

VersaStack Solution by Cisco and IBM with Oracle RAC, IBM FlashSystem V9000, and IBM Spectrum Protect

Jon Tate

Dharmesh Kamdar

Dong Hai Yu

Randy Watson



Storage



International Technical Support Organization

**VersaStack Solution by Cisco and IBM with Oracle
RAC, IBM FlashSystem V9000, and IBM Spectrum
Protect**

October 2016

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (October 2016)

This edition applies to the software and hardware described in Chapter 3, “Software configuration and revision guidelines” on page 25.

© Copyright International Business Machines Corporation 2016. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
Authors	x
Now you can become a published author, too!	xi
Comments welcome	xi
Stay connected to IBM Redbooks	xi
Chapter 1. Introduction	1
1.1 Easy, efficient, and versatile	2
1.2 Evolving data center requirements	2
1.3 Holistic approach	3
1.4 Hardware options	3
1.5 Related information	4
Chapter 2. Architecture	7
2.1 Architecture features	8
2.2 Environment overview	13
2.2.1 Cisco UCS C3260 Server	13
2.3 IBM Spectrum Protect server deployment on Cisco UCS C3260	14
2.3.1 Pre-installation Steps	15
Chapter 3. Software configuration and revision guidelines	25
3.1 Software revisions	26
3.2 Configuration guidelines	27
Chapter 4. Planning for Oracle Real Application Clusters	33
4.1 Design considerations	34
4.1.1 Database workload	34
4.1.2 Server virtualization	34
4.1.3 Database availability	34
4.1.4 Quality of service and network segregation	34
4.1.5 Network availability	35
4.2 Network and Storage Design	35
4.2.1 Network Design	35
4.2.2 Single Client Access Name (SCAN)	36
4.2.3 Storage Design	37
4.2.4 Oracle Automatic Storage Management (ASM) design	38
Chapter 5. Physical infrastructure	39
5.1 VersaStack cabling	40
5.2 Storage compatibility and interoperability	41
5.3 VersaStack system build process	42
Chapter 6. IBM FlashSystem V9000 storage configuration	43
6.1 Volume layout for four node Oracle-RAC	44
6.2 Volume creation and mapping	44

Chapter 7. Virtual infrastructure configuration	53
7.1 ESXi Configuration	54
7.1.1 Scan Disk Devices	54
7.1.2 Mark FlashSystem V9000 Disks as Flash Disks	56
7.1.3 Modify Path Selection Policy	58
7.1.4 Create datastore	60
7.1.5 Configure NIC Teaming	66
7.1.6 Enable Jumbo Frames	69
7.1.7 Add Host Networking	71
7.2 Create and modify virtual machines	74
7.2.1 Create virtual machines	75
7.2.2 Add network adapter	81
7.2.3 Create shared disks for virtual machines	83
7.2.4 Install VMware Tools	92
7.2.5 Disable VMware Tools Time Synchronization	95
7.2.6 Enable disk UUID	96
7.3 Considerations of installing Red Hat Enterprise Linux	100
7.3.1 Base Environment Selection during installation	100
7.3.2 Swap size	101
Chapter 8. Oracle RAC installation	103
8.1 Oracle Cluster network and storage planning	104
8.1.1 Network planning	104
8.1.2 Storage planning	104
8.2 Oracle RAC installation preparation	105
8.2.1 Check hardware configuration	105
8.2.2 Check Linux kernel and OS version	107
8.2.3 Setup hostname	107
8.2.4 Configure IP address	108
8.2.5 Check swap size	111
8.2.6 Enable jumbo frames support on private Ethernet interface	111
8.2.7 Configure DNS Server	112
8.2.8 Synchronize Time with Time Server	112
8.2.9 Install required RPM packages	114
8.2.10 Create users and groups	115
8.2.11 Configure firewall	116
8.2.12 Configure SELinux	118
8.2.13 Set Kernel parameters	119
8.2.14 Set user limits	120
8.2.15 Disable avahi-daemon service and Zero Configuration Networking	121
8.2.16 Disable RemoveIPC feature	122
8.2.17 Shared memory file system /dev/shm	123
8.2.18 Disable Transparent HugePages	123
8.2.19 Setting disk I/O scheduler on Linux	124
8.2.20 Create disk partitions	125
8.2.21 Install and Configure Oracle ASMLib	127
8.3 Oracle Grid Infrastructure installation	129
8.3.1 Get installation packages	129
8.3.2 Install Oracle Cluster Verification Utility	130
8.3.3 Oracle Grid Infrastructure installation	131
8.4 Oracle RAC Database installation	159
8.5 Create Oracle ASM disk groups	175

8.6 Oracle RAC Database creation	177
8.6.1 Oracle RAC Database creation	178
8.6.2 Update grid and oracle user profile	197
8.6.3 HugePages Configuration	198
8.7 Oracle RAC Database Management	200
8.7.1 Enable archive log mode	200
8.7.2 Add redo log group	203
8.7.3 Multiplex Oracle control files	204
8.7.4 Update processes and sessions parameters	206
8.7.5 Startup and shutdown of Oracle CRS and Database	207
8.7.6 OCR and voting disk management	212
Chapter 9. Using IBM Spectrum Protect	213
9.1 Environment overview	214
9.1.1 IBM Spectrum Protect Backup Server	214
9.1.2 IBM Spectrum Protect Blueprints	214
9.1.3 IBM Spectrum Protect Operations Center	215
9.1.4 IBM Spectrum Protect Data Protection for Databases: Oracle	215
9.2 IBM Spectrum Protect server deployment	215
9.2.1 Code installation: Step 1	216
9.2.2 Server configuration using automated Blueprint script: Step 2	216
9.3 DP for Oracle: Install and configure IBM Spectrum Protect on Oracle server node	219
9.3.1 BAC Installation	219
9.3.2 Install and configure Data Protection (DP) for Databases (TDPO)	221
9.3.3 TDPO.opt	223
9.3.4 TDPO dsm.opt	223
9.3.5 BAC API dsm.sys	223
9.3.6 Use RMAN to back up the Oracle Database using TDPO	225
9.3.7 IBM Spectrum Protect activity log for Oracle database backup	227
9.3.8 View of backup data from IBM Spectrum Protect Operations Center	227
9.3.9 View of Oracle server node from IBM Spectrum Protect Operations Center	228
9.3.10 TSM CLI query occ command after backups	228
Related publications	233
IBM Redbooks	233
Other resources	233
Online resources	234
Help from IBM	235

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM SmartCloud®	Redbooks®
DB2®	IBM Spectrum™	Redbooks (logo)  ®
Easy Tier®	IBM Spectrum Control™	Storwize®
FlashCopy®	IBM Spectrum Protect™	System Storage®
IBM®	MicroLatency®	Tivoli®
IBM FlashCore™	Passport Advantage®	
IBM FlashSystem®	Real-time Compression™	

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Dynamic organizations want to accelerate growth while reducing costs. To do so, they must speed the deployment of business applications and adapt quickly to any changes in priorities. Organizations today require an IT infrastructure that is easy, efficient, and versatile.

The VersaStack solution by Cisco and IBM® can help you accelerate the deployment of your data centers. It reduces costs by more efficiently managing information and resources while maintaining your ability to adapt to business change.

The VersaStack solution combines the innovation of Cisco UCS Integrated Infrastructure with the efficiency of the IBM Storwize® storage system. The Cisco UCS Integrated Infrastructure includes the Cisco Unified Computing System (Cisco UCS), Cisco Nexus and Cisco MDS switches, and Cisco UCS Director. The IBM FlashSystem® V9000 enhances virtual environments with its Data Virtualization, IBM Real-time Compression™, and IBM Easy Tier® features. These features deliver extraordinary levels of performance and efficiency.

The VersaStack solution is Cisco Application Centric Infrastructure (ACI) ready. Your IT team can build, deploy, secure, and maintain applications through a more agile framework. Cisco Intercloud Fabric capabilities help enable the creation of open and highly secure solutions for the hybrid cloud. These solutions accelerate your IT transformation while delivering dramatic improvements in operational efficiency and simplicity.

Cisco and IBM are global leaders in the IT industry. The VersaStack solution gives you the opportunity to take advantage of integrated infrastructure solutions that are targeted at enterprise applications, analytics, and cloud solutions.

The VersaStack solution is backed by Cisco Validated Designs (CVD) to provide faster delivery of applications, greater IT efficiency, and less risk.

This IBM Redbooks® publication is aimed at experienced storage administrators who are tasked with deploying a VersaStack solution with Oracle Real Application Clusters (RAC) and IBM Spectrum™ Protect.

Authors


This book was produced by a team of specialists from around the world working at the Cisco campus in Raleigh, North Carolina.



Jon Tate is a Project Manager for IBM Storage at the International Technical Support Organization (ITSO), San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2/3 support for IBM storage products. Jon has 30 years of experience in storage software and management, services, and support, and is an IBM Certified IT Specialist, an IBM SAN Certified Specialist, and a Project Management Professional (PMP). He also serves as the UK Chairman of the Storage Networking Industry Association.



Dharmesh Kamdar has been working in IBM Systems group for over 14 years as a Senior Software Engineer. He is working in the Open Systems Lab (OSL), where he focuses on interoperability testing of a range of IBM storage products with various vendor products, including operating systems, clustering solutions, virtualization platforms, volume managers, and file systems.



