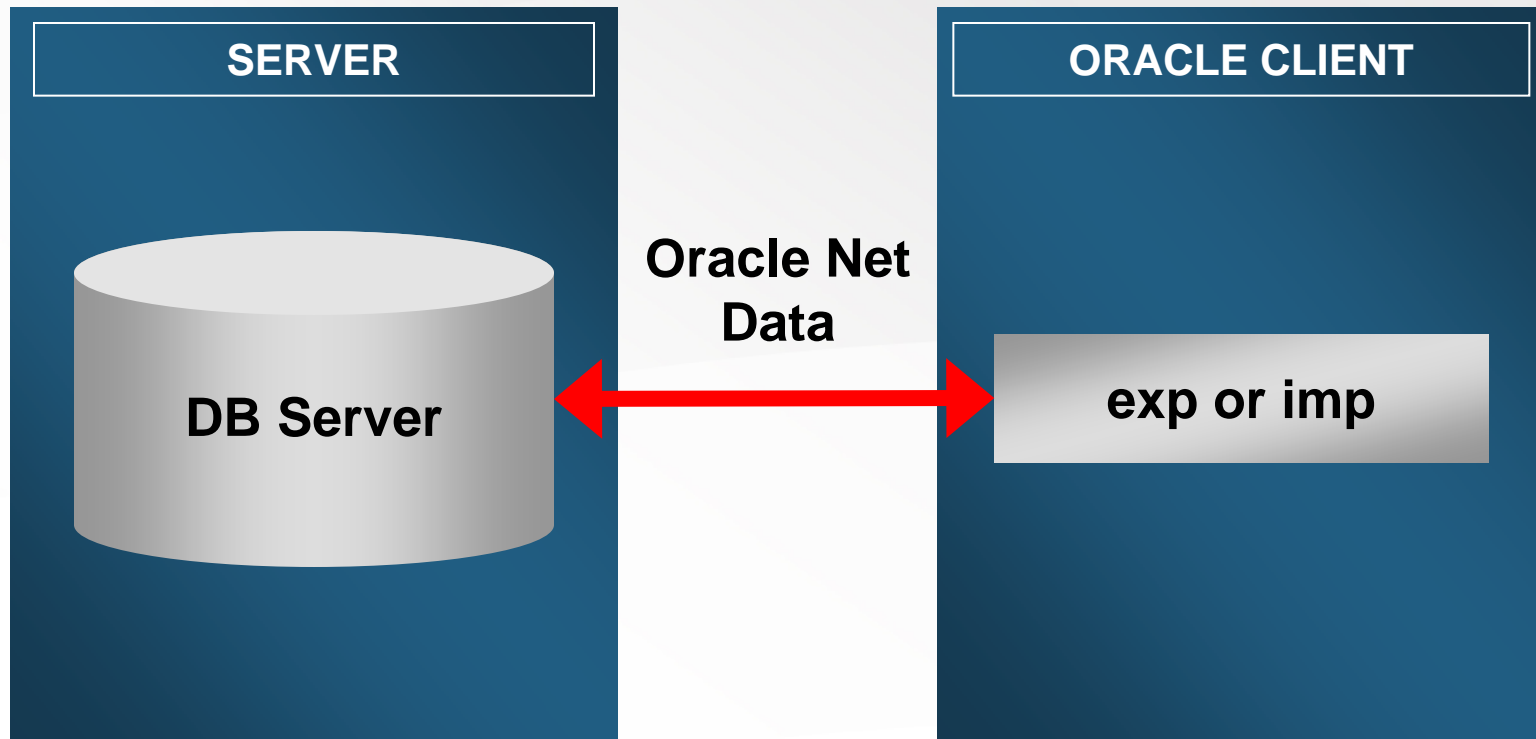
Cost, complexity, and implementation time
are factors to consider

Oracle Database High Availability Spectrum



One certainty: failures will occur.

