



# Oracle High Availability, Disaster Recovery, and Cloud Services

Explore RAC, Data Guard, and  
Cloud Technology

---

Y V Ravi Kumar  
Nassyam Basha  
Krishna Kumar K M  
Bal Mukund Sharma  
Konstantin Kerekovski

**Apress®**

# **Oracle High Availability, Disaster Recovery, and Cloud Services**

**Explore RAC, Data Guard,  
and Cloud Technology**

**Y V Ravi Kumar  
Nassyam Basha  
Krishna Kumar K M  
Bal Mukund Sharma  
Konstantin Kerekovski**

**Apress®**

## ***Oracle High Availability, Disaster Recovery, and Cloud Services***

Y V Ravi Kumar  
Chennai, Tamil Nadu, India

Nassyam Basha  
Riyadh, Saudi Arabia

Krishna Kumar K M  
Bangalore, Karnataka, India

Bal Mukund Sharma  
Belmont, CA, USA

Konstantin Kerekovski  
Pinellas Park, FL, USA

ISBN-13 (pbk): 978-1-4842-4350-3  
<https://doi.org/10.1007/978-1-4842-4351-0>

ISBN-13 (electronic): 978-1-4842-4351-0

Copyright © 2019 by Y V Ravi Kumar, Nassyam Basha, Krishna Kumar K M,  
Bal Mukund Sharma, Konstantin Kerekovski

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr  
Acquisitions Editor: Celestin Suresh John  
Development Editor: Matthew Moodie  
Coordinating Editor: Aditee Mirashi

Cover designed by eStudioCalamar

Cover image designed by Freepik ([www.freepik.com](http://www.freepik.com))

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail [orders-ny@springer-sbm.com](mailto:orders-ny@springer-sbm.com), or visit [www.springeronline.com](http://www.springeronline.com). Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail [rights@apress.com](mailto:rights@apress.com), or visit [www.apress.com/rights-permissions](http://www.apress.com/rights-permissions).

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at [www.apress.com/bulk-sales](http://www.apress.com/bulk-sales).

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at [www.apress.com/978-1-4842-4350-3](http://www.apress.com/978-1-4842-4350-3). For more detailed information, please visit [www.apress.com/source-code](http://www.apress.com/source-code).

Printed on acid-free paper

# Table of Contents

**About the Authors.....xvii**

**About the Technical Reviewer .....xxi**

**Acknowledgments .....xxiii**

**Chapter 1: Introduction to High Availability and Disaster Recovery  
with Cloud Technology..... 1**

    Why High Availability? ..... 3

    Best Practices for a High Availability Environment ..... 7

    Data Guard ..... 10

        Maximum Performance ..... 10

        Maximum Protection ..... 10

        Maximum Availability..... 10

    Importance of HA and DR..... 11

    Evolution of Cloud Technologies..... 11

    Advantages of Having a Cloud Environment ..... 15

        Centralized Location ..... 15

        Setup Cost ..... 15

        Operational Cost ..... 16

        Duration to Create the Setup ..... 17

        Scalability ..... 17

        Optimization and Modification..... 17

        Ease of Handling..... 18

        Migration ..... 18

        Security ..... 18

        Variety of Applications ..... 19

        Free from Operations Headaches: Licensing Expiration ..... 19

TABLE OF CONTENTS

Differences Between On-Premises, IaaS, PaaS, and SaaS.....	19
On-Premise.....	20
IaaS.....	20
PaaS .....	21
SaaS .....	21
Summary.....	21
<b>Chapter 2: Oracle Active Data Guard with Cloud Services.....</b>	<b>23</b>
Types of DR Databases .....	24
Physical Standby Databases .....	24
Logical Standby Databases .....	25
DR Activity Attributes .....	26
Log Destination.....	26
ARCH.....	26
Log Writer Process .....	27
LNS .....	27
RFS .....	27
Behavior of the DR Activity.....	28
Maximum Protection .....	28
Maximum Availability.....	28
Maximum Performance .....	28
Differences Between Modes .....	29
Switchover .....	29
Failover .....	32
Data Guard Broker.....	33
DG Setup in the Cloud .....	35
DBaaS Tool Update .....	35
dbaascli Commands to Data Guard .....	38
dbaascli Data Guard Switchover .....	40
dbaascli Data Guard Failover.....	42
dbaascli Data Guard Reinstate .....	44