Dong Hai Yu is a consulting IT specialist at IBM Greater China Group (GCG) Systems Lab Services storage team. His current role in Lab Services team is to deliver IBM storage-related implementation services and consulting services for customers in GCG. Dong Hai has been with IBM since 2006. He has 13 years of experience on storage products, and his areas of expertise include storage infrastructure optimization, IBM storage disaster recovery and resiliency solution implementation, and IBM virtualization storage support in heterogeneous environments. Dong Hai holds a Bachelor degree in microelectronics from Sichuan University in China.



Randy Watson is a member of the IBM Systems Storage SWAT team that focuses on assisting customers to achieve their goals with IBM storage software. He has over a decade of experience with implementation and performance consulting with IBM Spectrum Protect backup and restore. Randy focuses on IBM Spectrum Protect Snapshot and database protection.

Thanks to the following people for their contributions to this project:

Sreeni Edula
Chris O'Brien
Caela Dehaven
Cisco

Dave Gimpl
Troy Lee
Ravisankar Shanmugam

Rawley Burbidge
Doug Bloom
Mike Ault
Roberta Lockaby
Sally Neate
Kristy Ortega
Anne Fitzpatrick
Jason F Basler
Daniel Thompson
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Introduction

IBM and Cisco have a long history of working together to deliver technology that benefits our mutual clients. Expanding on this success, IBM and Cisco have jointly developed VersaStack, an innovative validated design that brings together IBM Storwize, IBM FlashSystem, and Cisco Unified Computing System (UCS) integrated infrastructure, allowing business partners and customers to create solutions that transform their business and reduce risk.

This collaboration incorporates IBM Storwize and IBM FlashSystem storage into the single pane of glass management environment that is provided by Cisco UCS Director. Future capabilities will deliver Application Centric Infrastructure (ACI) and Intercloud Fabric from Cisco and use within the IBM Cloud Builder Professional Services offering from IBM Global Services.

VersaStack will be backed by a series of Cisco Validated Designs (CVD) and IBM Redbooks developed in conjunction by Cisco and IBM. These materials will provide faster delivery of applications, and greater reliability and confidence for customers and business partners.

This chapter includes the following sections:

- ▶ Easy, efficient, and versatile
- ▶ Evolving data center requirements
- ▶ Holistic approach
- ▶ Hardware options
- ▶ Related information

1.1 Easy, efficient, and versatile

In today's environment, quick deployment and execution of business applications plus the versatility to adapt as business priorities change are essential for dynamic organizations that want to accelerate business growth while reducing costs. Organizations today require an IT infrastructure that is easy, efficient, and versatile. The VersaStack solution by Cisco and IBM helps accelerate data center infrastructure deployment, efficiently manage information and resources, and adapt to business change.

VersaStack includes technologies that complement and enhance virtual environments with built-in functions such as IBM Data Virtualization, Real-time Compression, and Easy Tier that deliver extraordinary levels of performance and efficiency.

Alternatively (and outside the intended scope of this book) for clients who require the combined capabilities to support block and file data, the IBM Storwize V7000 Unified storage product is also offered in VersaStack. This validated design also includes the performance and innovation of Cisco UCS integrated infrastructure, which includes the Cisco UCS, Cisco Nexus and Cisco MDS 9000 Family switches, and Cisco UCS Director, with the performance and efficiency of the IBM FlashSystem storage system.

VersaStack is backed by Cisco Validated Designs for faster delivery of applications and increased IT efficiency with less risk. VersaStack is supported by a broad range of services from IBM Business Partners and IBM Global Services.

VersaStack is also ready for Cisco ACI, enhancing business agility by allowing IT to build, deploy, secure, and maintain applications through a more agile framework. This capability, combined with Cisco Intercloud Fabric, can enable the creation of open and secure hybrid cloud-ready solutions that accelerate IT agility while delivering dramatic improvements in deployment, operation efficiency, and simplicity.

1.2 Evolving data center requirements

As anyone that has been around this industry knows, the data center industry is always evolving, and current trends make evolution more critical than ever. The data center has moved far beyond a simple repository for digital records, and way beyond just a vehicle for backup and restore.

Increasingly, its compute, storage, and networking facilities are being used to power complex analytical operations that are becoming essential for competitive advantage and business agility.

This trend is exemplified by the growth in demand for big data applications, and the Internet of Things. These applications involve data sets so large and complex that they cannot easily be processed using traditional computing tools.

Two other trends are making it easier to provision data center resources:

- ▶ Cloud computing, in which computing and storage assets are managed and allocated from a shared pool rather than from application-based silos, is rapidly becoming the standard for data center resource deployment.
- ▶ The advent of virtualization and software-defined networking (SDN), in which management is abstracted from lower-level functions, promises to make it easier than ever to allocate resources.

These trends are related because the scalability of big data and the simplicity implied by SDN help organizations manage the increased compute requirements of big data. And underpinning these trends are changes in hardware. Vendors are adapting specific data center components to address cloud, SDN, and big data requirements. IBM, for example, has changed its Storwize family of virtualized storage technologies specifically for software-defined environments.

Cisco, meanwhile, developed Cisco ACI to accelerate the configuration of infrastructure to match the needs of applications, and Cisco Intercloud Fabric technology to make it easier to move workloads between different cloud models.

Another significant development is the emergence of integrated infrastructure solutions for the data center. Previously, data center teams purchased computing, storage, and network building blocks separately and assembled, configured, and tested the various technologies with the hope that everything would work together. With integrated infrastructure, servers, networking resources, storage systems, and management systems are combined into a predesigned, tested, and supported solution. This approach massively simplifies asset purchasing, deployment, and management.

1.3 Holistic approach

So, to be absolutely clear on one thing, this is not about just bolting hardware and software together. Both IBM and Cisco are fully aware of the requirements of the enterprise today. With that in mind, it makes perfect sense to streamline and consolidate the traditional infrastructure into a full stack solution that is a new way to management efficiency and enhanced productivity. IT professionals the world over trust IBM and Cisco products, and this partnership takes this trust to a new level.

The VersaStack solution by Cisco and IBM is optimized for those IT professionals.

1.4 Hardware options

IBM FlashSystem V9000 offers full integration and is a comprehensive all-flash enterprise storage solution. IBM FlashSystem V9000 delivers the full capabilities of IBM FlashCore™ technology plus a rich set of storage virtualization features. It is optimized for flash storage with an upcoming release that will support a simple two-tier easy tier solution. IBM FlashSystem V9000 is ideal for migrating external storage into the new configuration and providing future flexibility.

FlashSystem V9000 uses a fully featured and scalable all-flash architecture that performs at up to 2.5 M IOPS with IBM MicroLatency®, is scalable to 19.2 GBps, and delivers up to 2.28 PB effective capacity. Leveraging its Flash-optimized design, IBM FlashSystem V9000 can provide response times of 200 microseconds. It delivers better acquisition costs than high-performance spinning disk for the same effective capacity while achieving five times the performance, making it ideal for environments that require extreme performance. For more information, see these websites:

<http://www.ibm.com/systems/uk/storage/flash/v9000/>

<http://www.redbooks.ibm.com/abstracts/tips1281.html?Open>

For customers who want to go outside the IBM FlashSystem V9000 solution, the FlashSystems 900 can go behind stand-alone SVC 2145-DH8 nodes, offering greater flexibility.

The IBM FlashSystem 900 can be added to a storage array and provide high performance and low latency to connected hosts, while taking advantage of the IBM storage management services. By using Spectrum Control, you can use advanced analytics to automatically tier I/O-intensive payloads to the FlashSystem.

The IBM FlashCore technology used in IBM FlashSystem 900 employs several new and patented mechanisms to achieve greater capacity and throughput. These mechanisms enable you to accelerate your mid-range storage solution by taking advantage of the extreme performance and low latency of the FlashSystem.

This option is also available with the existing Storwize V7000, and is as simple as adding the FlashSystem 900 to an existing pool.

For more information, see the following websites:

<http://www.ibm.com/systems/storage/flash/>

<http://www.redbooks.ibm.com/abstracts/tips1261.html?Open>

<http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg248271.html?Open>

1.5 Related information

This section provides links to other material related to VersaStack that might be of interest:

- ▶ VersaStack Solution - Cisco
<http://www.cisco.com/c/en/us/solutions/data-center-virtualization/versastack-solution-cisco-ibm/index.html>
- ▶ VersaStack Solution by Cisco and IBM
http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=PM&subtype=SP&htmlfid=TS03159USEN&appname=TAB_2_1_Appname
- ▶ VersaStack Designs (links to PDF download page)
<http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/versastack-designs.html>
- ▶ Video: New VersaStack Solution by Cisco and IBM
<https://www.youtube.com/watch?v=HHTgEABDYts>
- ▶ Video: High Level Business Value of VersaStack from IBM & CISCO
<https://www.youtube.com/watch?v=E0W4ggyN99o>
- ▶ Video: IBM and Cisco VersaStack - Introduction
<https://www.youtube.com/watch?v=mkg1fkpAKII>
- ▶ Video: IBM and Cisco VersaStack - Turbo Compression
https://www.youtube.com/watch?v=PR_Uir1mxXE
- ▶ Video: IBM and Cisco VersaStack - Data Virtualization
<https://www.youtube.com/watch?v=N-rNcokXzf0>
- ▶ Video: IBM and Cisco VersaStack - Flash Optimization and IBM EasyTIER
<https://www.youtube.com/watch?v=J7Rr13fEv0U>
- ▶ Video: IBM and Cisco VersaStack - Compression
<https://www.youtube.com/watch?v=xDbk4ddXzL0>

- ▶ Video: Talking VersaStack with Your Customers
<https://www.youtube.com/watch?v=UHANwo51ie0>
- ▶ Video: Client value of VersaStack
<https://www.youtube.com/watch?v=dvDG6UHMEuQ>
- ▶ Video: Growth Opportunities with VersaStack Solution
<https://www.youtube.com/watch?v=h32TsA2smLk>
- ▶ Video: Take 5 - VersaStack by Cisco & IBM
<https://www.youtube.com/watch?v=18mKR0sKQ3o>



Architecture

This chapter describes the features of the architecture that is implemented later in this book.