Oracle Export/Import Overview



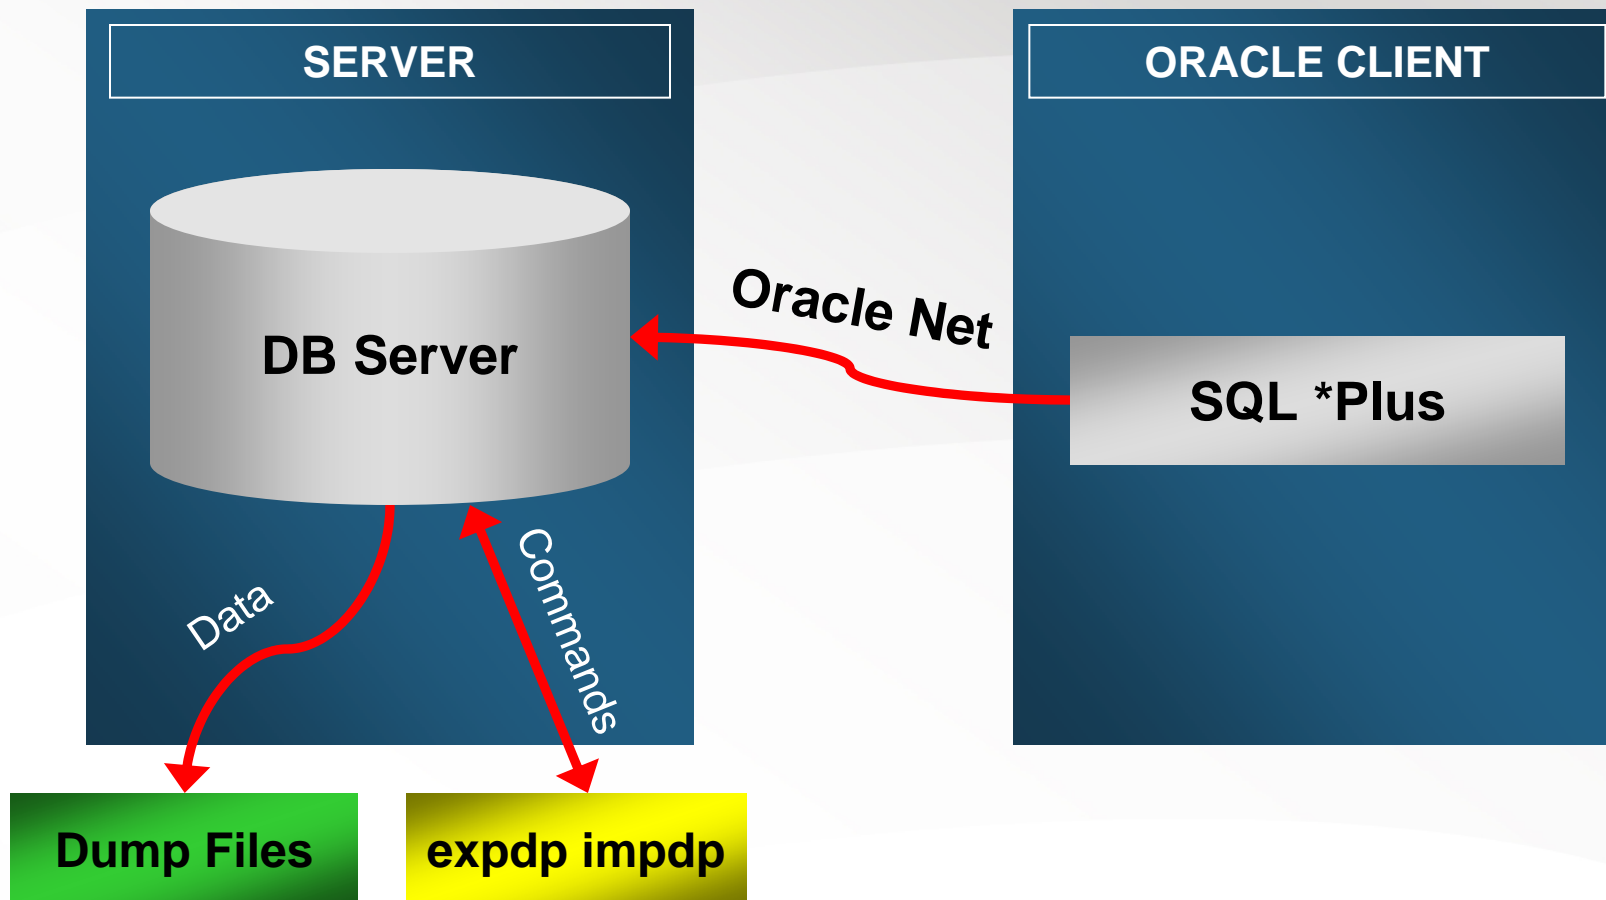
Oracle Export/Import Advantages

- ▶ Well-known, well-defined, stable
- ▶ Simple to implement, very few prerequisites
- ▶ Cross-platform, multiple version compatibility

Oracle Export/Import Disadvantages

- ▶ Time consuming
- ▶ Resource intensive
- ▶ Large files produced
- ▶ No "roll forward" capability
- ▶ Does not meet most availability requirements

