



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

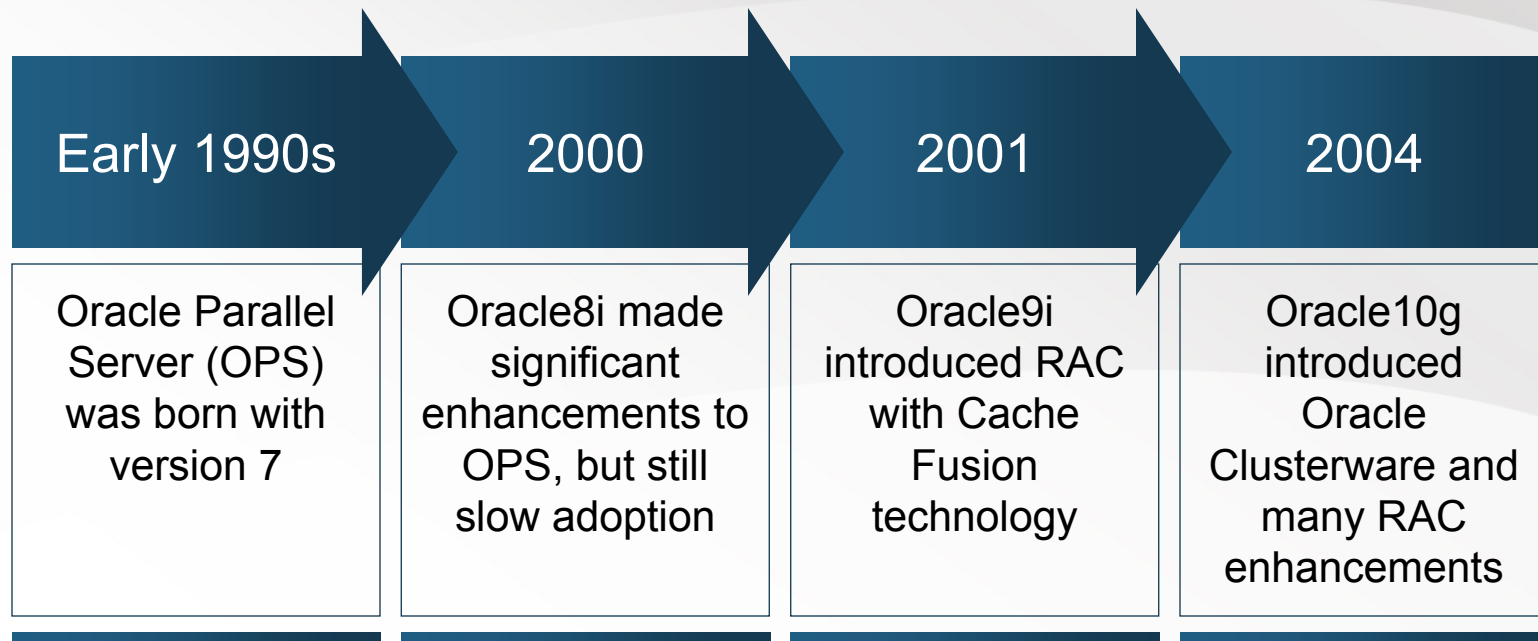
Next Steps & References

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

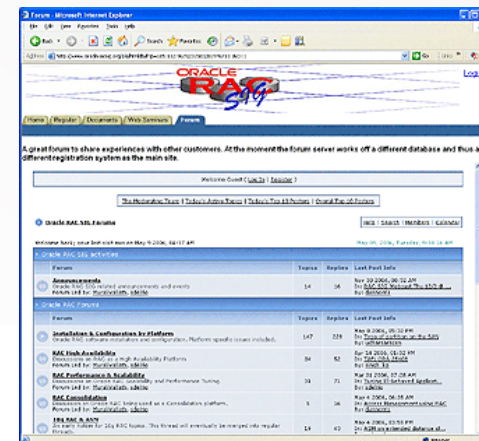
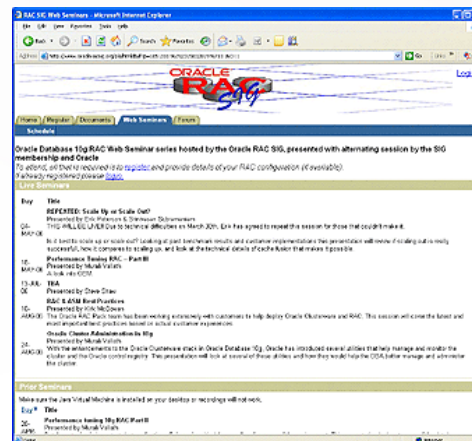
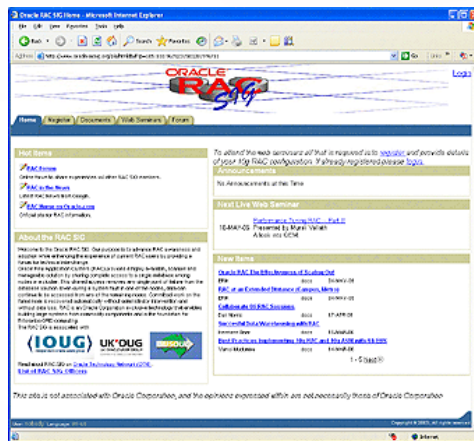
- ▶ Other clusters providing high(er) availability than standalone, single-instance database servers:
 - ▶ Microsoft Cluster Server (w/ Oracle Failsafe)
 - ▶ Veritas Cluster Server
 - ▶ PolyServe Matrix Server (now Novell's)
 - ▶ Red Hat Cluster Server
 - ▶ HP MC ServiceGuard
 - ▶ IBM HACMP

RAC SIG Events


▶ See www.oracleracsig.org for details

- ▶ Live Webcasts
- ▶ Online Forums
- ▶ Great Document Repository

▶ Join the RAC SIG at www.oracleracsig.org!



Save the Date!



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

Submit to present for the IOUG!

- ▶ Share your expertise with the greater Oracle community. Solidify your reputation as an Oracle expert! The IOUG is looking for presentations in the following tracks: **Architecture, Database, Development, and Middleware.**
- ▶ Submit your abstracts no later than **November 17.**
- ▶ If selected, you will receive a FREE COLLABORATE 07 conference registration, industry and peer recognition and much more!

Submit today at www.ioug.org



COLLABORATE07
Technology and Applications Forum for the Oracle Community



Sign Out

- Home
- Content Catalog
- Exhibitor Catalog
- Demo Catalog
- Updates
- My Schedule
- Session Surveys**
- Conferences

The following sessions are available for survey. Please click on session title to take session survey.

ID	Session Title	Speakers
S281248	Oracle Real Application Clusters Performance Experts Reveal All	Barb Lundhild , Oracle; Michael Zoll , Oracle
S281696	Identity Management Implementation Best Practices: Real-World Lessons	Harish Gaur , Oracle; Aaron Johnson , Aegon USA; Dave Bennett , Nulli Secundus; Michael Brengs , Optimal IDM; John Donovan , NetApp
	Oracle Real Application Clusters	Su Tang , Oracle; Arthur Britz , Absa; Kirti Bhansuhalia , ...



Thank You!



RAC For Beginners: The Basics

An IT Convergence presentation by Dan Norris

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

Only Oracle OpenWorld has been granted permission to reprint and distribute this presentation. Others may request redistribution permission from dnorris@itconvergence.com.