

THE EXPERT'S VOICE® IN ORACLE

Pro Oracle Database 10g RAC on Linux

Installation, Administration, and Performance

*Work with Oracle Database 10g RAC on Linux
to build robust grid-computing environments*

Julian Dyke and Steve Shaw

Apress®

Contents

About the Authors	xxi
About the Technical Reviewers	xxii
Acknowledgments	xxiii
Introduction	xxv
CHAPTER 1 Introduction to RAC	1
What Is Real Application Clusters?	1
Clustered Database Concepts	2
Cluster Terminology	2
High Availability	4
Scalability	5
Disaster Recovery	7
Why Deploy RAC?	7
Cost of Ownership Advantages and Benefits	8
High Availability Advantages and Disadvantages	10
Scalability Advantages and Disadvantages	12
Manageability Advantages and Disadvantages	14
Transparency Advantages and Disadvantages	15
Alternatives to RAC	15
Single Instance Databases	16
Data Guard	16
Third-Party Clustering Solutions	16
Differences Between Single-Instance and RAC Databases	16
Differences Between Oracle 9i and Oracle 10g	17
History of Real Application Clusters	19
Summary	20
CHAPTER 2 Introduction to Linux	21
History of Linux	21
UNIX	21
Free Software	22
Linux	23
Open Source	24

Oracle, Open Source, and Linux	25
Unbreakable Linux	25
Red Hat Enterprise Linux	26
SUSE Linux Enterprise Server	27
Asianux	29
Summary	29

CHAPTER 3

RAC Concepts	31
Oracle Clusterware	31
Oracle Cluster Registry	32
Voting Disk	33
Node Applications	33
Listener	33
Oracle Notification Service (ONS)	33
Virtual IP (VIP)	34
Global Services Daemon (GSD)	34
Database Structure	34
Datafiles	34
Control Files	35
Online Redo Log Files	35
Undo Tablespaces	36
Archived Logs	36
Recovery Area	36
Instances	36
Global Cache Service (GCS)	37
Global Enqueue Service (GES)	37
Parameters	37
Initialization Parameter File (PFILE)	38
Server Parameter File (SPFILE)	39
Global Parameters	40
Instance-Specific Parameters	40
RAC-Specific Parameters	40
Background Processes	42
RDBMS Instance Background Processes	42
ASM Instance Background Processes	47
Workload Management	48
Services	48
Connection Load Balancing	49
Transparent Application Failover (TAF)	49
Administration	50
Enterprise Manager	50
Server Control Utility (SRVCTL)	51
SQL*Plus	51
CRSCTL	51

Monitoring	51
Operating System Utilities	51
Enterprise Manager	52
Automatic Workload Repository (AWR)	52
Automatic Database Diagnostic Monitor (ADDM)	52
Active Session History (ASH)	52
Dynamic Performance Views	52
Statistics	53
Tools and Utilities	53
Cluster Verification Utility	53
Oracle Universal Installer (OUI)	53
Virtual IP Configuration Assistant (VIPCA)	54
Database Creation Assistant (DBCA)	54
Database Upgrade Assistant (DBUA)	54
Network Configuration Assistant (NETCA)	54
Oracle Interface Configuration Tool (OIFCFG)	54
OCR Configuration Tool (OCRCONFIG)	55
Backup and Recovery	55
Recovery Manager (RMAN)	55
Storage-Based Backup and Recovery	55
Disaster Recovery	55
Data Guard	56
Stretch Clusters	57
Hardware Disaster Recovery Solutions	57
Summary	57

CHAPTER 4	RAC Design	59
	Business Requirements	59
	Technical Requirements	59
	Business Constraints	60
	Upgrade Policy	61
	Architecture	62
	Development and Test Systems	62
	Disaster Recovery	62
	Stretch Clusters	63
	Reporting Systems	64
	Spare Servers	65
	Storage Requirements	65
	ASM and Cluster File Systems	65
	Shared Oracle Homes vs. Local Oracle Homes	66
	Backup and Recovery	66
	Archived Redo Logs	68
	Hardware Requirements	70