Oracle Data Pump Overview



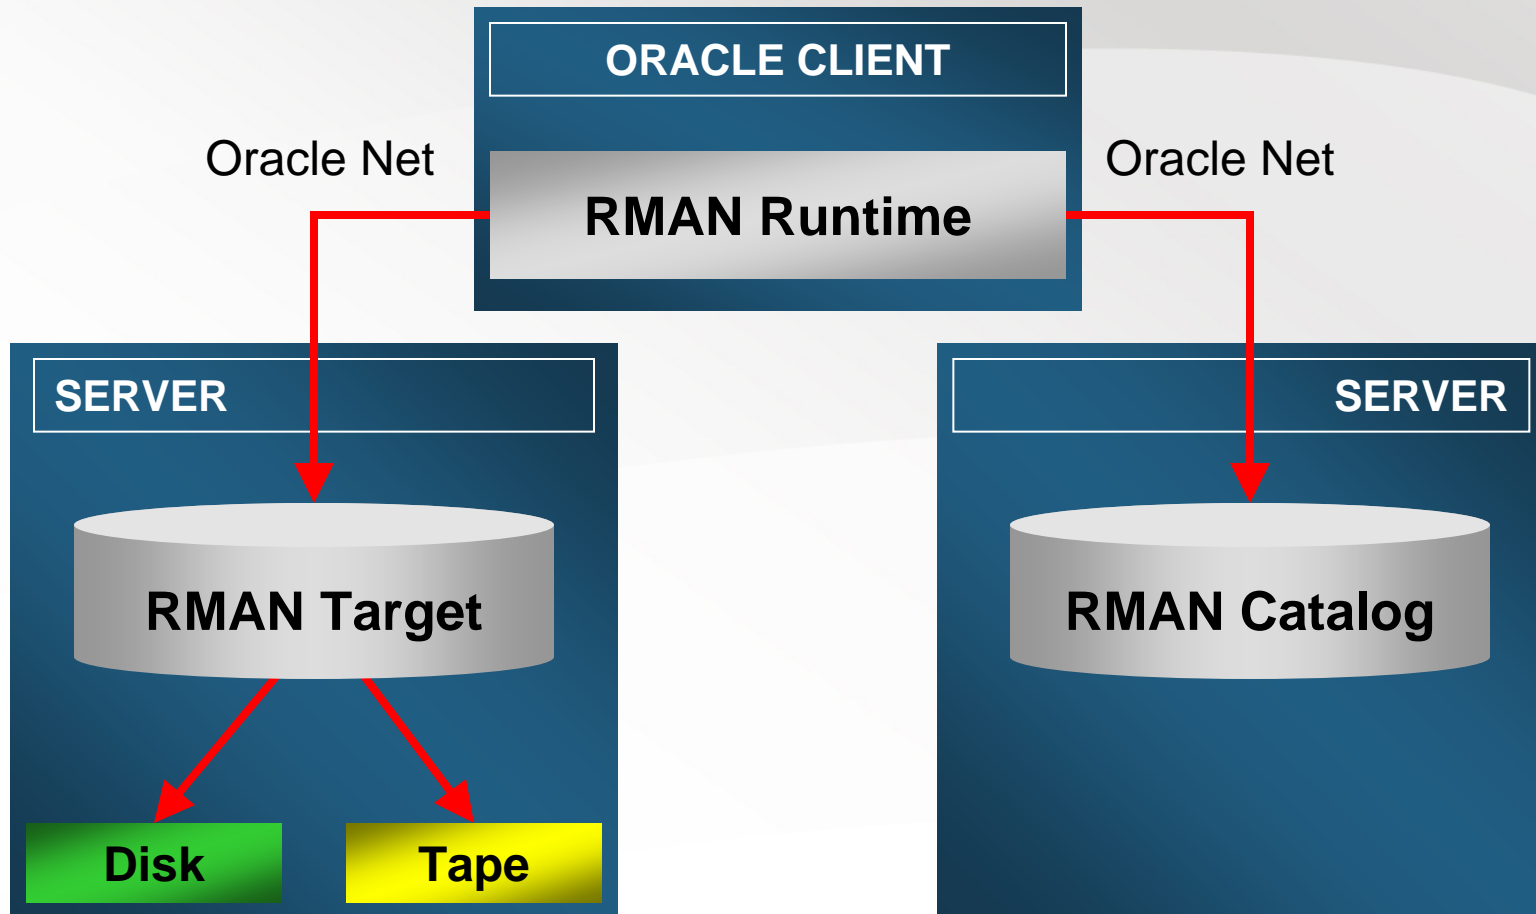
Oracle Data Pump Advantages

- ▶ Programmatic API: DBMS_DATAPUMP
- ▶ Cross-platform compatibility
- ▶ Parallelism capability for speedy operations
- ▶ Architecture allows detach and reattach to long-running jobs

Oracle Data Pump Disadvantages

- ▶ Server-side only, no client/server
- ▶ Resource intensive
- ▶ Large files produced
- ▶ Only available in 10g and higher (new technology)
- ▶ Does not meet most availability requirements