This chapter includes the following sections:

- ▶ Architecture features
- ▶ Environment overview
- ▶ IBM Spectrum Protect server deployment on Cisco UCS C3260

2.1 Architecture features

The Oracle database on VersaStack design combines a high availability Oracle-RAC cluster configuration running on VersaStack with IBM FlashSystem V9000, as shown in Figure 2-1.




Figure 2-1 Architecture diagram

Figure 2-1, illustrates a four node Oracle-RAC built on VersaStack components and the network connections for a configuration with an IBM FlashSystem V9000 storage system.

The VersaStack architecture is highly modular, and has sufficient architectural flexibility and design options to scale as required, providing investment protection. The platform can be scaled up (adding resources to existing VersaStack units) or out (adding more VersaStack units).

Specifically, this VersaStack is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VersaStack All-Flash includes IBM FlashSystem V9000, Cisco networking, Cisco UCS, Cisco MDS Fibre Channel switches, and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations.

One benefit of the VersaStack architecture is the ability to meet any client’s capacity or performance needs in a cost effective manner. The Converged Infrastructure system is capable of serving multiple protocols across a single interface, which allows for customer choice and investment protection because it is wire-once architecture.

This architecture references relevant criteria that pertain to resiliency, cost benefit, and ease of deployment of all components including IBM FlashSystem V9000 storage.

The architecture for this solution, which is shown in Figure 2-2, uses two sets of hardware resources:

- ▶ Common Infrastructure services on redundant and self-contained hardware
- ▶ VersaStack PoD with Oracle-RAC




Figure 2-2 Oracle-RAC built on VersaStack

Common infrastructure services such as Active Directory, DNS, DHCP, vCenter, Nexus 1000v virtual supervisor module (VSM), and UCS Performance Manager can be deployed on redundant and self-contained hardware in a Common Infrastructure Pod along with the VersaStack Pod. Because these services are integral to the deployment and operation of the platform, you need to adhere to preferred practices in their design and implementation. These practices include such features as high-availability, choosing the appropriate RAID setup, and considering performance and scalability because these services might need to be extended to multiple PoDs. At a client’s site, depending on whether this is a new data center, there might not be a need to build this infrastructure from scratch.

Figure 2-2 illustrates Oracle-RAC built on VersaStack components and the network connections for a configuration with a IBM FlashSystem V9000 storage system.

The reference hardware configuration includes the following items:

- ▶ Two Cisco Nexus 9396 or 9372 switches.
- ▶ Two Cisco UCS 6248UP Fabric Interconnects.

- ▶ Two Cisco MDS 9148S Fibre Channel switches.
- ▶ Cisco UCS C3260 Server.
- ▶ Support for 32 Cisco UCS C-Series servers without any additional networking components.
- ▶ Support for eight Cisco UCS B-Series servers without any additional blade server chassis.
- ▶ Support for 160 Cisco UCS C-Series and B-Series servers by way of extra fabric extenders and blade server chassis.
- ▶ Two IBM FlashSystem V9000 control enclosures, and one FlashSystem V9000 Storage enclosure. Support for up to 12 flash modules of the same capacity within storage enclosures.

For server virtualization, the deployment includes VMware vSphere. Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, extra disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features.

This book guides you through the low-level steps for deploying the base architecture. These procedures cover everything from physical cabling to network, compute, and storage device configurations, including Oracle-RAC deployment.

For more information about the design of VersaStack, see the Design guide at:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/Versastack_vmw6_flash_design.html

IBM Spectrum Protect™ has been deployed on a Cisco UCS C3260 server, which has support for high capacity onboard storage that is used as backup repository in the solution. The Cisco UCS C3260 addresses the need for highly scalable computing with high-capacity local storage and is ideal for small to large-scale IT environments.

The IBM Spectrum Protect software embedded in Cisco UCS C3260 is simple to deploy and manage. Together, IBM and Cisco UCS C3260 create the perfect staging area for backups by reducing backup ingest bottlenecks and providing faster backups by using parallel processing.

The Cisco UCS C3260 offers these main benefits:

- ▶ Fully modular chassis
- ▶ Innovative airflow design for a compact 4RU, 31.8-inch-deep, industry-standard rack server
- ▶ Ease of upgradeability

Figure 2-3 shows the architecture.




Figure 2-3 Cisco UCS C3260 Architecture

In addition, IBM Spectrum Protect Backup and Replication provides these advantages:

- ▶ Advanced data protection and flexible recovery options for virtual machines, file servers, email, databases, enterprise resource planning (ERP) systems, mainframes, and desktops
- ▶ Application-aware and virtual machine optimized options enable customized data protection for critical workloads
- ▶ A single platform for virtual, physical, software- defined, and cloud backups
- ▶ Auto- discovery of new virtual backups helps ensure that all data in the virtualized environment is protected
- ▶ Application- consistent backup helps ensure complete application backups and simplifies restore processes
- ▶ Instant access/ instant restore reduces the user impact to nearly zero during virtual machine restores

IBM and Cisco offer the correct solution for performance, flexibility, and reliability, providing a highly efficient and cost effective data protection solution for your virtual and application environments on VersaStack.

Figure 2-4 illustrates a high-level overview of VersaStack All-Flash environment with integrated backup infrastructure.




Figure 2-4 VersaStack data protection with IBM Spectrum Protect on Cisco UCS C3260

The VM and application data from the VersaStack gets backed up to the IBM Spectrum Protect appliance that is deployed on Cisco UCS C3260 server. The Cisco UCS C-Series servers can be seamlessly integrated into the VersaStack environment by being connected to the Cisco 9k Nexus Unified Fabric, providing organizations with a built-in path to unified computing. The main benefits of the VersaStack integrated data protection solution are reduced TCO, improved IT agility and flexibility, investment protection, and future readiness.

IBM Spectrum Protect Snapshot, formerly IBM Tivoli® Storage FlashCopy® Manager, delivers high levels of protection for key applications and databases by using advanced integrated application snapshot backup and restore capabilities.

VersaStack data protection environment can be further optimized with IBM Spectrum Protect snapshot integration. It lets you perform and manage frequent, near-instant, nondisruptive, application-aware backups and restores using integrated application and VM-aware snapshot technologies in both IBM and non-IBM storage systems.

However, this document does not cover the implementation details of IBM Spectrum Protect Snapshot, but the customer's can chose this optional deployment for their data protection needs.

IBM Spectrum Protect Snapshot has these characteristics:

- ▶ Easy to install, configure, and deploy while supporting a wide range of applications, operating systems, and storage.
- ▶ Improves application availability and service levels through high-performance, near-instant restore capabilities that reduce downtime.
- ▶ Manages application-aware snapshots on VersaStack storage.
- ▶ Enables fast, simple recovery of individual files, volumes, and mailboxes.
- ▶ Allows “instant” restore for VMware datastores and simplifies database cloning.
- ▶ Enables off-site copying for enhanced data protection using either IBM Spectrum Protect (Tivoli Storage Manager) or IBM Metro Global Mirror.

Figure 2-5 illustrates IBM Spectrum Protect Snapshot in a VersaStack Environment.




Figure 2-5 VersaStack data protection with IBM Spectrum Protect Snapshot

2.2 Environment overview

This section details the IBM Spectrum Protect server components that are deployed in the example Oracle DB server on VersaStack environment.

IBM Spectrum Protect has these core components:

- ▶ Hardware:
 - Cisco UCS C3260 Server
- ▶ Software:
 - IBM Spectrum Protect backup server
 - Spectrum Operations Center
 - IBM Spectrum Protect Data Protection for Databases: Oracle

2.2.1 Cisco UCS C3260 Server

As described earlier, the C3260 is a modular, high-storage-density rack server for storage-driven and high availability use cases. The rack server offers the highest levels of drive density and availability.

The C3260 is a dual-node, high-density, bare-metal, x86-based enterprise storage server. Its features include the following:

- ▶ Enterprise-class redundancy with full featured Redundant Array of Independent Disks (RAID) plus just a bunch of disks (JBOD)
- ▶ Stand-alone management interface (Cisco Integrated Management Controller)
- ▶ No data migration required when replacing or upgrading server nodes
- ▶ No need for extended depth racks

The C3260 supports:

- ▶ Network File System (NFS)
- ▶ Internet Small Computer System Interface (iSCSI)
- ▶ Fibre Channel over Ethernet (FCoE)
- ▶ SMB
- ▶ SMB Direct

The following are the specifications at a glance:

- ▶ Supports up to 360 TB of modular storage capacity
- ▶ Optimized for high throughput performance, high capacity, and a small footprint
- ▶ Enterprise-class redundancy with full featured RAID plus JBOD
- ▶ Stand-alone management interface (Cisco Integrated Management Controller)
- ▶ Up to 512 GB of memory per server node
- ▶ Up to 62 drive bays
- ▶ Up to 4 GB of RAID cache

The UCS C3260 server can be customized based on the initial and predicted growth of the protected data on a VersaStack. Also, the performance and capacity needs of IBM Spectrum Protect can be met with the varying drive options supported on the C3260.