Data Guard Switchover Through the GUI.....	46
Data Guard Failover Through the GUI .....	50
Hybrid DR .....	53
Automatic Approach .....	54
Manual Approach.....	62
Backup Utility for the Primary Database in Cloud .....	68
Backing Up to Cloud Storage .....	68
Backing Up to Disk and Cloud Storage .....	69
Managing DR Instances Through DBaaS Monitor/Oracle SQL Developer Web.....	71
Enable the Login Schema .....	73
Using the Schema .....	77
Using the GUI .....	77
Summary.....	89
<b>Chapter 3: High Availability for Cloud Control 13c .....</b>	<b>91</b>
OEM Agent High Availability Setup .....	91
Directory Structures and Their Purpose .....	91
Checking the Current Agent High Availability Configuration .....	92
Changing the Agent's High Availability Configuration .....	93
HA and MAA Best Practices for the OEM 13c Repository Database.....	95
Creating the Repository in a 12.2 Primary Database on a RAC Database .....	95
Setting Recommended Parameters for the Primary Database.....	96
Creating a Data Guard Standby of the OEMREPO Database with OEM 13.2 Cloud Control .....	97
OMS High Availability Setup.....	129
Adding a Second OMS Server.....	130
Configuring the Application Servers for Disaster Recovery.....	166
OEM Always-On Monitoring .....	178
Configuring AOM.....	178
Starting AOM .....	185

## TABLE OF CONTENTS

Bug 25132376 .....	185
Enabling Notifications.....	186
Summary.....	186
<b>Chapter 4: GI and RAC Options .....</b>	<b>187</b>
ASM Filter Driver.....	187
Configuring the ASM Filter Driver .....	188
Manual Configuration of the ASM Filter Driver .....	189
ACFS Snapshot Replication.....	194
ACFS Creation on the Source Cluster .....	194
ACFS Creation on the Standby.....	196
Public Keys .....	198
Initiating Replication on the Standby.....	201
Initiating Replication on the Primary .....	204
Testing the Replication .....	205
HANFS Filesystems (NFS v3).....	207
HANFS (NFS v3) Server-Side Setup .....	207
HANFS (NFS v3) Client-Side Setup .....	215
Oracle Flex Clusters .....	217
Adding a Leaf Node to a Freshly Installed Cluster .....	219
Flex Redundancy Disk groups.....	232
Creating a Flex Disk group .....	232
Filegroup .....	234
Rapid Home Provisioning .....	240
Configuring the RHP Server Resource .....	241
Using the Existing RDBMS and GI Installations on the RHP Server as Golden Images .....	243
Provisioning a Working Copy of the GI Home.....	244
ORACLE_HOME for 12.2 GI Provisioning.....	248
Converting the Target Cluster into an RHP Cluster .....	250
Provisioning a Working Copy of the RDBMS Home.....	251

Provisioning a Database on the Client Cluster.....	253
Database SYS/SYSTEM Passwords .....	255
Summary.....	255
<b>Chapter 5: Troubleshooting and Conducting Health Checks for High Availability .....</b>	<b>257</b>
Cluster Verification Utility.....	257
Software Home Checks .....	258
File System Space Checks .....	260
Cluster Verification Utility Health Checks.....	262
Database Health Checks.....	262
Cluster Health Checks .....	270
Cluster Verification Utility Baselines.....	273
Orachk.....	276
Upgrading Orachk.....	277
Orachk/Exachk and Oracle RESTful Data Services.....	277
Recommended Settings for Orachk Daemons.....	279
Trace File Analyzer .....	281
Upgrading TFA to Include the Support Tools Bundle.....	281
Using Trace File Analyzer to Manage Logfiles .....	283
New Health Check and Troubleshooting Features in 12.2.....	290
Cluster Health Advisor.....	290
Summary.....	292
<b>Chapter 6: Best Practices in Oracle Data Guard with Tips and Techniques.....</b>	<b>293</b>
Creating a Standby Database Using DBCA.....	293
Expected Error Messages per Limitations.....	293
Example of a DBCA-Created Standby .....	294
Alternative Ways to Supply the SYS Password.....	295
Diagnostic Pack on Active Data Guard .....	295
The SYS\$UMF User .....	296
Creating the DB Links.....	297



## TABLE OF CONTENTS

Configuring the Remote Management Framework Topology.....	298
Viewing the Topology information .....	299
Taking AWR Snapshots on ADG Databases.....	300
Creating AWR Reports for the Standby Database .....	300
SQL Tuning Advisor in ADG Databases .....	302
Creating a Database Link for the SQL Tuning Advisor .....	302
Tuning a Query.....	303
RMAN Support for NONLOGGED BLOCK recovery.....	306
Creating a Table with the NOLOGGING Clause .....	306
The Effect of NOLOGGING Operations on the Physical Standby.....	308
Fixing the Problem.....	309
Data Guard Support for Multiple Observers .....	311
Starting Multiple Observers.....	312
Determining the Master Observer .....	312
Manually Changing the Master Observer .....	314
Data Guard Multi-instance Apply .....	314
Broker Example .....	315
SQLPLUS Example .....	315
Selectively Applying Redo for PDBs in Data Guard .....	315
Effects of the ENABLED_PDBS_ON_STANDBY Parameter .....	316
Disabling Redo Apply for an Existing PDB .....	319
Data Guard Database Compare .....	319
Using DBMS_DBCOMP to Detect Nologging Operations.....	320
Contents of the dbcomp.txt File.....	321
Password File Change Synchronization .....	323
Checksums Prior to the SYS Password Change .....	324
Changing the Password.....	324
Checksums After to the SYS Password Change .....	324
In-Memory Columnar Store for ADG Instances.....	325
Role Transition Connection Preservation.....	326