Oracle RMAN Overview



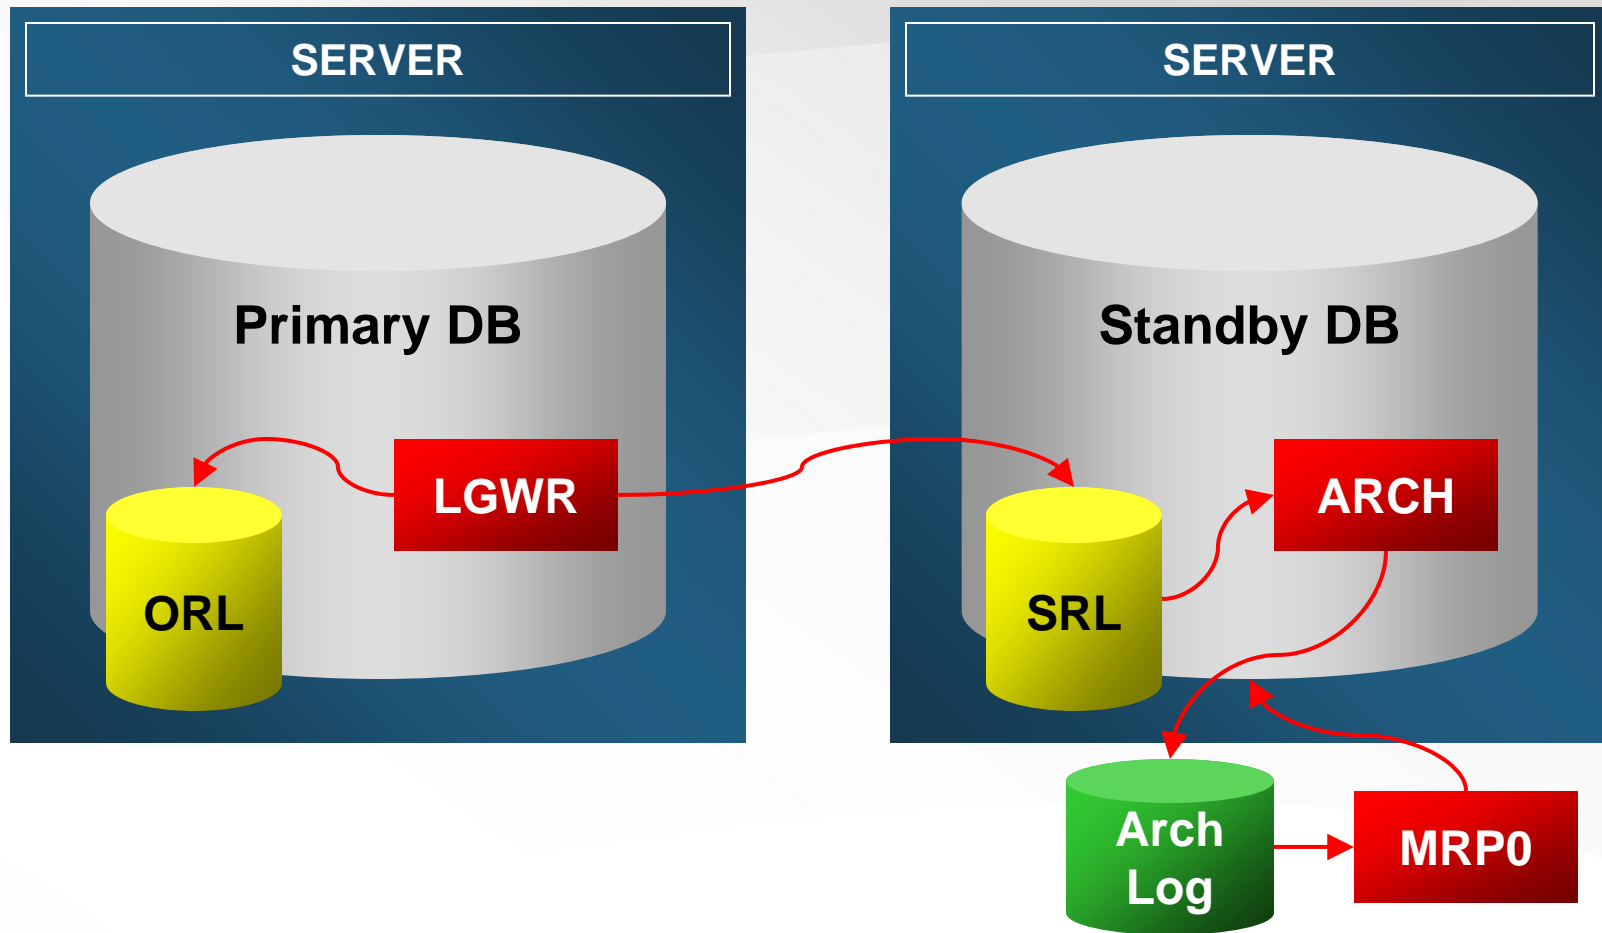
Oracle RMAN Advantages

- ▶ Common interface for all platforms
- ▶ Incremental backup functionality
- ▶ OS-independent scripting language
- ▶ Backup automation with very few commands
- ▶ Parallelism built in to speed up all operations

Oracle RMAN Disadvantages

- ▶ Separate catalog database needed for some features
- ▶ Proprietary syntax
- ▶ Learning RMAN can be challenging for some DBAs
- ▶ Compatibility can be tricky between releases