For more information about IBM Spectrum Protect scale points and sizing guidelines, see the blueprints at the following URL:

<https://ibm.biz/BdHc6b>

2.3 IBM Spectrum Protect server deployment on Cisco UCS C3260

This section describes the Cisco UCS C3260 server deployment. Note that this does not cover specific customer needs and is a generic deployment with suggestions. See the UCS C-Series and IBM Spectrum Protect deployment documents for more information and best practices:

<http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c3260-rack-server/model.html>

2.3.1 Pre-installation Steps

Perform the following steps before starting the installation.

Connect Server to the Switches

The UCS C3260 Server supports two VIC cards and can be ideally connected to the Cisco Nexus Switches in the VersaStack as shown in Figure 2-6. Connect each port on the adapter to each of the Nexus 9k switches. This setup provides high availability during a switch failure. During this validation, we used one VIC card only within the server that was connected to both the Nexus Switches for high availability.

Figure 2-6 shows the Nexus 9k switches to Cisco C3260 connectivity.




Figure 2-6 Nexus 9k switches to Cisco C3260 connectivity

Connect the Management IP Interface

Connect the keyboard, video, and mouse (KVM) dongle. Connect the CIMC management port to the Cisco Nexus 9000 Series Switch to access through InBand management network or to the dedicated management switch according to your specific VersaStack deployment. Set the management interface.

Upgrade to the Latest Firmware and Patches

Log in to <http://software.cisco.com> and navigate to **Downloads Home** → **Products** → **Servers—Unified Computing** → **UCS C-Series Rack-Mount Standalone Server Software** → **UCS C3260 Rack Server Software**. Upgrade to the latest and the suggested firmware.

Create Additional vNIC Interfaces

Log in to the Cisco Integrated Management Controller (IMC) user interface and open the Inventory tab under Server and Cisco VIC adapters. Add interfaces to adapter. The validation environment consisted of only one VLAN (11) that was used for management and backup connectivity. Two VLANs are suggested: One for the public network (VLAN 11), and one for the backup network (VLAN 3176).

Figure 2-7 shows adding vNICs on Adapter 1.




Figure 2-7 Adding vNIC's on Adapter 1

The Cisco UCS C3260 IBM Spectrum Protect appliance is connected to the VersaStack environment by using the Cisco Virtual Interface Card (VIC) adapters on the server.

The server can connect at 40 Gbps to the switches if needed. In the lab environment, we connected the server at 10 Gbps speeds to the two upstream Nexus 9k switches that were part of VersaStack.

Configuring Interfaces and creating Virtual Port channel on the Nexus 9k switches

See the VersaStack V9000 CVD for detailed Nexus Configuration information and for the variables that need to be defined:

VLAN ID 3176 can be additionally defined for a backup network if a dedicated network is of interest to you.

Cisco Nexus 9000 A

Complete these steps to configure the Cisco Nexus 9000 A:

1. Define a description for the port-channel connecting to UCS C3260:

```
interface Po15
description Spectrum-Protect-PC
```
2. Make the port-channel a switchport, and configure a trunk to allow in-band management, NFS, VM traffic, and the native VLANs:

```
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_backup_vlan_id>>
```
3. Make the port channel and associated interfaces spanning tree edge ports:

```
spanning-tree port type edge trunk
```

4. Set the MTU to 9216 to support jumbo frames:


```
mtu 9216
```
5. Make this a VPC port-channel and bring it online:


```
vpc 15
no shutdown
```
6. Define a port description for the interface that connects to UCS C3260, NIC1:


```
interface Eth1/47
description Spectrum-Protect-NIC1
```
7. Apply it to a port channel and open the interface:


```
channel-group 15 force mode active
no shutdown
copy run start
```

Cisco Nexus 9000 B

Complete these steps to configure the Cisco Nexus 9000 B:

1. Define a description for the port-channel connecting to UCS C3260:


```
interface Po15
description Spectrum-Protect-PC
```
2. Make the port-channel a switchport, and configure a trunk to allow in-band management, NFS, VM traffic, and the native VLANs:


```
switchport
switchport mode trunk
switchport trunk native vlan <<var_native_vlan_id>>
switchport trunk allowed vlan <<var_ib-mgmt_vlan_id>>, <<var_backup_vlan_id >>
```
3. Make the port channel and associated interfaces spanning tree edge ports:


```
spanning-tree port type edge trunk
```
4. Set the MTU to be 9216 to support jumbo frames:


```
mtu 9216
```
5. Make this a VPC port-channel and bring it up:


```
vpc 15
no shutdown
```
6. Define a port description for the interface connecting to UCS C3260 NIC2:


```
interface Eth1/47
description Spectrum-Protect-NIC2
```
7. Apply it to a port channel and open the interface:


```
channel-group 15 force mode active
no shutdown
copy run start
```

UCS C3260 Storage configuration using Cisco IMC

The Cisco UCS C3260 RAID controller supports various RAID configurations. The disk drives can be allocated to the operating system by configuring them using hardware RAID. The drives can also be presented in pass-through mode to the operating system optionally. This solution validation presents the drives to Red Hat OS in pass-through mode and did not create RAID groups.

In client production environments, create RAID groups for high availability and better performance. The RAID type and the drive capacity can be determined based on the IBM suggestion for IBM Spectrum Protect installation, and that in turn depends on protected environment size and specific customer requirements.

The following procedure details the configuration of RAID drives:

1. Open up a web browser and enter the IP for the Cisco IMC. You are directed to the login window where you can enter your credentials for access as shown in Figure 2-8.



Figure 2-8 Cisco Integrated Management Controller

- After you are logged in, click the menu options in the upper left corner in Figure 2-9. Expand the **Storage** menu from the left side pane and select the controller.




Figure 2-9 Menu options

- Under the Physical Drive Info tab, select each drive individually and click **Set State as Unconfigured Good**. Repeat this for all physical drives available, then click the Controller Info tab.
- Click **Create Virtual Drive from Unused Physical Drives** and a configuration box opens (Figure 2-10). The selections made here result in a new virtual drive that is used as a IBM Spectrum Protect backup repository.




Figure 2-10 Create virtual drive

For more information about IBM Spectrum Protect scale points and sizing guidelines, see to the blueprints at the following URL:

<https://ibm.biz/BdHc6b>

Note: IBM DB2® as part of IBM Spectrum Protect installation has a limitation on disk sector sizes. DB2 only supports storage devices with a sector size of 512 bytes. The 6 TB drives supported in UCS C3260 have a sector size of 4096 bytes, so these drives are not supported for DB2 database in an IBM Spectrum Protect environment. Lower capacity drives need to be considered for DB2. SSDs provide the best performance and therefore accelerate backup and restore operations.

Install the operating system

Complete these steps to install the OS:

1. Load the OS installation disk into your CD/DVD drive, or copy the disk image files to your computer.
2. If Cisco IMC is not open, log in.
3. In the Navigation pane, click the **Compute** menu.
4. In the **Compute** menu, click **Server 1** or **Server 2**.
5. In the Server pane, click the Remote Management tab.
6. In the Remote Management pane, click the Virtual KVM tab.
7. In the Actions area, click **Launch KVM Console**. The KVM Console opens in a separate window.
8. From the KVM console, click the VM tab.
9. In the VM tab, map the virtual media by using either of the following methods:
 - Select the **Mapped** check box for the CD/DVD drive containing the OS installation disk.
 - Click **Add Image**, navigate to and select the OS installation disk image, click **Open** to mount the disk image, and then select the **Mapped** check box for the mounted disk image.

You must keep the VM tab open during the OS installation process.

10. Reboot the server and select the virtual CD/DVD drive as the boot device.

When the server reboots, it begins the installation process from the virtual CD/DVD drive. See the installation guide for the Red Hat OS to guide you through the rest of the installation process at www.redhat.com.

Perform NIC Bonding

Log in to the server after Red Hat OS installation and create bond interfaces: bond0 for public and client traffic, and bond1 for the backup network. Separate backup network was not created for validation of the solution. However, generally have a separate backup network to segregate the production, application, and management traffic from the backup traffic.

Complete these steps to create NIC bonding by using the GUI:

1. Click **Application** and select the **System Tools and Settings**.
2. Double-click **Network** under the Hardware section.
3. Click the **plus** symbol to open the selection list. Select **Bond**. The Editing Bond connection 1 window opens.

4. On the Bond tab, click **Add** and select the type of interface you want to use with the bond connection. Click the **Create** button. Note that the dialog to select the slave type only comes up when you create the first slave. After that, it will automatically use that same type for all further slaves.
5. The Editing bond0 slave 1 window appears. Use the **Device MAC address** menu to select the MAC address of the interface to be bonded. The first slave's MAC address is used as the MAC address for the bond interface.
6. Select **802.3ad** as the failover mode and leave all the other entries to default.
7. Click the **Save** button to add further slaves as shown in Figure 2-11.