Establishing the Connection .....	326
Performing the Role Transition .....	327
Checking the Existing Session .....	327
PDB Migrate .....	328
ALTCDB Broker Configuration .....	329
Broker Configuration .....	329
The Migration .....	330
Source Database Alert Log Entry .....	331
Target Database Alert Log Entry .....	331
State of ALTCDB Post-Migration .....	332
Summary.....	333
<b>Chapter 7: Oracle Sharding.....</b>	<b>335</b>
Components of the Sharding Architecture .....	337
Methods of Sharding.....	339
CREATE TABLE Syntax in SDB.....	340
Installation and Configuration .....	341
Add the DNS or Hosts File Entries .....	343
Add a Disk for Database Storage, Create LVM on It, and Mount It.....	343
Calculate the Size of /dev/shm.....	344
Disable iptables and Firewalls on All Hosts.....	344
Disable IPv6.....	345
Install the Oracle RDBMS Pre-installation RPM on All Hosts .....	345
Create the Necessary Directories on Each Sharded Database (shdb1 and shdb2) and Catalog/GSM Nodes (shcat) .....	345
Install the Database Software in Silent Mode on Each Sharded Database (shdb[1-4]) and Catalog/GSM Node (shcat).....	346
Create a Noncontainer Database Using DBCA and a Local Listener Using NETCA on the Catalog/GSM Node Only .....	347
Install the GSM Software on the Catalog/GSM Node (shcat) .....	352
Create a Shard Catalog Schema in the SHRCAT Database in the SHCAT Node .....	357
Start the Scheduler Agents on Sharding Nodes .....	360

TABLE OF CONTENTS

Using GSDCTL to Create the Shard Catalog .....	363
Create the Shard Group, Director, and Shards .....	365
Deploy Shards Using GSDCTL .....	368
Create the Global Service .....	373
Disable the FSFO-Observer (Optional) .....	374
Create Tablespaces, Schemas, and Tables .....	378
Testing the Methodology .....	383
Troubleshooting .....	393
High Availability with Oracle Sharded Databases .....	396
Backup and Recovery for SDBs .....	397
Summary.....	398
<b>Chapter 8: Oracle Engineered Systems for High Availability .....</b>	<b>399</b>
Oracle Database Appliance Performance Optimization.....	402
Technical Specifications of the ODA x7-2 Family.....	402
Oracle Database Appliance x7-2-HA Architecture and Storage Specifications.....	403
ODA Deployment Models.....	406
ODA Bare-Metal Components.....	406
ODA Virtualized Components .....	407
Options for ODA Virtualized Environments.....	409
Deployment in ODA .....	410
The Flow of Deployment in ODA .....	411
ODA Deployment in a Bare-Metal Environment.....	411
ODA Deployment in a Virtualized Environment.....	418
Purging LOGs in ODA.....	423
Management and Diagnostics of ODA.....	424
Auto Service Request (ASR) .....	424
ODA ORACHK .....	424
OEM ODA Plug-In .....	425
Diagnostics Collection .....	425
odaadmcli manage diagcollect.....	425

System Check for ODA Environments.....	426
Details of the ODA Assessment Report.....	429
ODA Patching .....	429
Prerequisites .....	431
Applying the 12.1.2.10.0 Patch.....	431
Using ODACLI Commands to Patch a System and Upgrade Databases.....	444
Updating the Repository of Grid Infrastructure and Databases Using the ODACLI Command .....	446
Updating the Databases to the Latest PSU .....	449
Upgrading Database Patchsets or to Major Releases.....	450
ODACLI Commands .....	451
Commands for Creating Databases .....	451
Appliance Commands.....	451
CPU Core Commands.....	452
Database Commands.....	452
DB Home Commands.....	453
DB Storage Commands.....	453
ODAADMCLI Commands.....	454
Storage Commands .....	454
Hardware Monitoring.....	454
Resize the ACFS File System in ODA.....	455
Nondatabase Volumes .....	456
Database Volumes .....	456
Alert Received from EM Cloud Control .....	457
Using the OAKCLI Command Utility with Databases .....	458
Creating a Single-Instance Database .....	459
Creating a RAC Database.....	460
Upgrading the Database from Oracle 11g R2 (11.2.0.4.0) to Oracle 12c R2 (12.1.0.2.0) .....	462
Creating a Snapshot Database .....	463
Dropping a Snapshot Database .....	465