Oracle Data Guard Overview



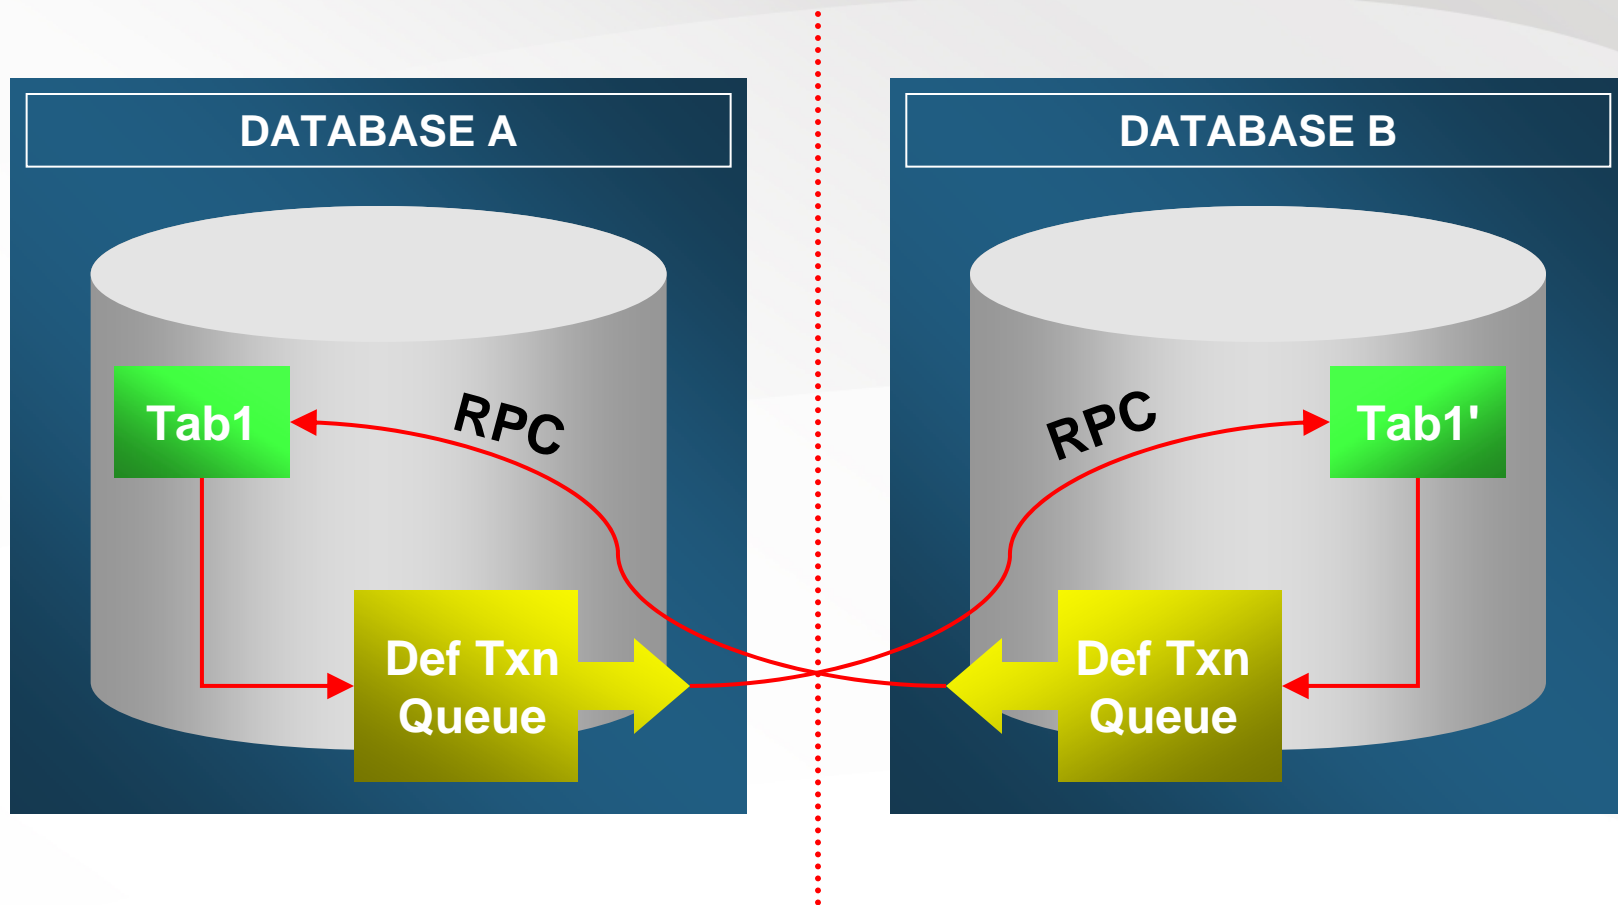
Oracle Data Guard Advantages

- ▶ Geographically dispersed sites
- ▶ Useful for logical data corruptions if lag behind used
- ▶ Flexible configuration options for protection level
- ▶ Reporting and backups can be diverted to standby
- ▶ Automatic resync for failed primary

Oracle Data Guard Disadvantages

- ▶ Same platform and OS required
- ▶ Additional database copies require storage and maintenance
- ▶ Specific knowledge required, training may be needed
- ▶ Typically a DR option, not "true" HA

Oracle Advanced Replication Overview (Asynchronous)