Editing Bond connection 1

Connection name:

General **Bond** IPv4 Settings IPv6 Settings

Interface name:

Bonded connections:

bond0 slave 1	Add
bond0 slave 2	

Edit

Delete

Mode:

Link Monitoring:

Monitoring frequency: ms

Link up delay: ms

Link down delay: ms

MTU: bytes

Cancel Save

Figure 2-11 Editing Bond connection 1

8. Click the IPv4 Settings tab, select **Manual Method**, and assign the management IP address details of the server as shown in Figure 2-12.

Editing Bond connection 1

Connection name:

General Bond **IPv4 Settings** IPv6 Settings

Method:

Addresses

Address	Netmask	Gateway
192.168.161.45	255.255.252.0	192.168.160.1

DNS servers:

Search domains:

DHCP client ID:

Require IPv4 addressing for this connection to complete

Figure 2-12 Assigning the management IP address details

9. Repeat the above steps to create another bond interface if dedicated backup network is deployed as shown in Figure 2-13.

Editing Bond connection 2

Connection name: Bond connection 2

General **Bond** IPv4 Settings IPv6 Settings

Interface name: bond1

Bonded connections:

- bond1 slave 1
- bond1 slave 2

Buttons: Add, Edit, Delete

Mode: 802.3ad

Link Monitoring: MII (recommended)

Monitoring frequency: 1 ms

Link up delay: 0 ms

Link down delay: 0 ms

MTU: automatic bytes

Buttons: Cancel, Save

Figure 2-13 Edit Bond Connection 2



Software configuration and revision guidelines

This chapter describes the software revisions and versions that are used, and the configuration that is deployed.

This chapter includes the following sections:

- ▶ Software revisions
- ▶ Configuration guidelines

3.1 Software revisions

Table 3-1 describes the software revisions that are used for validating various components of the Cisco Nexus 9000 based VersaStack architecture at the time of writing.

For the latest supported versions, see the following IBM and Cisco support matrix links:

- ▶ IBM System Storage® Interoperability Center:
<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>
- ▶ Spectrum Control Interoperability Matrix:
<http://www.ibm.com/support/docview.wss?uid=swg21386446>
- ▶ IBM Spectrum Protect Interoperability Matrix:
<http://www.ibm.com/support/docview.wss?uid=swg21243309>
- ▶ FlashCopy Manager Interoperability Matrix:
<http://www.ibm.com/support/docview.wss?uid=swg21829854>
- ▶ Cisco UCS Interoperability Matrix:
<http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>

After the software versions are validated, it is necessary to validate the Cisco Drivers:

- ▶ To validate your ENIC version, run `ethtool -i vmnic0` by using the command-line interface (CLI) of the ESX host.
- ▶ To validate your FNIC version, run `vmkload_mod -s fnic` by using the CLI of the ESX host.

Table 3-1 Software revisions

Layer	Device	Version/Release	Details
Compute	Cisco UCS fabric interconnect 6248	3.1(1e))	Embedded management
	Cisco UCS 5108 Blade Server Chassis	N/A	Software runs on FI
	Cisco UCS B 200 M4	3.1(e)	Software bundle release
	Cisco ENIC	2.3.0.6	Ethernet driver for Cisco VIC
	Cisco FNIC	1.6.0.24	FCoE driver for Cisco VIC
Network	Cisco Nexus 9000 c9372PX	6.1(2)I3(5)	Operating system version
	Cisco MDS 9148S	6.2(13b)	FC switch firmware version
Storage	IBM FlashSystem V9000	7.6.1.4	Software version

Layer	Device	Version/Release	Details
Software	Cisco UCS hosts	VMware vSphere ESXi 6.0u1	Operating system version
	Oracle	12.1.0.2	Built-in server for vCenter
	VMware vCenter	6.0 u1	Software version
	Linux Server	Red Hat Enterprise Linux Server v7.2	Operating system version
	IBM Spectrum Control™ (IBM SmartCloud® Virtual Storage Center)	5.2.6	Software version
	IBM Spectrum Protect	7.1.6	Software version

3.2 Configuration guidelines

This document provides details about configuring a fully redundant, highly available VersaStack unit with IBM FlashSystem V9000 storage. This process is covered by *VersaStack for Data Center with All-Flash Storage and VMware vSphere 6.0 Deployment Guide*, which is available here:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/Versastack_vmw_6_flash.html

The Cisco UCS fabric Interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts. These hosts are identified sequentially (vm-host-Oracle-01, vm-host-Oracle-02, and so on).

Finally, to indicate that you should include information pertinent to your environment in a step, <text> appears as part of the command structure.

The following example shows the **network port vlan create** command parameters:

```
network port vlan create ?
```

Where:

[-node] <nodename>	Node
{ [-vlan-name] {<netport> <ifgrp>}	VLAN name
-port {<netport> <ifgrp>}	Associated network port
[-vlan-id] <integer> }	Network switch VLAN identifier

Example 3-1 shows an example of the command.

Example 3-1 Example of the network port command

```
network port vlan -node <node01> -vlan-name i0a-<vlan id>
```

You can use this book to configure the VersaStack PoD in the environment. Various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, and to record the appropriate MAC addresses.

Table 3-2 describes the VLANs that are necessary for deployment, as outlined in this book.

Table 3-2 Necessary VLANs

VLAN name	VLAN ID used in validating this document	Purpose
DevMgmt	11	All structure management in this VLAN
vMotion	3173	VMware vMotion traffic
InterConnect	3175	Oracle RAC heartbeat traffic

Table 3-3 lists the virtual machines (VMs) necessary for deployment, as outlined in this book.

Table 3-3 VMware virtual machine created

Virtual machine description	Quantity
Red Hat Enterprise Linux v7.2 x86-64 VM	4
vCenter Server	1

Table 3-4 lists the configuration variables that are used throughout this book. This table can be completed based on the specific site variables used in implementing the document configuration steps. These variables are referenced at various places within this book.

Table 3-4 Configuration variables

Variable	Description	Customer value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01	
<<var_node01_mgmt_mask>>	Out-of-band management network netmask	
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02	
<<var_node02_mgmt_mask>>	Out-of-band management network netmask	
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_cluster_mgmt_ip>>	Out-of-band management IP for cluster	
<<var_cluster_mgmt_mask>>	Out-of-band management network netmask	
<<var_cluster_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_password>>	Global default administrative password	
<<var_dns_domain_name>>	DNS domain name	

Variable	Description	Customer value
<<var_nameserver_ip>>	DNS server IPs	
<<var_timezone>>	VersaStack time zone (for example, America/New_York)	
<<var_global_ntp_server_ip>>	NTP server IP address	
<<var_email_contact>>	Administrator email address	
<<var_admin_phone>>	Local contact number for support	
<<var_mailhost_ip>>	Mail server host IP	
<<var_country_code>>	Two-letter country code	
<<var_state>>	State or province name	
<<var_city>>	City name	
<<var_org>>	Organization or company name	
<<var_unit>>	Organizational unit name	
<<var_street_address>>	Street address for support information	
<<var_contact_name>>	Name of contact for support	
<<var_admin>>	Secondary admin account for storage login	
<<var_nexus_A_hostname>>	Cisco Nexus A host name	
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_B_hostname>>	Cisco Nexus B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_devmgmt_vlan_id>>	In-band management network VLAN ID	
<<var_vmotion_vlan_id>>	VMware vMotionVLAN ID	
<<var_orarac_vlan_id>>	Oracle RAC heartbeat traffic	
<<var_backup_vlan_id>>	Backup traffic for storage	

Variable	Description	Customer value
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name	
<<var_ucsa_mgmt_ip>>	Cisco UCS Fabric Interconnect (FI) out-of-band management IP address	
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address	
<<var_vsan_a_id>>	VSAN ID for Fabric A (101 is used)	
<<var_vsan_B_id>>	VSAN ID for Fabric B (102 is used)	
<<var_fabric_a_fcoe_vlan_id >>	Fabric ID for Fabric A (101 is used)	
<<var_fabric_b_fcoe_vlan_id >>	Fabric ID for Fabric B (102 is used)	
<<var_In-band_mgmtblock_net >>	Block of IP addresses for KVM access for UCS	
<<var_vmhost_Oracle_01_ip>>	VMware ESXi host 01 in-band management IP	
<<var_vmhost_Oracle_01_2nd_ip>>	VMware ESXi host 01 secondary in-band management IP	
<<var_vmotion_vlan_id_ip_host-01>>	vMotion VLAN IP address for ESXi host 01	
<<var_vmotion_vlan_id_mask_host-01>>	vMotion VLAN netmask for ESXi host 01	
The last four variables should be repeated for all ESXi hosts.		

Example 3-2 shows the volumes that are mapped to VMware hosts from IBM FlashSystem V9000.