TABLE OF CONTENTS

Dropping a Database .....	465
Dismounting an ACFS Filesystem .....	466
gDBClone Package for Oracle Engineered Systems .....	471
gDBClone Package RPM Installation .....	471
Using gDBClone .....	474
Oracle Exadata Database Machine .....	475
Oracle Exadata Hardware Generation Advances from v1 to X7 .....	476
Oracle Exadata Database Machine x7-2 Component Specifications.....	476
Oracle Exadata Unique Smart Database Software Features.....	478
Enhancements and Features of Oracle Exadata x7-2 .....	479
Oracle Exadata x7 High vs. Normal Redundancy .....	480
Oracle Exadata Deployment Assistant .....	480
M.2 Hot-Swappable Boot Drives from Oracle Exadata x7-2.....	481
Finding the Model of the Oracle Exadata.....	482
Information on Oracle Exadata Database Machine Models.....	483
Improved Oracle Database Deployment Assistant.....	483
Oracle Exadata x7: Unique Fault-Tolerant In-Memory Databases .....	483
Quorum Disks in Oracle Exadata .....	484
Checking the Volumes (/ and /u01) in All Oracle Exadata Compute Nodes.....	485
Configuring Hugepages in Oracle Exadata .....	487
Generating Cell Events on Exadata Storage Server .....	488
Configuring SSH Equivalence from Oracle Exadata Compute Nodes to All Exadata Storage Servers and IB Nodes for the Root User .....	490
Disk Scrubbing Feature for Oracle Exadata Storage Server .....	492
Deleting Trace Files from Oracle Exadata Storage Servers .....	498
Oracle Exadata x7-2: Do-Not-Service LED Property .....	500
Online Flash Disk Replacement in Oracle Exadata x7 .....	502
RAID 0 on Flash Disks in Oracle Exadata x7 .....	503
Oracle Exadata x7: Disk Controller CacheVault Flash Module Monitoring.....	504
Exafusion Parameter on the Oracle Exadata Database Machine.....	505

Oracle Exadata: Flash Cache Compression .....	507
Flash Card Changes in Oracle Exadata x4-2/x6-2 .....	508
Oracle Exadata: Patch Backups Taking Up a Lot of Disk Space.....	512
Snap Clone Feature in Oracle Exadata .....	513
Summary.....	515
<b>Chapter 9: Oracle Cloud Overview .....</b>	<b>517</b>
General Cloud Concepts: What Is a Private Cloud? .....	518
What Is a Public Cloud? .....	519
What Is a Hybrid Cloud? .....	520
What Is a Community Cloud?.....	520
What Is the Oracle Public Cloud? .....	520
What Is Oracle's Software as a Service? .....	520
What Is Oracle's Platform as a Service? .....	520
What Is Oracle's Infrastructure as a Service? .....	521
Oracle Infrastructure as a Service (IaaS).....	522
Oracle Cloud Infrastructure Concepts.....	523
Oracle Cloud Infrastructure Key Concepts.....	524
Signing Up for Oracle Cloud Infrastructure Services .....	528
Signing In to the Console .....	530
Virtual Cloud Network .....	530
Subnet .....	531
Internet Gateway .....	531
Dynamic Routing Gateway.....	531
Security List.....	532
Route Table .....	533
Creating a Compute Instance in the Public Cloud .....	537
Oracle Database Cloud Service Options.....	552
Database Standard Edition .....	553
Database Enterprise Edition .....	553

TABLE OF CONTENTS

Database Enterprise Edition High Performance .....	553
Database Enterprise Edition Extreme Performance .....	554
Database Cloud Deployment Scenarios .....	554
Bronze Category .....	554
Silver Category .....	555
Gold Category .....	555
Platinum Category .....	555
Significance of a Maximum Availability Architecture .....	556
Disaster Recovery Deployments with the Oracle Public Cloud .....	557
Moving an On-Premise Database to the Oracle Database Cloud Service .....	571
Summary.....	574
<b>Chapter 10: Exadata Cloud Provisioning .....</b>	<b>575</b>
Exadata Cloud Service on Oracle Cloud Infrastructure .....	576
Exadata Cloud Service: Provisioning.....	577
Using the Oracle Cloud Infrastructure Dashboard .....	578
Recommended Configuration for Security Rules.....	588
Creating the Exadata Database System .....	594
Summary.....	599
<b>Chapter 11: Migrating Databases to Exadata Cloud Service Using DIPC.....</b>	<b>601</b>
Creating a DBCS Service.....	603
Creating a DIPC Service (18.x.x) .....	608
Generating a Secure Shell (SSH) Public/Private Key Pair on Linux.....	611
Generate SSH Key Pair on Windows Using the PuTTYgen Program .....	612
DIPC Remote Agent Deployment.....	616
Configuring a Remote Agent for the On-Premise Host .....	619
Configuration On-Premise .....	641
Setting Up and Starting the SOCKS5 Proxy Tunnel for the PUMP Process .....	645
Steps to Access DIPC VM and DBCS/EXACS.....	646
Connecting to the DIPC VM via VNC.....	650
Steps to connect to DBCS/EXACS from the DIPC VM.....	651

Configuring the DIPC VM.....	652
Starting the OGG Pump and Replicate Process .....	660
Validating.....	660
Gathering Stats.....	661
Summary.....	662
<b>Chapter 12: Managing Exadata Cloud Service.....</b>	<b>663</b>
Agent Deployment on Exadata Cloud Service .....	665
Creating the Credentials .....	666
Deploying the Oracle Enterprise Manager Cloud Control Agent .....	667
Standard Agent Deployment on ExaCS with a VPN.....	667
Summary.....	668
<b>Index.....</b>	<b>669</b>



# About the Authors



**Yenugula Venkata Ravi Kumar (YVR)** is an Oracle ACE Director and Oracle Certified Master (OCM) with 21+ years of experience in the banking, financial services, and insurance (BFSI) verticals. He has played various roles such as vice president (DBA), senior database architect, and senior specialist production DBA. He is also an OCP in Oracle 8i to 12c and Oracle certified in GoldenGate, RAC, performance tuning, and Oracle engineered systems (Exadata and ODA).