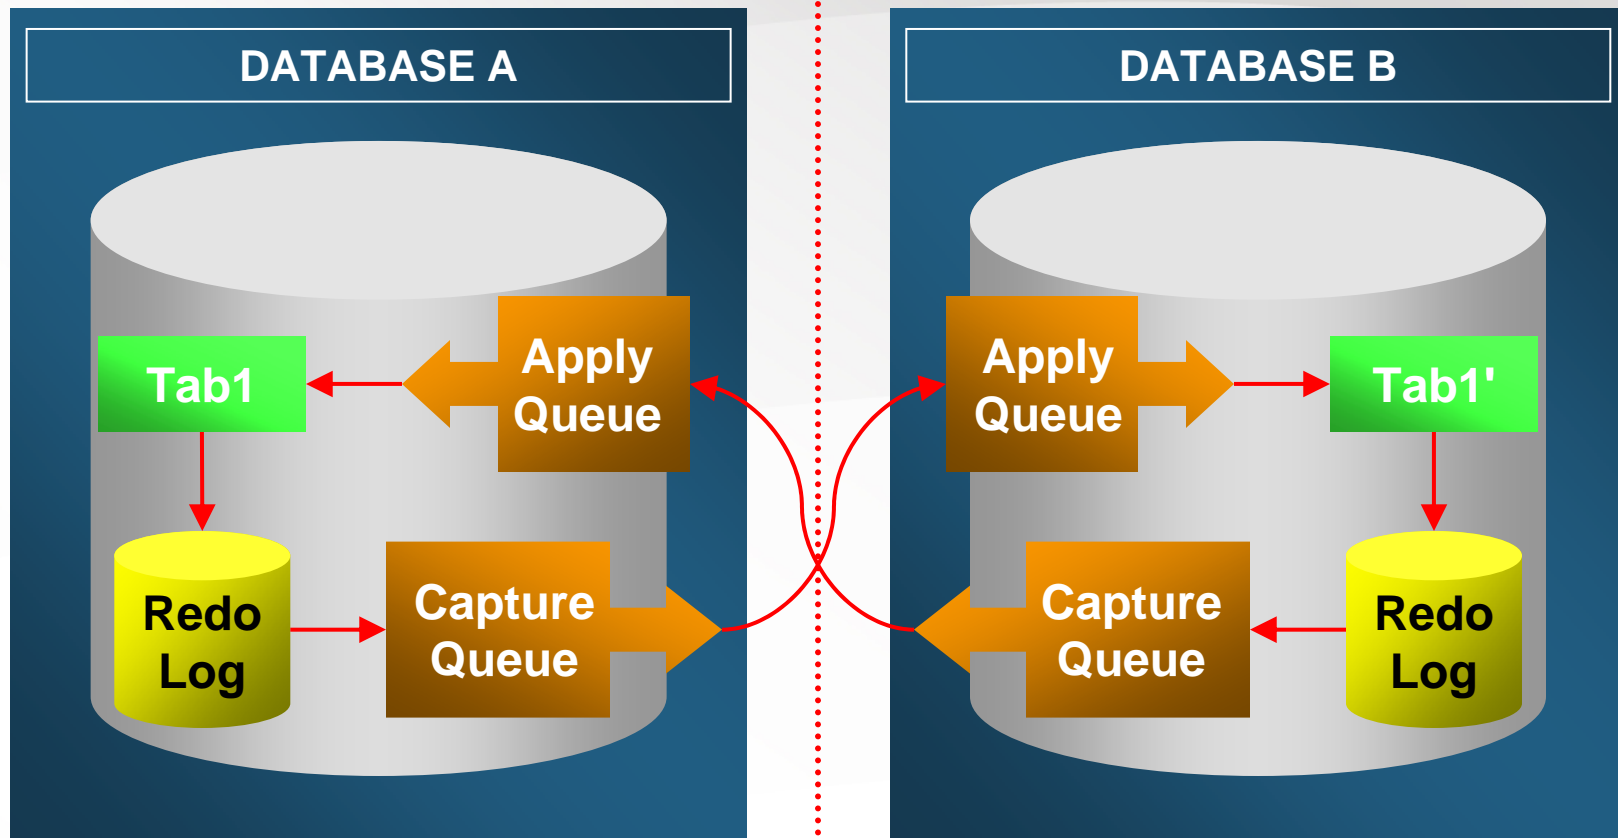
Oracle Advanced Replication Advantages

- ▶ Geographically dispersed sites
- ▶ Long history, well-deployed
- ▶ Any platform, any OS, most DB versions
- ▶ Relatively simple architecture

Oracle Advanced Replication Disadvantages

- ▶ Triggers on tables cause overhead
- ▶ Not all DBAs have experience with Adv Rep
- ▶ No easy way to fail over users to remote site
- ▶ Conflict resolution can be challenging
- ▶ Replication delay typically at least 60 seconds

Oracle Streams Overview




Oracle Streams Advantages

- ▶ Geographically dispersed sites
- ▶ Included with Enterprise Edition
- ▶ Source site continues even if capture is disabled
- ▶ Typically quick propagation (seconds)
- ▶ Oracle is encouraging Streams usage

Oracle Streams Disadvantages

- ▶ New technology, not as much knowledge available
- ▶ Relatively complex architecture
- ▶ Conflict resolution can be challenging
- ▶ Learning curve can be steep
- ▶ Troubleshooting can be challenging

Failover Cluster Overview



Failover Cluster Advantages

- ▶ Straightforward implementation, no application changes
- ▶ Clusterware vendors support Oracle directly
- ▶ Oracle Clusterware is included with DB license
- ▶ Automatic monitoring and failover actions


Failover Cluster Disadvantages

- ▶ Additional testing required
- ▶ Some additional training and management planning needed
- ▶ Clusterware can be expensive
- ▶ Troubleshooting can be more challenging
- ▶ Outage is typically 2 to 5 minutes

Failover Cluster Products

- ▶ Microsoft Cluster Server (w/ Oracle Failsafe)
- ▶ Veritas Cluster Server
- ▶ PolyServe Matrix Server (now Novell's)
- ▶ Red Hat Cluster Server
- ▶ HP MC ServiceGuard
- ▶ IBM HACMP

Oracle Real Application Clusters Overview



Oracle RAC Advantages

- ▶ Most applications require no changes
- ▶ Incredible scalability and zero downtime capable
- ▶ Well-defined, mature product with active community of support (www.oracleacsig.org)
- ▶ No third-party software required
- ▶ No special hardware required

Oracle RAC Disadvantages

- ▶ RAC license can be expensive (except for SE)
- ▶ Some applications require additional exception handling
- ▶ Extensive testing required for configuration changes
- ▶ Third-party application vendors have certified RAC
- ▶ DBA training necessary for success