Example 3-2 Volumes mapped

```

IBM_FlashSystem:VersaStack:superuser>lshostvdiskmap 6
id name          SCSI_id vdisk_id vdisk_name          vdisk_UID
IO_group_id IO_group_name
6 VM-Host-Oracle-01 0 29 VM-Host-Oracle-01
600507680C8181138800000000000021 0 io_grp0

```

```
6 VM-Host-Oracle-01 1 33 ITS0_repository
600507680C818113880000000000002C 0 io_grp0
6 VM-Host-Oracle-01 2 35 ITS0_Infra_datastore1
600507680C8181138800000000000027 0 io_grp0
6 VM-Host-Oracle-01 3 36 ITS0_Infra_datastore2
600507680C8181138800000000000028 0 io_grp0
6 VM-Host-Oracle-01 4 37 ITS0_Infra_datastore3
600507680C8181138800000000000029 0 io_grp0
6 VM-Host-Oracle-01 5 38 ITS0_Infra_datastore4
600507680C818113880000000000002A 0 io_grp0
6 VM-Host-Oracle-01 6 34 ITS0_Oracle_redo_log
600507680C818113880000000000002D 0 io_grp0
6 VM-Host-Oracle-01 7 39 ITS0_OracleVM_boot_01
600507680C8181138800000000000032 0 io_grp0
6 VM-Host-Oracle-01 8 43 ITS0_Infra_datastore5
600507680C8181138800000000000036 0 io_grp0
6 VM-Host-Oracle-01 9 44 ITS0_Infra_datastore6
600507680C8181138800000000000037 0 io_grp0
6 VM-Host-Oracle-01 10 46 big_test
600507680C8181138800000000000039 0 io_grp0
IBM_FlashSystem:VersaStack:superuser>
```



Planning for Oracle Real Application Clusters

This chapter describes some of the considerations and assumptions that are followed during the design of the Oracle Real Application Clusters (RAC) installation on VersaStack.

This chapter includes the following sections:

- ▶ Design considerations
- ▶ Network and Storage Design

4.1 Design considerations

The goal of this section is to come up with a simple and efficient Oracle RAC database design that is suited for a VersaStack solution. The major design considerations of the suggested architecture are described in the following subsections. These assumptions are influenced by several factors, including the status of the technology and the specific business requirements that drive each specific solution.

The upcoming sections detail the design considerations from different layers of the architectural stack.

4.1.1 Database workload

The entire architecture is designed for an online transaction processing (OLTP) workload, which is characterized by small number of random I/Os. Log I/O is the most critical component, as it directly affects the transaction latency. Memory mitigates the I/O pressure on the storage system. However, beyond a certain threshold, increasing memory might not yield any noticeable benefit. Certain OLTP workloads have reporting or End-Of-Day (EOD) consolidation jobs in the mix. For these reporting and EOD jobs, I/O capacity must be carefully evaluated to ensure that such workloads are not affecting regular production OLTP transactions. Many of the reporting and batch jobs use temporary database space. To provide optimal performance for these workloads, use disks with highest performance to store redo log files in Oracle database.

4.1.2 Server virtualization

The database deployment is built on server virtualization by using VMware ESXi. This design provides an efficient and flexible back end for hosting Oracle RAC database transactional workloads. Each of the virtual machines (VMs) hosting the Oracle RAC database instances should be configured with the optimal computational and storage resources to suit the workload. Typical OLTP workloads are not CPU-intensive. For a virtualized database platform, you can start with four vCPUs and scale when the aggregate usage of those vCPUs crosses the threshold that is set by internal IT practices.

4.1.3 Database availability

Oracle RAC enables continuous and uninterrupted database operations by providing multiple database instances running on different nodes. If one of the database instances fails, the Oracle server moves the services from the failed instance to a surviving alternate instance. Meanwhile, the VMware hypervisor back end provides a rich medium for VM high availability and optimal performance by using the VMware HA and Distributed Resource Scheduler (DRS) features. However, on the Oracle VMs, anti-affinity rules are set to prevent VMs from migrating under the HA/DRS feature. This feature ensures that VMs are not placed on the same ESXi host, and that VMs are not migrated to a different ESXi host.

4.1.4 Quality of service and network segregation

The network traffic within the proposed architecture is segregated to ensure maximum bandwidth availability. Each of the network interfaces that are defined follow a quality of service (QoS) policy, which is assumed to provide the intended performance and functions.

4.1.5 Network availability

All the networking elements in the architecture have a high amount of redundancy. All the network paths are configured to ensure aggregated bandwidth for the traffic and resiliency against individual failures.

4.2 Network and Storage Design

This section introduces the network design and storage design for a sample Oracle RAC database environment to be built later following the considerations mentioned in 4.1, “Design considerations” on page 34.

The sample Oracle RAC database environment is built on four virtual machines, running on four ESXi hosts, and the hardware of each ESXi host is a Cisco UCS B200-M4 Blade.

4.2.1 Network Design

Oracle RAC database has the following network requirements and suggestions on each node in the cluster:

- ▶ Each node must have at least two network adapters: One for the public network interface and the other for the private network interface or interconnect.
- ▶ The public interface ports and private interface ports from each node connect to the public switch and private switch, respectively.
- ▶ It is suggest to use redundant network cards, bonded as one Ethernet port, to serve as a public or private network interface.
- ▶ The switches and network interface adapters must be at least 1 GbE, with 10 GbE suggested.
- ▶ Public network and private network are on different VLANs.

Following Oracle’s requirements, and the considerations mentioned in 4.1, “Design considerations” on page 34, the designed network topology in this environment is as shown in Figure 4-1 on page 36:

- ▶ Configure NIC teaming on each ESXi host to provide network high availability in case one of the Ethernet adapters goes down.
- ▶ Create two host networkings on each ESXi host: One for public network and one for private network, with different VLAN IDs configured.
- ▶ Create two Ethernet adapters on each virtual machine: One adapter connects to public network, and the other adapter connects to private network.
- ▶ Enable jumbo frames on private Ethernet interfaces on each node to improve performance, and enable jumbo frames on virtual switch on each ESXi host.




Figure 4-1 Network design

4.2.2 Single Client Access Name (SCAN)

Single Client Access Name (SCAN) is a feature used in Oracle RAC environments that provides a single name for clients to access any Oracle database running in a cluster. One of the benefits of SCAN is that the connect information does not need to change while Oracle RAC scales to more or less nodes in the cluster. For more information about SCAN, see the following white paper:

<http://www.oracle.com/technetwork/products/clustering/overview/scan-129069.pdf>

SCAN is used in this Oracle RAC database environment. SCAN also supports at least one and up to three IP addresses with one single name. This environment uses three IP addresses. Therefore, registering these three IP addresses with one name in the DNS server is required.

4.2.3 Storage Design

IBM FlashSystem V9000 supports various volumes, as shown in Figure 4-2, and each volume type has its own usage scenarios. Three types of volumes are used in this environment.




Figure 4-2 Flexible volumes in FlashSystem V9000

A generic volume is a fully allocated volume. It is used for Oracle redo log files and OCR files because these files are frequently accessed, but the capacity growth of these files is flat.

A thin-provisioned volume has a virtual capacity and a real capacity. Virtual capacity is the volume storage capacity that is available to a host. Real capacity is the storage capacity that is allocated to this volume from a storage pool. The virtual capacity can be much larger than the real capacity. Thin provisioning eliminates almost all white space, helping to avoid the poor usage rates that occur in the traditional storage allocation method where large pools of storage capacity are allocated to individual servers but remain unused (not written to). Thin-provisioned volume is used for storing operating system data of each virtual machine in this environment.

A compressed volume is a special type of volume where data is compressed as it is written to disk, saving additional space. Database information is stored in table space files. It is common to see high compression ratios in database volumes. Examples of databases that can greatly benefit from Real-time Compression are IBM DB2, Oracle, and Microsoft SQL Server. Expected compression ratios are 50 - 80%. Compressed volume is used for storing Oracle data files in this environment.

The storage design in this environment is shown in Figure 4-3.




Figure 4-3 Storage design

4.2.4 Oracle Automatic Storage Management (ASM) design

Oracle Automatic Storage Management (ASM) is a volume manager and file system that is designed for Oracle database files, which is integrated in the Oracle Grid Infrastructure. Oracle ASM provides three redundancy types while managing disks:

- ▶ External redundancy: ASM relies on the storage system to provide redundancy. All disks must be located to successfully mount the disk group.

Note: Use external redundancy on SAN disks to save disk space, and it is not suggested to use it on local disks because it has a data loss risk during any local disk member failure.

- ▶ Normal redundancy: ASM provides two-way mirroring. By default, all files are mirrored across different disks so that there are two copies of every data extent.
- ▶ High redundancy: ASM provides triple mirroring by default.

Because all data is protected by IBM Flashsystem V9000 RAID, use external redundancy while configuring disk groups in Oracle ASM.

Oracle ASM supports different allocation unit sizes while creating disk groups. This value can be 1, 2, 4, 8, 16, 32, or 64 MB, depending on the specific disk group compatibility level. Usually larger allocation unit sizes provide performance advantages for data warehouse applications that use large sequential reads. The allocation unit size of ASM disk groups in this environment is set to 4 MB.



Physical infrastructure

This chapter covers the physical infrastructure for VersaStack.

This chapter includes the following sections:

- ▶ VersaStack cabling
- ▶ Storage compatibility and interoperability
- ▶ VersaStack system build process

5.1 VersaStack cabling

The information in this section is provided as a reference for cabling the equipment in a VersaStack environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

This book assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces are used in various configuration steps.

Be sure to follow the cabling directions in this section. Failure to do so will result in changes to the deployment procedures that follow because specific port locations are mentioned here.

It is possible to order IBM FlashSystem V9000 storage systems in a different configuration by using Cisco Nexus 9000 and IBM FlashSystem V9000. Before starting, be sure that the configuration matches the descriptions in the tables and the diagrams in this section.