He has written 100+ articles for the Oracle Technology Network (in Spanish, Portuguese, and English), OraWorld, UKOUG, OTech Magazine, and Redgate. He is an Oracle speaker at @OOW, @IOUG, @NYOUG, @OTN, @AIOUG, @Sangam, and @SOUG. He designed, architected, and implemented the core banking system (CBS) database for the central banks of two countries, India and Mahe, Seychelles. He also coauthored *Oracle Database Upgrade and Migration Methods*, published by Apress.



**Nassyam Basha** is a technology expert with 11 years of Oracle Database experience. He holds a master's degree in computer science. In 2014 he became an Oracle Database 11g Certified Master. Nassyam was honored with the title of Oracle ACE Director in 2015 based on his extraordinary contributions to the Oracle community. He was the second Oracle ACE Director in India in the area of database management. He is board member of the Oracle RAC SIG team. He actively participates in Oracle-related forums, such as OTN, Oracle Lists (with the status of Super Hero), and Oracle Support (with the title of Guru), and acts as an OTN moderator. He has deep knowledge in the business fields of finance, retail, health, manufacturing, and consulting.

## ABOUT THE AUTHORS

He started his career with dBase and FoxPro, and participated in several projects with FoxPro and Oracle Database from Oracle 7. He currently works as a database expert with eProseed. Nassyam co-wrote his first book on Oracle Data Guard 11gR2 in 2013 and co-wrote *Oracle Database Upgrade and Migration Methods* in 2017. He has written more than 50 articles and published in OTN, IOUG, Oracle Scene, and various other portals on technologies such as Data Guard, EM 13c, RAC, Oracle Cloud, 12c DB, GDS/GSM and sharding technologies.

In 2012 Nassyam started presenting at various user group conferences, such as OTN Yathra, AIOUG Tech Days, and SANGAM, and appeared at international conferences such as OOW, TROUG, and OUGF. Nassyam maintains an Oracle technology blog ([www.oracle-ckpt.com](http://www.oracle-ckpt.com)). Feel free to contact him at [nassyambasha@gmail.com](mailto:nassyambasha@gmail.com).



**K M Krishna Kumar** is a database administrator with 14 years of experience with Oracle Database (8i to 18c). He has a master's degree in business administration from Anna University. He started his career in 2004 as an Oracle database administrator and has been working with Oracle Corporation for the last eight years. He has expertise in Oracle Database installation, patching, upgrades, migration, Data Guard, RMAN, RAC, backup, and recovery. He is also proficient in Oracle public, private, and hybrid cloud environments. He is an Oracle 11g and 12c certified professional on an Oracle Cloud infrastructure. He has delivered presentations to customers through various channels. He actively participates in Oracle-related forums such as OTN communities. He is a Oracle speaker at @Oracle Open World, @OTN, @AIOUG, and @Sangam. He coauthored the book *Oracle Database Upgrade and Migration Methods*, published by Apress. He is reachable at [kmkittu2005@gmail.com](mailto:kmkittu2005@gmail.com).



**Bal Sharma** works for Oracle America as a product manager of Oracle Cloud infrastructure. He is an Oracle Certified Professional (OCP) with 14 years of experience across various industries such as telecom, healthcare, financial services, and insurance; he has played various roles such as database manager, technical lead DBA, and content architect. He holds several certifications across the Oracle Cloud infrastructure and Oracle 10g/11g/12c, is certified in RAC and Oracle Exadata, and has a deep understanding of information technology and public cloud services. In his recent role, he advises customers on the adoption of cloud platforms, digital transformation, and database technologies such as high availability and engineered systems on the Oracle Cloud infrastructure. He is also setting up a database and infrastructure strategy for the largest banks and telcos in India and the United States and has led various complex implementations at top banks and telcos in India.

Bal is an avid technologist and believes customer satisfaction is of prime importance. His work with customers such as the State Bank of India, HDFC, Idea Cellular, and the Reliance Group is a testimonial for this.

Besides this, Bal enjoys cooking for friends and family.



**Konstantin Kerekovski** is an Oracle professional with more than six years of experience in the financial services industry. He is a member of the Independent Oracle User Group, Suncoast Oracle User Group, and New York Oracle User Group. He has worked as a senior and lead database administrator with production support duties in the financial services industry. He was a speaker at the OOW 17/18 IOUG 17/18, SOUG, and NYOUG meetings. He is an expert at Oracle Real Application Cluster technologies, ASM, GoldenGate, and Oracle 11g and 12c.