	Application Design	70
	Bind Variables	70
	Sequences	70
	Database Design	72
	Cursor Sharing	72
	Optimizer Statistics	72
	Histograms	73
	Dynamic Sampling	73
	System Statistics	73
	Locally Managed Tablespaces	74
	Automatic Segment Space Management (ASSM)	74
	Reverse Key Indexes	75
	Partitioning	76
	Summary	81
CHAPTER 5	Benchmarking Concepts	83
	Load Testing and Benchmarks	83
	TPC Benchmarks	86
	The TPC-C Benchmark	86
	The TPC-H Benchmark	91
	Summary	98
CHAPTER 6	Hardware	99
	Oracle Availability	100
	Server Architecture	101
	Processors	101
	Memory	112
	Additional Platform Features	115
	Network and Private Interconnect	117
	Network I/O	118
	Private Interconnect Selection	119
	Fully Redundant Ethernet Interconnects	121
	Summary	124
CHAPTER 7	Storage	125
	RAC I/O Characteristics	125
	Read Activity	125
	Write Activity	126
	Forced Reads and Writes	128
	Asynchronous I/O	128
	Direct I/O	129

Storage Fundamentals	130
Disk Drive Performance	130
RAID	131
RAID 0 Striping	131
RAID 1 Mirroring	132
RAID 10 Striped Mirrors	133
RAID 0+1 Mirrored Stripes	133
RAID 5	134
Storage Cache	135
RAID Summary	136
Intelligent Storage	136
Storage Protocols for Linux	138
SCSI	139
Fibre Channel	140
iSCSI	141
S-ATA	142
Using Block-Based Storage	143
Linux I/O Scheduling	143
NAS	144
SAN and NAS Compared	144
SAN Storage Example	146
PowerPath	147
HBA and Driver Configuration	149
Fibre Channel Switch Configuration	151
EMC CLARiiON Configuration	152
NFS and iSCSI Storage Example	158
NetApp Filer RAID Configurations	158
Aggregates and Volumes	159
NFS	160
Snapshot Backup and Recovery	164
iSCSI	166
Summary	170
CHAPTER 8 Oracle Cluster File System (OCFS)	171
Cluster File Systems and OCFS	171
OCFS Version 1	172
OCFS Design Considerations	173
Installation and Configuration	175
Tools and Utilities	179
Debugging and Tracing OCFS	186
Operating System Utilities	188

OCFS Version 2 (OCFS2)	188
Installation and Configuration	188
Tools and Utilities	197
Summary	200
CHAPTER 9 Automatic Storage Management (ASM)	201
Introducing ASM	201
ASM Installation	202
configure	203
start	204
stop	204
status	204
createdisk	205
scandisks	205
listdisks	205
querydisk	205
deletedisk	205
ASM Instance Configuration	206
Adding an ASM Instance Using DBCA	206
Adding an ASM Instance Manually	210
Deleting an ASM Instance Using DBCA	211
Deleting an ASM Instance Manually	211
Re-Creating the ASM Instance	212
Using Files Instead of Devices	212
ASM Components	212
ASM Instance	213
ASM Files	217
ASM Administration	220
ASM Administration Using DBCA	221
ASM Administration Using SQL*Plus	221
ASM Administration Using SRVCTL	221
ASM Administration Using Enterprise Manager	222
ASM Administration Using FTP	223
ASM Administration Using ASMCMD	224
Monitoring ASM	231
Dynamic Performance Views	231
Fixed Views	231
Summary	233
CHAPTER 10 Installing Linux	235
Linux Software Selection	235
Hardware Requirements	236
Networking Requirements	236