Figure 5-1 shows the cabling diagram for a VersaStack configuration that uses the Cisco Nexus 9000 and IBM FlashSystem V9000 storage system. For more information about IBM FlashSystem V9000 enclosure cabling information, see the following website:

https://www.ibm.com/support/knowledgecenter/STKMVQ_7.6.0/com.ibm.storage.vflashsystem9000.7.6.doc/FlashSystem_V9000_welcome.htm?lang=en




Figure 5-1 VersaStack cabling diagram

Figure 5-2 shows the VersaStack management cabling. IBM FlashSystem V9000 has redundant management connections. One path is through the dedicated out-of-band management switch, and the secondary path is through the in-band management path going up through the IBM FlashSystem V9000 to the production network.




Figure 5-2 VersaStack management cabling

The details of all these connections have been tabulated and described in the section titled “VersaStack Cabling” in *VersaStack for Data Center with All-Flash Storage and VMware vSphere 6.0 Deployment Guide: Design Guide for Cisco Unified Computing System 3.1 and IBM FlashSystem V9000 with VMware vSphere 6.0 Update 1a*, which is available at:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/Versastack_vmw_6_flash_design.html

5.2 Storage compatibility and interoperability

The IBM System Storage Interoperation Center (SSIC) provides information about supported external hardware and software for the specific IBM FlashSystem V9000 version.

Make sure that the hardware and software components are supported by the version of Storwize V7000 storage system that you plan to install by checking the SSIC. Click **IBM System Storage Enterprise Flash**, then click **FlashSystem V9000 Host Attachment or FlashSystem V9000 Storage Controller Support**.

Software and hardware limitations for IBM FlashSystem V9000 running firmware v7.6.x can be found at:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1005242>

Detailed information about supported hardware, device driver, firmware, and software level information can be found at:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1005419>

5.3 VersaStack system build process

Figure 5-3 depicts the VersaStack build process for the environment used for this book.




Figure 5-3 VersaStack build process



IBM FlashSystem V9000 storage configuration

This chapter describes the steps that are necessary to create and configure the storage volumes for Oracle RAC (four nodes) from IBM FlashSystem V9000 storage system in the VersaStack environment.

As a prerequisite, the user must have configured the IBM FlashSystem V9000 as described in the “Storage Configuration” section of the document *VersaStack for Data Center with All-Flash Storage and VMware vSphere 6.0 Deployment Guide: Design Guide for Cisco Unified Computing System 3.1 and IBM FlashSystem V9000 with VMware vSphere 6.0 Update 1a* available here:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/Versastack_vmw_6_flash_design.html

This chapter includes the following sections:

- ▶ Volume layout for four node Oracle-RAC
- ▶ Volume creation and mapping

6.1 Volume layout for four node Oracle-RAC

For the Oracle-RAC configuration, a total of six volumes were created. Table 6-1 describes the type, size, and purpose of each volume created.

Table 6-1 Volumes layout for Oracle-RAC configuration

Volume Name	Volume Type	Volume Size	Volume Purpose	Volume Quantity
ITSO_VMOS_vol	Thin-provisioned	1 TB	To store guest OSs for VMware environment	1
ITSO_DB_vol_<x>	Compressed	2 TB each	For Oracle database files	3
ITSO_DB_LOG_vol	Regular	100 GB	For Oracle Redo logs	1

6.2 Volume creation and mapping

This section illustrates the steps for volume creation and mapping them to VMware hosts. The sequence and steps provided here are for one volume. However, it can be repeated by specifying the relevant parameters, such as volume size, volume type, and volume name, for the wanted number of volumes.

1. Open a browser and go to https://<V9000_Management_IP>.
2. Log in as “superuser” with the password for that user.
3. As shown in Figure 6-1, click **Volumes**.




Figure 6-1 Volumes pane from the initial window

- Select **Volumes**, which shows all the volumes that exist in the system as shown in Figure 6-2.

Name	State	Pool	UID	Host Mappings	Capacity
IOmeter_Compressed4	✓ Online	mdiskgrp0	600507680C81811388000000000019	No	20.00 GiB
IOmeter_Compressed5	✓ Online	mdiskgrp0	600507680C8181138800000000001A	No	20.00 GiB
IOmeter_Compressed6	✓ Online	mdiskgrp0	600507680C8181138800000000001B	No	20.00 GiB
IOmeter_Compressed7	✓ Online	mdiskgrp0	600507680C8181138800000000001C	No	20.00 GiB
IOmeter_regular0	✓ Online	mdiskgrp0	600507680C8181138800000000000D	Yes	20.00 GiB
IOmeter_regular1	✓ Online	mdiskgrp0	600507680C8181138800000000000E	Yes	20.00 GiB
IOmeter_regular2	✓ Online	mdiskgrp0	600507680C8181138800000000000F	Yes	20.00 GiB
IOmeter_regular3	✓ Online	mdiskgrp0	600507680C81811388000000000010	Yes	20.00 GiB
IOmeter_regular4	✓ Online	mdiskgrp0	600507680C81811388000000000011	Yes	20.00 GiB
IOmeter_regular5	✓ Online	mdiskgrp0	600507680C81811388000000000012	Yes	20.00 GiB
IOmeter_regular6	✓ Online	mdiskgrp0	600507680C81811388000000000013	Yes	20.00 GiB
IOmeter_regular7	✓ Online	mdiskgrp0	600507680C81811388000000000014	Yes	20.00 GiB
ITSO_DB_LOG_vol	✓ Online	mdiskgrp0	600507680C81811388000000000049	Yes	100.00 GiB
ITSO_DB_vol_1	✓ Online	mdiskgrp0	600507680C81811388000000000046	Yes	2.00 TiB
ITSO_DB_vol_2	✓ Online	mdiskgrp0	600507680C81811388000000000047	Yes	2.00 TiB
ITSO_DB_vol_3	✓ Online	mdiskgrp0	600507680C81811388000000000048	Yes	2.00 TiB
ITSO_VMOS_vol	✓ Online	mdiskgrp0	600507680C81811388000000000045	Yes	1.00 TiB
ITSO_repository	✓ Online	mdiskgrp0	600507680C8181138800000000002C	Yes	200.00 GiB
Infra-ESXi-iSCSI-ACI-05	✓ Online	mdiskgrp0	600507680C8181138800000000001D	Yes	10.00 GiB
Infra-ESXi-iSCSI-ACI-06	✓ Online	mdiskgrp0	600507680C8181138800000000001E	Yes	10.00 GiB
Infra_datastore1	✓ Online	mdiskgrp0	600507680C81811388000000000005	Yes	2.00 TiB
Infra_datastore2	✓ Online	mdiskgrp0	600507680C81811388000000000006	Yes	2.00 TiB
Test	✓ Online	mdiskgrp0	600507680C8181138800000000004B	Yes	35.00 GiB
VDBench_compressed	✓ Online	mdiskgrp0	600507680C8181138800000000000B	Yes	2.00 TiB
VDBench_regular	✓ Online	mdiskgrp0	600507680C8181138800000000000C	Yes	2.00 TiB
VM-Host-Infra-01	✓ Online	mdiskgrp0	600507680C81811388000000000000	Yes	40.00 GiB
VM-Host-Infra-02	✓ Online	mdiskgrp0	600507680C81811388000000000001	Yes	40.00 GiB
VM-Host-Infra-03	✓ Online	mdiskgrp0	600507680C81811388000000000002	Yes	40.00 GiB
VM-Host-Infra-04	✓ Online	mdiskgrp0	600507680C81811388000000000003	Yes	40.00 GiB
VM-Host-Oracle-01	✓ Online	mdiskgrp0	600507680C81811388000000000021	Yes	40.00 GiB
VM-Host-Oracle-02	✓ Online	mdiskgrp0	600507680C81811388000000000022	Yes	40.00 GiB
VM-Host-Oracle-03	✓ Online	mdiskgrp0	600507680C81811388000000000023	Yes	40.00 GiB
VM-Host-Oracle-04	✓ Online	mdiskgrp0	600507680C81811388000000000024	Yes	40.00 GiB
infra_swap	✓ Online	mdiskgrp0	600507680C81811388000000000007	Yes	500.00 GiB