# About the Technical Reviewer



**Arup Nanda** has been working in database management for 25 years and counting. With 6 books, more than 700 published articles, and 500 presentations in 22 countries, he is well known in the database area in general and Oracle technologies in particular. Today he is the chief data officer of Priceline.com in the New York area. He blogs at [arup.blogspot.com](http://arup.blogspot.com).

# Acknowledgments

I am grateful to God who gave us all the strength, courage, perseverance, and patience in this sincere and honest attempt of knowledge sharing. This second book of mine would not have been possible without the following people:

Shri Yenugula Venkata Pathi and Smt. Yenugula Krishna Kumari, my parents, instilled in me good thoughts and values. My uncle Shri K L Janardhan Rao (Tallapudi, AP) always encouraged me to work hard and be a step ahead in learning new things.

B. Suresh Kamath (founder of LaserSoft and Patterns), my mentor, my guru, my strength, and my guide, has inspired me for the last 21 years. He is an immensely talented and technically sound individual. He taught me how to be well-read with no compromises. He led by example in being content yet hungry of knowledge. He motivated me to go that extra mile in attempting and experimenting with newer technologies/environments and in being regularly out of my comfort zone.

Anitha Ravi Kumar, my wife, was immensely tolerant with me. “Behind every successful man, there is a good woman,” as they say. I believe she is the embodiment of this well-known belief. Special thanks to my daughter, Sai Hansika, and my son, Sai Theeraz, for giving me time to write a second book.

I would like to thank Arup Nanda for accepting to be the technical reviewer for this book.

I would like to thank two people (Sateesh Prabakaran, SVP, and Puneet Dham, VP) from Raymond James Financial Inc. who extended their support in giving me expertise on the production side of Oracle engineered systems (Oracle Exadata, ODA, and ZFS systems).

Thanks to Mariami Kupatadze (senior solutions architect at FlashGrid Inc.) for excellent support when writing the Oracle sharding chapter.

Special thanks to two people, Jonathan Gennick (assistant editorial director) and Celestin Suresh John (senior manager, editorial acquisition), at Apress for giving me an opportunity to write a second book.

Thank you to the reader for picking up this book. We have attempted to be as simple and straightforward as possible when sharing this knowledge, and we truly believe that it will help you to steadily deep dive into various interesting concepts and procedures.

## ACKNOWLEDGMENTS

I would also like to express my gratitude to the Oracle professionals Murali Vallath, Binay Rath, Mohit Singh (my RAC guru), Nassyam Basha, Krishna Kumar, Konstantin Kerekovski, Jennifer Nicholson (OTN), Bal Sharma, Archana Durai, and all of my LaserSoft colleagues.

**—Y V Ravi Kumar**

I would like to acknowledge all the people who have helped me write this book, both directly and indirectly.

First, I would like to thank my mother and father, Snejana and Roman Kerekovski, for their love, support, and constant encouragement.

Many thanks to Miladin Modrakovic for introducing me to Oracle and teaching me the trade from the time that I was a junior Oracle DBA. I would also like to thank Maryann Marnell for her professional guidance throughout my career.

Lastly, I would like to thank attorney Leonardo B. Perez III for protecting my legal interests.

**—Konstantin Kerekovski**

I would like to thank the Apress team who worked patiently until the end to produce an excellent book. This book is the result of continuous hours of work, and I hope you will find it extremely helpful. I thank the entire team for making it possible.

I would like to express my gratitude to Arup Nanda for the technical review of the entire book.

I would like to convey my sincere thanks to Jaikishan Tada and Vikash Paliseti for their great support, and special thanks to my management team for their support and encouragement.

Above all, I want to thank my wife, Swetha Ramesh, for standing beside me throughout my career and helping to complete this book. Her support and feedback helped me a lot to complete this book. I also want to thank my parents, Mohanram and Lalitha, and the rest of my family members, especially my brother, Ramkumar, who supported and heartened me to complete this book.

Also, I want to thank my friend Koushik Suresh who has been my motivation for continuing to improve my knowledge and move my career forward. I would like to thank my friend Velu Natarajan for his guidance in my DBA carrier.

Finally, last but not least, this book is dedicated to the Oracle database administrators around the world who work 24/7 in production environments.

All kinds of feedback, comments, and ideas are warmly welcomed. Any suggestions regarding the improvement of this book will be acknowledged.

**—Krishna Kumar**

I would like to thank the Apress team, my coauthors, and the technical reviewers for their great support and patience throughout the writing process. I must start by acknowledging the many IT professionals I have had the privilege of working alongside.

To start, to the team at Oracle USA, I would like to convey my sincere thanks to Oracle product manager and my mentor Rohit Rahi for technical guidance and enablement and to the Oracle Technology Sales Consulting Management group (Riz Mithani, Rich Hua, Steve Richardson, Carl Griffin, Tim Cowan, and Prasuj Loganathan) for encouraging me. To the team at Oracle India, thanks to Sandesh Naik and Dinesh Parekh for driving me to be better in life.