Installing Red Hat Enterprise Linux 4 AS	237
Starting the Installation	237
Installation Media Check	238
Anaconda Installation	238
Language Selection	238
Keyboard Configuration	238
Upgrade Examine	238
Disk Partitioning	239
Boot Loader Configuration	243
Network Configuration	243
Firewall Configuration	244
Additional Language Support	244
Time Zone Selection	244
Root Password Configuration	245
Package Installation Defaults	245
Package Group Selection	245
Start Installation	247
Installing Packages	247
Installation Complete	247
Welcome	247
License Agreement	247
Date and Time	248
Display	248
Red Hat Login	249
Why Register?	249
System User	249
Additional CDs	249
Finish Setup	249
Manual Package Installation	249
Installing SUSE Linux Enterprise Server	250
Starting the Installation	250
License Agreement	251
Select Your Language	251
Previous Installation	251
Installation Settings	251
Suggested Partitioning	251
Preparing Hard Disk	251
Expert Partitioner	253
Software Selection	253
Time Zone	255
Warning	256
Root Password	256
Network Configuration	256
Test Internet Connection	257

Service Configuration	257
User Authentication Method	257
Add a New Local User	257
Release Notes	258
Summary	258
CHAPTER 11	
Configuring and Verifying Linux	259
Operating System and RPM Package Checks	259
Hostnames and Name Resolution	260
NTP	261
Hangcheck-Timer	263
Kernel Parameters	264
Shared Memory	265
Semaphores	268
Network	269
Open Files	271
Kernel Parameter Example	271
Oracle User Configuration	272
Creating the dba and oinstall Groups	272
Creating the oracle User	272
Setting the Password for the oracle User	273
Setting Environment Variables	273
Creating the Oracle Software Directories	278
Setting Security Limits	278
Secure Shell Configuration	279
Shared Storage Configuration	281
Partitioning	281
Preparing the Oracle Clusterware Partitions	285
Channel Bonding	287
Cluster Verification	290
Java Runtime Environment	291
Syntax	292
Stage Checks	292
Component Checks	295
Summary	299
CHAPTER 12	
Installing and Configuring Oracle Software	301
Preparing to Install	301
Installation Media	302
Oracle Clusterware	303
Preinstallation Tasks	303
Verifying the Configuration	304

Running the Installer	305
Verifying the Oracle Clusterware Installation	314
Oracle Database Software Installation	315
Verifying the Configuration	315
Running the Installer	317
Troubleshooting	320
Cluster Verification Utility	321
Common Problems	321
Other Information Sources	322
Configuration Files	323
Inventory	323
OCR	324
/etc/inittab	324
Deinstallation	324
Summary	325
CHAPTER 13	Creating and Configuring Databases
	327
Using DBCA As a GUI Tool	327
Configuring ASM	329
Creating a Database	332
Deleting a Database	345
Creating a Template	345
Managing Instances	349
Managing Services	349
Using DBCA at the Command Line	350
Creating a Database Using Scripts	351
Primary node	352
Secondary Nodes	360
Running the Scripts	361
Summary	362
CHAPTER 14	Installing and Configuring Enterprise Manager
	363
Manageability Infrastructure	364
Database Control	365
Grid Control	366
Planning	367
Preparation	371
Management Service Installation and Configuration	372
Starting and Stopping the Management Service	380
Management Agent Installation	381
Summary	384

CHAPTER 15	Linux Administration	385
	Run Levels	385
	Services	387
	Terminal Sessions	388
	Manual Pages	388
	Bash Shell	389
	Commands	392
	Packages	393
	Kernel Parameters	394
	Swap Space and File Systems	395
	Swap Space	395
	General File Systems	396
	NFS	397
	Booting the System	398
	Bootloaders	398
	Initial Ramdisk Image	399
	Log Files, Tracing, and Debugging	399
	Log Files	400
	Tracing	400
	Debugging	400
	Summary	400
CHAPTER 16	Oracle Clusterware	401
	Oracle Clusterware Components	401
	Cluster Synchronization Services (CSS)	402
	Cluster Ready Services (CRS)	402
	Event Manager (EVM)	402
	Oracle Clusterware Files	402
	Oracle Cluster Registry (OCR)	402
	Voting Disk	403
	Oracle Clusterware High Availability Framework	403
	Oracle Clusterware Installation and Configuration	404
	Virtual Internet Protocol Configuration Assistant (VIPCA)	404
	OIFCFG	407
	Administering Oracle Clusterware	408
	CRSCTL	408
	CRS_STAT	413
	OCRCONFIG	416
	OCRCHECK	421
	OCRDUMP	422
	Administering Voting Disks	422