Agenda

Database High Availability Spectrum

Oracle Export/Import

Oracle 10g Data Pump

Oracle RMAN

Oracle Data Guard

Oracle Advanced Replication

Oracle Streams

Failover Clusters

Oracle Real Application Clusters

Summary



RAC For Beginners: The Basics

An IT Convergence presentation by Dan Norris

Agenda

Terminology

Brief History of RAC

What RAC is NOT

Single-instance Databases vs. RAC

Licensing

Installation Process

Services & Workload Management

Tuning RAC

RAC Backup and Recovery

New RAC DBA Recommendations

Vendor-Provided Applications

High Availability Alternatives

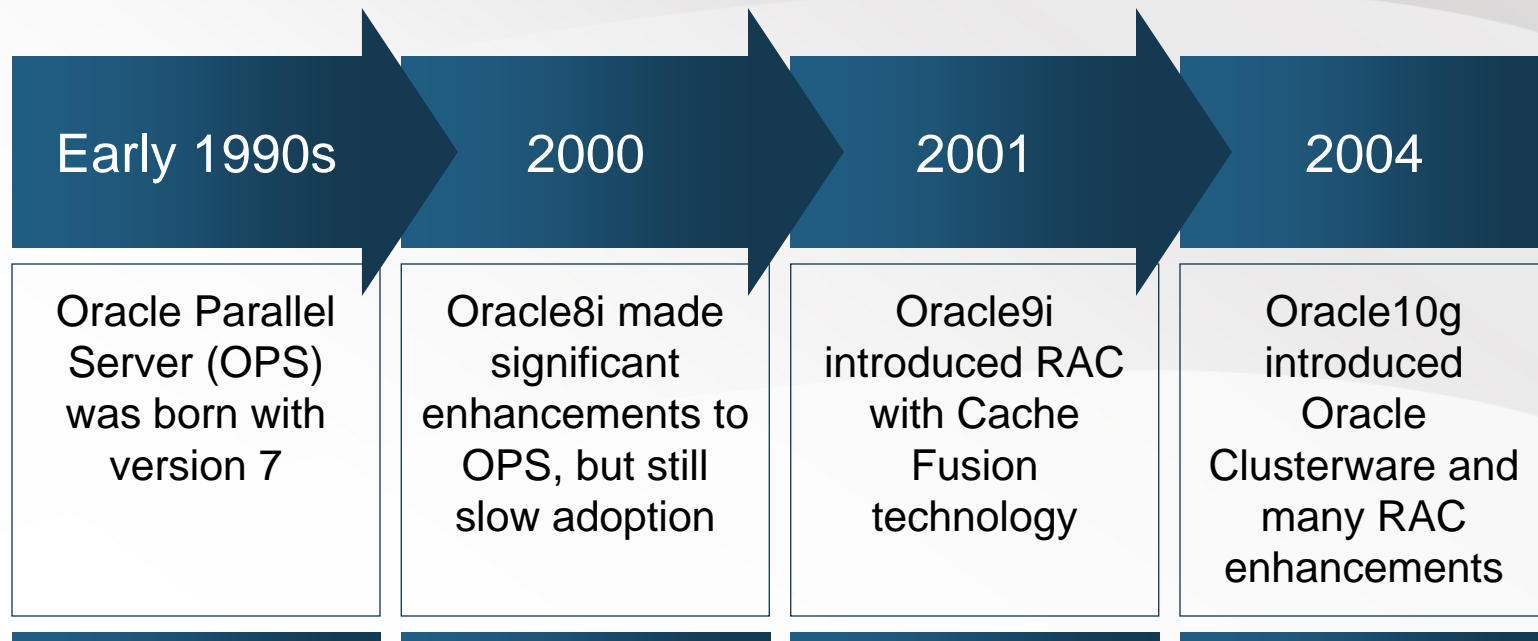
Next Steps & References

Terminology



Database	Instance	Clusterware	Storage Area Network (SAN)
	Local Storage, Shared Storage	Raw Device, Cluster Filesystem	Automatic Storage Management (ASM)
		Single-instance DB, Multi-instance DB	Oracle Services

Brief History of RAC



What RAC is NOT

NOT

NOT always the best choice for the given situation


NOT “production only” technology

NOT something easy to learn only by doing
(without also studying)

NOT a “set it and forget it” environment

NOT a transparent change for some applications

Single-Instance vs. RAC: System Architecture



Single-Instance vs. RAC: Database Configuration

- ▶ Shared Database Components
 - ▶ Control Files, Temp Tablespace, Application Tablespaces, spfile

- ▶ Unshared Database Components (still reside in the database)
 - ▶ Online Redo Logs, Undo Tablespaces, Rollback Segments (if not using AUM)

- ▶ Locally-Managed Tablespaces & Automatic Segment Space Management (ASSM)

- ▶ Server Parameter File

Single-Instance vs. RAC: Database Access Considerations

- ▶ Services should be designed and thought out
- ▶ There are benefits to using multiple services for a single database
- ▶ Do we load-balance all users on all nodes, or segregate workload to groups of nodes?
- ▶ Can we use runtime connection load balancing and/or FCF?
- ▶ Will we configure Transparent Application Failover (TAF)?

Licensing

- ▶ Included with Oracle Database Standard Edition
 - ▶ Maximum 4 CPUs per cluster
 - ▶ Must use ASM for all database storage
 - ▶ Must use *only* Oracle Clusterware (no 3rd party clusterware)

- ▶ An option added to Oracle Enterprise Edition
 - ▶ Required for higher CPU counts
 - ▶ Required to use EE-only features with RAC

Installation Process

- ▶ First, prepare the hardware environment

- ▶ There are four major installation tasks:
 - ▶ Plan the installation, particularly storage
 - ▶ Install Clusterware
 - ▶ Install ASM
 - ▶ Install RDBMS

Installation Process: Prepare

- ▶ Hardware needed usually includes:
 - ▶ More than one private interconnect network
 - ▶ Shared storage, usually fibre channel or iSCSI
 - ▶ Multiple servers, same OS, same packages and patches

- ▶ Storage can be configured as:
 - ▶ Oracle Cluster Filesystem (OCFS)
 - ▶ ASM
 - ▶ Raw
 - ▶ Other 3rd-party cluster filesystems (i.e. GPFS, VxFS)

- ▶ Determine storage for each component

- ▶ Study the installation guide for your platform

Installation Process: Clusterware

- ▶ The OUI checks node connectivity
- ▶ For UNIX hosts, use SSH key pairs to allow transparent logins to remote hosts
- ▶ For Windows hosts, use the same username/password on all hosts
- ▶ The OUI prompts for location of Oracle Cluster Registry (OCR) copies. You can have 1 or 2 of these
- ▶ The OUI also asks for the voting disk location(s). You can have 1 or 3 of these