I would take a moment to thank my co-captain in life, my wife, my best friend, Suchitra, for standing beside me. Without her I would not have found my voice! I also want to thank my kids, Ishaanvi and Viransh; my parents; my brothers, Abhishek Kumar Suman and Rajesh Kumar Suman; and the rest of my family members.

Throughout my carrier I have been surrounded by an amazing set of mentors and colleagues. All have contributed in a big way, and this list would not be complete without Rashmi Ranjan Panda, Amit Kumar, and Kripa Sindhu. Thank you for being my friends as well as great colleagues.

**—Bal Sharma**

Above and beyond all others, I thank the Almighty and my parents, N.Abdul Aleem and Rahimunnisa.

Without their support, I wouldn't have been able to be what I am today. A special thanks to my family: my wife, Ayesha, and my daughter and my lucky charm, Yashfeen Fathima. I would also like to thank my brother, Nawaz Basha who has supported me all the time. Thanks to my family members Zaheer Ahmed, Farhana, Riyana, Aiman Fathima, Yaseen, Saad, and Maaz.

I would also like to express my gratitude to the Oracle Professionals Chinar Aliyev, Syed Jaffar, my first mentor Shahbaz sir, Emre Baransel, Michale Seberg, Franciso Munez Alvarez, Bjoern Rost, Andy Klock, Simon Pane, Javid Ur Rehman, Asif Momen, Syed Zaheer, and Monowar Mukul, and all the Oracle Community contributors across the world.

I thank my clients and colleagues who provided me with invaluable opportunities to expand my knowledge and shape my career.

## ACKNOWLEDGMENTS

My heartfelt appreciation goes to the technical reviewer of this book, Arup Nanda, for the time he spent reviewing this book. I would like to thank the Apress team for giving us the opportunity to write this book and their support throughout. Last but not least, I would like to give a big thanks to my friends and coauthors for their contribution to the team and great direction.

—**Nassyam Basha**



## CHAPTER 1

# Introduction to High Availability and Disaster Recovery with Cloud Technology

Welcome to the world of high availability and disaster recovery. It is high time for us to learn about database software high availability concepts. Software has become an essential element in today's world. Almost all businesses in every industry have started using software to complete their business tasks more efficiently. In fact, some professions are completely dependent on software; they get all of their tasks done using software. On the other hand, there are some industries such as the manufacturing sector that use software for only a portion of their life cycle. But anyhow all businesses should be trying to utilize software to automate their work as much as possible.

Consider how the banking system has improved using software. In the past we had to visit a bank in person and conduct a transaction with the help of bank staff members. It took a lot of manual effort and time to conduct a transaction. Now, the banking system has been completely reformed and automated with most of its functions being done using software. Their tasks are automated with automatic teller machines (ATMs), which in turn has helped to build high-quality banking products. All banks now have websites that work 24/7, and customers are able to complete banking transactions without the involvement of bank staff members. All this has become possible because of the evolution of software.

Another example is the manufacturing industry where human attention is mandatory at each phase of production work, but the industry uses software for data analysis, management, auditing, and wherever else it is possible.

Now we know about the influence of software in almost all businesses, but software runs on hardware. So, how much can an industry trust the availability of the software it is using? If something goes wrong with the hardware and that affects the software, then how will the business survive? If something happens to the environment where the software and hardware are hosted, then what will happen to the application? These questions insist on the necessity of Hardware & software's *high availability* and *disaster recovery*. With the booming software influence, Achieving high availability and disaster recovery is mandatory for any industry. Having high availability and disaster recovery denotes the capability of running a business continuously without any disruption or irrespective of any damages or planned maintenance in the environment.

Let's discuss an example explaining the necessity of high availability and disaster recovery. Suppose a bank needs to migrate its data to a new data center. We know that migration may require downtime, but the bank cannot afford downtime; that will affect its business a lot. This is not only for banks. In this competitive world, none of the industry can afford downtime to its business. Customers look for vendors that are capable of providing continuous service around the clock. So, the bank has to look for options to do migrations without any downtime.

Along with maintenance tasks, there could be a chance of other, unpredicted disruptions that may affect the software's availability. Natural disasters like earthquakes may affect the site where the hardware is located, and the building that has the physical server in it may collapse. This might be rare, but we cannot completely ignore this factor. When a disaster happens, we will not be able to access the physical servers; hence, the Application is not accessible, and the business cannot continue. That is unacceptable! Our environment design should be robust to handle these kind of situations. To be proactive, we need to set up alternate plans to be able to overcome the impact of an earthquake or any fire or server crash that affects your physical servers located in the data center. Remember when the World Trade Center in the United States was bombed in 2001. The companies located on that campus lost their physical servers, but that did not stop them from running their businesses. They came online and started serving customers within a couple of hours. How was that possible? What did they do to overcome such a situation? We can get answers to all these questions in this book. We will be discussing Oracle Database and its high availability concepts.

In addition, we will find many best practices that can be followed in our industry to make our data safe. This book also covers the Oracle public cloud. Enjoy reading!!!