	Oracle Clusterware Logging	423
	Summary	423
CHAPTER 17	Application Development	425
	RAC-Specific Development Considerations	425
	Instances and Database Services	425
	Multiple SGAs	426
	Local and Shared Storage	427
	Node Affinity	428
	Transparent Application Failover	430
	Fast Application Notification	437
	Oracle Notification Service	439
	Server-Side Callouts	441
	Fast Connection Failover	442
	Oracle Clusterware High Availability Framework	446
	Example Application	447
	Implementing the HA Framework	452
	HA Framework Commands	454
	Summary	463
CHAPTER 18	Using Administrative Tools	465
	Using EM	465
	Starting and Stopping Databases and Instances	467
	Setting Parameters	470
	Performance Options	471
	Administrative Options	472
	Maintenance Options	474
	Using SRVCTL	475
	SRVCTL Syntax	476
	Administering Databases	477
	Administering Instances	480
	Administering Node Applications	483
	Administering the Listener	486
	Using SQL*Plus	487
	Setting the SQL*Plus Prompt	487
	Connecting to a RAC Instance Using SQL*Plus	487
	Starting an Instance	488
	Stopping an Instance	488
	Setting Parameters in an Initialization Parameter File (PFILE)	489
	Setting Parameters in a Server Parameter File (SPFILE)	490
	Summary	492

CHAPTER 19	Workload Management	493
	Workload Distribution	493
	Client-Side Connection Balancing	494
	Server-Side Connection Balancing	494
	Database Services	497
	Implementing Database Services	499
	Administering Database Services	501
	Monitoring Services	511
	Summary	519
CHAPTER 20	Adding and Deleting Nodes and Instances	521
	Running GUI Tools	521
	Adding a Node	522
	Plan the Installation	522
	Install and Configure the New Hardware	523
	Configure the Network	523
	Install the Operating System	523
	Configure the Storage	524
	Install Oracle Clusterware	525
	Install Oracle Database Software	528
	Configure the Listener	530
	Adding an Instance	535
	Deleting an Instance	540
	Deleting a Node	544
	Delete ASM Instance	544
	Delete the Listener Process	544
	Delete the Oracle Database Software	547
	Update Inventories on Remaining Hosts	548
	Remove Node-Specific Interface Configuration	548
	Disable Oracle Clusterware Applications	549
	Delete the Node from the OCR	550
	Delete the Oracle Clusterware Software	550
	Update the Inventories on the Remaining Hosts	551
	Verify Node Deletion Using CLUVFY	552
	Summary	552
CHAPTER 21	Trace and Diagnostics	553
	Trace File Locations	553
	Alert Log	553
	Trace Files	554

DBMS_MONITOR	555
Tracing Applications	556
Tracing Multitier Applications	557
DBA_ENABLED_TRACES	558
ORADEBUG	559
LKDEBUG	563
Events	572
Summary	574

CHAPTER 22 RAC Internals

Global Enqueue Services	576
Background Processes	576
Resources and Enqueues	576
Lock Types	577
Library Cache Locks	577
Row Cache Locks	580
Tracing GES Activity	581
Optimizing Global Enqueues	581
Global Cache Services	581
Global Resource Directory (GRD)	582
Cache Coherency	584
Cache Fusion	584
Cache Fusion Examples	586
Disk Writes	599
System Change Numbers (SCNs)	603
Optimizing the Global Cache	604
Instance Recovery	604
Summary	605

CHAPTER 23 Benchmarking Tools: Hammerora and Swingbench

Hammerora	607
Installation	608
Bespoke Application Load Tests	610
TPC-C Simulation	622
Swingbench	633
Installation and Configuration	633
Calling Circle	635
Developing Swingbench Benchmarks	644
Summary	647