Installation Process: ASM

- ▶ The ASM installation should reside in a separate ORACLE_HOME per best practices
- ▶ The OUI should recognize the cluster
- ▶ Configure as much storage as you'll need in order to avoid rebalancing operations later
- ▶ Use at least two disk groups: data & flash recovery area
- ▶ If using Oracle Standard Edition, ASM is required for all database data

Installation Process: RDBMS

- ▶ The OUI should recognize the cluster
- ▶ Shared or local ORACLE_HOME?
- ▶ Complete the install without database creation (software only)
- ▶ Install patches (patch sets, CPU); OPatch is cluster-aware
- ▶ With all patches applied, then use DBCA to create the database

Installation Process: Testing

- ▶ Failure testing
- ▶ OS crash
- ▶ Backup & recovery testing
- ▶ Disaster recovery testing (Data Guard, storage replication, etc)
- ▶ Client failover testing (TAF, FCF, load balancing)
- ▶ Don't cut corners on testing time

Services & Workload Management

- ▶ A service is an entity to which users connect
- ▶ Usually designates a module or application used by a specific group of users
- ▶ Technically, a service is listed in the `service_name` parameter for an instance (Note: You should not edit the `service_name` parameter in a RAC environment.)
- ▶ Clusterware processes alter the `service_name` parameter on the fly to relocate services (according to policies)
- ▶ Stats in 10g are also gathered per service

Services & Workload Management



- ▶ Services can be available via one or more instances
- ▶ Failover policies are set per service

Tuning RAC

- ▶ Good news! The first step is to do all the normal single-instance tuning
- ▶ Tuning RAC can pose special challenges, but don't abandon what you know
- ▶ Examining bottlenecks on the interconnect is the most common RAC-specific activity
- ▶ Statspack, ADDM, and AWR are RAC-aware and RAC-friendly
- ▶ OEM 10g Database Control has good tuning information

RAC Backup & Recovery

- ▶ Multiple threads of redo means media recovery is more complicated
- ▶ There is just one database (often useful to remember in recovery situations)
- ▶ A cluster filesystem can help simplify the backup & recovery processes
- ▶ If using ASM, RMAN is your best friend. ASMCMD does not currently offer backup capabilities
- ▶ Test, test, test

New RAC DBA Recommendations

- ▶ Read and **understand** the concepts guide (not RAC-specific, but a good foundation for understanding RAC)
- ▶ RAC is much easier to debug and troubleshoot if you understand concepts, not procedures
- ▶ Metalink and OTN have several “how to” articles on RAC and they are excellent guides for beginners
- ▶ Take advantage of the RAC SIG and other online sources for information. The pool of RAC knowledge online is growing quickly.

RAC Recommendations for Managers

- ▶ "Grid is not RAC, RAC is not grid."
<http://tkyte.blogspot.com/2006/02/so-what-was-answer-part-iii.html>
- ▶ RAC requires additional DBA training (or at least training time)
- ▶ It is dangerous to have a production-only RAC environment without a non-production environment

RAC and Vendor-Provided Applications

- ▶ The most important mentionable here is that vendor-provided applications usually certify RAC as a separate database platform
- ▶ Don't presume that when a vendor supports Oracle 10.1.0.4.0, that they also support RAC 10.1.0.4.0
- ▶ RAC is specifically certified for PeopleSoft, SAP, Oracle EBS, Siebel, and many other applications as well
- ▶ Some vendors require additional patching

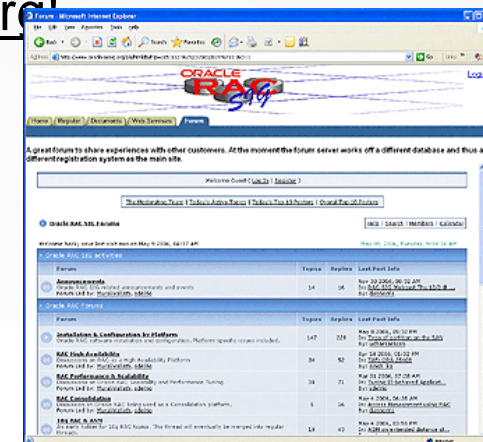
High-Availability Alternatives to RAC

- ▶ Other clusters providing high(er) availability than standalone, single-instance database servers:
 - ▶ Microsoft Cluster Server (w/ Oracle Failsafe)
 - ▶ Veritas Cluster Server
 - ▶ PolyServe Matrix Server (now Novell's)
 - ▶ Red Hat Cluster Server
 - ▶ HP MC ServiceGuard
 - ▶ IBM HACMP

RAC SIG Events

- ▶ See www.oracleacrsig.org for details
- ▶ Monday, Apr 16, 4:45p, RAC SIG Customer Panel
- ▶ Tuesday, Apr 17, 12:15p, RAC SIG Expert Panel
- ▶ Wednesday, Apr 18, 12:15p, RAC SIG Expert Panel
- ▶ More live webcasts on our website!

▶ Join the RAC SIG at www.oracleacrsig.org



ITConvergence

Be there!



April 15 - 19, 2007
Mandalay Bay Resort and Casino
Las Vegas, Nevada



Thank You!

Dan Norris

Technology Services Practice Manager

dnorris@itconvergence.com

800-675-0032 x2395

Legal

The information contained herein should be deemed reliable but not guaranteed. The author has made every attempt to provide current and accurate information. If you have any comments or suggestions, please contact the author at:

dnorris@itconvergence.com

You may request redistribution permission from
dnorris@itconvergence.com.