## Why High Availability?


Before proceeding into high availability, let's discuss what we are trying to achieve with high availability. Businesses in all industries rely on data, be it the software, medical, or insurance industry. Take, for example, a hospital. A patient's record is precious data. In the past, the data was seen just as a record, and it was used for querying purposes and to derive an annual progress report. Now data is not seen like that; it undergoes a data mining process, and many manipulations are done using the mining results. *Data mining* refers to extracting and discovering knowledge from large amounts of data stored in a database. For a hospital, data is required not only to calculate an annual financial sheet but to analyze how it has performed in the past and to predict how it will be doing in the future. This analysis can proactively find out whether any existing policy requires changes to make a better future. In fact, many industries mine their historical data and generate statistics using various algorithms to predict future business opportunities.

Another good example is the telecommunications industry, which mines its historical data and finds out customers' favorite phone plan, customers' calling behavior, ways to improve the quality of the service, and ways to achieve better customer retention and satisfaction. You probably have noticed many phone providers have introduced wallets recently to avail of their service. This is one of the promotions they invented after mining their historical data.

Another best example is e-commerce sites. When we purchase or review any item on a web site, it records the details in the background, and the next time we log into that site, we see items related to our past searches. Customers' favorites are the data.

Having discussed the importance of the data, now we know why the data has to be secured. That means the data has to be highly available in any circumstance.

Let's discuss the high availability options one by one. At first when industry required high availability option, they started looking at backup and restore strategies for their data. If they lose data, they can restore it from a backup. They had invested in creating remote backup storage, creating network connectivity for backup storage, taking backups of data regularly to the storage, and then periodically validating the stored backups. See Figure 1-1.



**Figure 1-1.** *Data backup to tape drive*

Low-cost media tape devices are commonly used as remote storage to store backups for restoring. In addition to being used for restore purposes, the backup is used in the development or preproduction environment for testing purposes.

Even now the backup/restore method is a standard practice that is used across many industries for securing data. But there are some challenges while dealing with backup/restore. Suppose there is a hardware failure or disaster and the current system crashes, then the data is not available, so a backup needs to be restored on the same or different system or in a different data center, based on the conditions. If you choose a different system, then an additional set of storage, memory, and other resources needs to be purchased, and then backup can be restored there. But it will take time to configure the environment. We need to install all the dependent operating system packages for our databases and then restore the backup. Also, the new environment should have access to the backup storage. Here the main concern is the time required for restoring the backup. Restoration will take time based on the volume of the data. But practically any organization cannot afford downtime to restore a backup while it is undergoing a disaster. Consider any international bank—their business cannot afford a single minute

of downtime to customers. So, they need an alternate solution. In addition to downtime, there are many challenges while dealing with backups.

- Suppose your database size is huge. Then the backup size will also be huge. It will take more time to restore that big database.
- If there is a failure with the restoration process, then you need to start it again from the beginning, which will increase the restoration time.
- The restoration process performance should be appropriate.
- There is a chance that a backup piece might have gotten corrupted physically, which means we don't have valid backup.
- When our backup storage is at a remote location and network bandwidth with the remote location is poor, then the restore will take a long time to transfer.
- To access the backup storage, we may need help from storage administrators and system administrators along with database administrators.

In addition, at regular intervals, we should perform restore and recovery testing to ensure the integrity of the backup else we might end up in trouble. What if you don't validate your backup and it doesn't restore properly when we need it. In other words, the backup is not fulfilling its purpose. Another issue with backups is storage media failure. Suppose the site with the backup has met a disaster. Then we don't have any backup for our database, which is an unsafe position.

The next challenging point with backups is determining a retention period. How much backup can we store for your database? Suppose we have the need to restore data up to one week before, as an undesirable change has been made a week ago, and we need to restore the backup to get your database to that point. In that case, we should have backups from a week ago. Remember, the database size will keep increasing every day, so the backup size increases daily. So, we need to do a thorough study about the environment and make a proper judgment about the backup retention period. Having a high retention period will increase the need for more storage.

In addition, we have to choose the right backup media. Storing the backup on disk is expensive. We need to look for low-cost storage like an SBT tape. Currently many customers have started moving backups to cloud object storage.

This discussion shows that we cannot fully trust a backup. We need something beyond the backup to confirm that our environment is safe. Here comes the high availability (HA) and disaster recovery (DR) concepts as our savior. In this method, a similar environment as production/live will be maintained at a remote location, and it will play a role if a disaster strikes. This means we will have a copy of the data at another location, and any changes made in production will get replicated to that secondary site. When there is a disaster situation, we don't need to perform a restore/recovery from backup; instead, we can connect the replicated environment and starting working on that.

By having a replicated high availability environment, we will get the following benefits:

- The secondary environment can be switched to the live environment at any time. The replicate environment is similar to production from a hardware and software configuration perspective. So, it can serve the same workload as production was serving. See Figure 1-2.



**Figure 1-2.** *Data replication*

- During a nondisaster period, the replicate environment can contribute as a reporting database as it has live data within it. This reduces the workload for the production database.