CHAPTER 24	Performance Monitoring	649
	Oracle Performance Monitoring	649
	Performance Monitoring with Enterprise Manager	649
	AWR Reports	652
	Active Session History	660
	Automatic Database Diagnostic Monitor	661
	Performance Monitoring Using SQL*Plus	662
	GV\$ Views	663
	System Statistics	663
	Segment Statistics	663
	Global Cache Services	664
	Global Enqueue Service	669
	Library Cache	669
	Dictionary Cache	670
	Lock Conversions	671
	Linux Performance Monitoring	672
	ps	672
	free	673
	top	675
	vmstat	676
	Summary	681
CHAPTER 25	Backup and Recovery	683
	Backup and Recovery Strategy	683
	Recovery Manager (RMAN)	685
	RMAN Utility	686
	RMAN Repository	689
	Backup Sets	692
	Tags	693
	Parameters	693
	Flashback Recovery Area	698
	Performing Backups	700
	Image Copies	703
	Incremental Backups	705
	Block Change Tracking	706
	Incrementally Updated Backups	707
	Housekeeping	707
	Performing a Restore	710
	Performing Recovery	712
	RAC Considerations	714
	Archived Redo Logs	715
	Summary	716

CHAPTER 26	Disaster Recovery	717
	Oracle Data Guard	718
	Data Protection Modes	719
	Maximum Performance Mode	719
	Maximum Availability Mode	719
	Maximum Protection Mode	719
	Setting the Data Protection Mode	720
	Redo Transport Services	720
	ARCn Background Process	721
	LGWR Background Process	721
	Asynchronous Network I/O (ASYNC)	721
	Synchronous Network I/O (SYNC)	722
	Standby Redo Logs	722
	Log Apply Services	723
	Role Management	723
	Read-Only Mode	723
	Switchover	723
	Failover	724
	RAC and Physical Standby	724
	RMAN	725
	NOLOGGING Attribute	725
	Archive Log Gaps	726
	Initialization Parameters	726
	LOG_ARCHIVE_DEST_n	726
	LOG_ARCHIVE_DEST_STATE_n	727
	LOG_ARCHIVE_START	727
	LOG_ARCHIVE_FORMAT	727
	DB_UNIQUE_NAME	728
	LOG_ARCHIVE_CONFIG	728
	REMOTE_ARCHIVE_ENABLE	728
	STANDBY_ARCHIVE_DEST	728
	STANDBY_FILE_MANAGEMENT	729
	FAL_CLIENT	729
	FAL_SERVER	729
	Creating a Physical Standby Database	729
	Enable Archiving on the Primary Database	730
	Create Password Files on Primary Nodes	731
	Force Logging on the Primary Database	732
	Back Up the Primary Database	732
	Set Parameters on the Primary Database	733
	Create Directories on Standby Nodes	733
	Create Password Files on Standby Nodes	734
	Create Server Parameter File for Standby Database	734

Create Initialization Parameter Files on Standby Nodes	734
Copy the RMAN Backup from Primary to Standby	735
Update /etc/oratab on Standby Nodes	735
Add the Standby Database to the OCR	735
Update the Listener Configuration Files on Standby Nodes	735
Update Oracle Net Configuration Files on All Nodes	736
Set Parameters on the Standby Database	737
Create Standby Database	737
Enable Managed Recovery on the Standby Database	738
Check the Standby Configuration	738
Verify Log Transportation	738
Role Management	739
Read-Only Mode	739
Switchover	739
Failover	740
Summary	741

APPENDIX	Third-Party Clustering Solutions	743
	Clusterware	743
	Third-Party Certification Programs	745
	PolyServe Matrix Server	745
	Central Cluster Management Console	747
	SAN Management Layer	747
	Sophisticated Fencing	747
	Dynamic Volume Manager	748
	CFS	748
	Matrix Server ODM	748
	Summary	751
INDEX		753