Figure 6-2 List of volumes

5. Select **Create Volumes** as shown in Figure 6-3.

+ Create Volumes Actions Filter					
Name	State	Pool	UID	Host Mappings	Capacity
ACI-IOM	✓ Online	mdiskgrp0	600507680C8181138800000000000038	Yes	100.00 GiB
ACI-IOM-1	✓ Online	mdiskgrp0	600507680C818113880000000000003A	Yes	20.00 GiB
ACI-IOM-2	✓ Online	mdiskgrp0	600507680C818113880000000000003B	Yes	20.00 GiB
ACI-IOM-3	✓ Online	mdiskgrp0	600507680C818113880000000000003C	Yes	20.00 GiB
ACI-IOM-4	✓ Online	mdiskgrp0	600507680C818113880000000000003D	Yes	20.00 GiB
ACI-IOM-5	✓ Online	mdiskgrp0	600507680C818113880000000000003E	Yes	20.00 GiB
ACI-IOM-6	✓ Online	mdiskgrp0	600507680C818113880000000000003F	Yes	20.00 GiB
ACI-IOM-7	✓ Online	mdiskgrp0	600507680C8181138800000000000040	Yes	20.00 GiB
ACI-IOM-8	✓ Online	mdiskgrp0	600507680C8181138800000000000041	Yes	20.00 GiB
ACI-Infra-Datastore	✓ Online	mdiskgrp0	600507680C8181138800000000000020	Yes	1.00 TiB
ACI-Infra-Swap	✓ Online	mdiskgrp0	600507680C818113880000000000001F	Yes	200.00 GiB
IOmeter_Compressed0	✓ Online	mdiskgrp0	600507680C8181138800000000000015	No	20.00 GiB
IOmeter_Compressed1	✓ Online	mdiskgrp0	600507680C8181138800000000000016	No	20.00 GiB
IOmeter_Compressed2	✓ Online	mdiskgrp0	600507680C8181138800000000000017	No	20.00 GiB
IOmeter_Compressed3	✓ Online	mdiskgrp0	600507680C8181138800000000000018	No	20.00 GiB
IOmeter_Compressed4	✓ Online	mdiskgrp0	600507680C8181138800000000000019	No	20.00 GiB
IOmeter_Compressed5	✓ Online	mdiskgrp0	600507680C818113880000000000001A	No	20.00 GiB
IOmeter_Compressed6	✓ Online	mdiskgrp0	600507680C818113880000000000001B	No	20.00 GiB
IOmeter_Compressed7	✓ Online	mdiskgrp0	600507680C818113880000000000001C	No	20.00 GiB
IOmeter_regular0	✓ Online	mdiskgrp0	600507680C818113880000000000000D	Yes	20.00 GiB
IOmeter_regular1	✓ Online	mdiskgrp0	600507680C818113880000000000000E	Yes	20.00 GiB
IOmeter_regular2	✓ Online	mdiskgrp0	600507680C818113880000000000000F	Yes	20.00 GiB
IOmeter_regular3	✓ Online	mdiskgrp0	600507680C8181138800000000000010	Yes	20.00 GiB
IOmeter_regular4	✓ Online	mdiskgrp0	600507680C8181138800000000000011	Yes	20.00 GiB
IOmeter_regular5	✓ Online	mdiskgrp0	600507680C8181138800000000000012	Yes	20.00 GiB
IOmeter_regular6	✓ Online	mdiskgrp0	600507680C8181138800000000000013	Yes	20.00 GiB
IOmeter_regular7	✓ Online	mdiskgrp0	600507680C8181138800000000000014	Yes	20.00 GiB
ITSO_repository	✓ Online	mdiskgrp0	600507680C818113880000000000002C	Yes	200.00 GiB
Infra-ESXi-ISCST-ACI-05	✓ Online	mdiskgrp0	600507680C818113880000000000001D	Yes	10.00 GiB
Infra-ESXi-ISCST-ACI-06	✓ Online	mdiskgrp0	600507680C818113880000000000001E	Yes	10.00 GiB
Infra_datastore1	✓ Online	mdiskgrp0	600507680C8181138800000000000005	Yes	2.00 TiB
Infra_datastore2	✓ Online	mdiskgrp0	600507680C8181138800000000000006	Yes	2.00 TiB
VDBench_compressed	✓ Online	mdiskgrp0	600507680C818113880000000000000B	Yes	2.00 TiB
VDBench_regular	✓ Online	mdiskgrp0	600507680C818113880000000000000C	Yes	2.00 TiB

Figure 6-3 Options for Create Volumes

6. Select **Custom** as shown in Figure 6-4.




Figure 6-4 Custom volume creation

7. Specify the **Quantity**, **Capacity**, **Capacity savings**, and **Name** parameters as shown in Figure 6-5.




Figure 6-5 Volume specifications

8. Click **Create**, which completes the task of creating the volume as shown in Figure 6-6.




Figure 6-6 Volume Creation completion

9. From the Volumes window, select the volume that was just created as shown in Figure 6-7.

Name	State	Pool	UID	Host Mappings	Capacity
ACI-IOM	✓ Online	mdiskgrp0	600507680C8181138800000000000038	Yes	100.00 GiB
ACI-IOM-1	✓ Online	mdiskgrp0	600507680C818113880000000000003A	Yes	20.00 GiB
ACI-IOM-2	✓ Online	mdiskgrp0	600507680C818113880000000000003B	Yes	20.00 GiB
ACI-IOM-3	✓ Online	mdiskgrp0	600507680C818113880000000000003C	Yes	20.00 GiB
ACI-IOM-4	✓ Online	mdiskgrp0	600507680C818113880000000000003D	Yes	20.00 GiB
ACI-IOM-5	✓ Online	mdiskgrp0	600507680C818113880000000000003E	Yes	20.00 GiB
ACI-IOM-6	✓ Online	mdiskgrp0	600507680C818113880000000000003F	Yes	20.00 GiB
ACI-IOM-7	✓ Online	mdiskgrp0	600507680C8181138800000000000040	Yes	20.00 GiB
ACI-IOM-8	✓ Online	mdiskgrp0	600507680C8181138800000000000041	Yes	20.00 GiB
ACI-Infra-Datastore	✓ Online	mdiskgrp0	600507680C8181138800000000000020	Yes	1.00 TiB
ACI-Infra-Swap	✓ Online	mdiskgrp0	600507680C818113880000000000001F	Yes	200.00 GiB
IOMeter_Compressed0	✓ Online	mdiskgrp0	600507680C8181138800000000000015	No	20.00 GiB
IOMeter_Compressed1	✓ Online	mdiskgrp0	600507680C8181138800000000000016	No	20.00 GiB
IOMeter_Compressed2	✓ Online	mdiskgrp0	600507680C8181138800000000000017	No	20.00 GiB
IOMeter_Compressed3	✓ Online	mdiskgrp0	600507680C8181138800000000000018	No	20.00 GiB
IOMeter_Compressed4	✓ Online	mdiskgrp0	600507680C8181138800000000000019	No	20.00 GiB
IOMeter_Compressed5	✓ Online	mdiskgrp0	600507680C818113880000000000001A	No	20.00 GiB
IOMeter_Compressed6	✓ Online	mdiskgrp0	600507680C818113880000000000001B	No	20.00 GiB
IOMeter_Compressed7	✓ Online	mdiskgrp0	600507680C818113880000000000001C	No	20.00 GiB
IOMeter_regular0	✓ Online	mdiskgrp0	600507680C8181138800000000000000	Yes	20.00 GiB
IOMeter_regular1	✓ Online	mdiskgrp0	600507680C818113880000000000000E	Yes	20.00 GiB
IOMeter_regular2	✓ Online	mdiskgrp0	600507680C818113880000000000000F	Yes	20.00 GiB
IOMeter_regular3	✓ Online	mdiskgrp0	600507680C8181138800000000000010	Yes	20.00 GiB
IOMeter_regular4	✓ Online	mdiskgrp0	600507680C8181138800000000000011	Yes	20.00 GiB
IOMeter_regular5	✓ Online	mdiskgrp0	600507680C8181138800000000000012	Yes	20.00 GiB
IOMeter_regular6	✓ Online	mdiskgrp0	600507680C8181138800000000000013	Yes	20.00 GiB
IOMeter_regular7	✓ Online	mdiskgrp0	600507680C8181138800000000000014	Yes	20.00 GiB
ITSO_DB_LOG_vol	✓ Online (formatting)	mdiskgrp0	600507680C8181138800000000000049	No	100.00 GiB
ITSO_DB_vol_1	✓ Online	mdiskgrp0	600507680C8181138800000000000046	No	2.00 TiB
ITSO_DB_vol_2	✓ Online	mdiskgrp0	600507680C8181138800000000000047	No	2.00 TiB
ITSO_DB_vol_3	✓ Online	mdiskgrp0	600507680C8181138800000000000048	No	2.00 TiB
ITSO_WMOS_vol	✓ Online	mdiskgrp0	600507680C8181138800000000000045	No	1.00 TiB
ITSO_repository	✓ Online	mdiskgrp0	600507680C818113880000000000002C	Yes	200.00 GiB
Infra-ESXi-iSCSI-ACI-05	✓ Online	mdiskgrp0	600507680C818113880000000000001D	Yes	10.00 GiB

Figure 6-7 Volume listing

10. Right-click the volume name that was just created as shown in Figure 6-8.




Figure 6-8 Selected volume

11. Click **Map to Host** as shown in Figure 6-9.




Figure 6-9 Map to host

12. Select all the hosts that this volume needs to be mapped to as shown in Figure 6-10.




Figure 6-10 Select hosts that need to have this volume mapped

13. Click **Map** as shown in Figure 6-11.




Figure 6-11 Map the volumes to hosts

14. Mapping of the selected volume to the hosts is completed as shown in Figure 6-12.




Figure 6-12 Mapping completed

For the Oracle-RAC configuration, repeat the steps 1 - 14 for all the volumes as listed in Table 6-1 on page 44. After following the steps described above, all the wanted VMware hosts that are going to be part of the Oracle-RAC configuration will have the required volumes mapped.

Complete these steps from the IBM FlashSystem V9000 GUI to verify the host-to-volume mapping:

1. Click the **Home** icon on the GUI.
2. Click **Hosts** as shown in Figure 6-13.




Figure 6-13 Hosts window

3. Click **Volumes by Host** as shown in Figure 6-14.




Figure 6-14 Volumes by Host window

4. Select the host from under the Host Filter pane as shown in Figure 6-15.




Figure 6-15 Host Filter pane

5. After selecting the host, the right pane displays the volumes that are mapped to that host as shown in Figure 6-16.




Figure 6-16 Volumes mapped to VMware host for Oracle-RAC

Repeat steps 1 - 5 for all the VMware hosts planned for Oracle-RAC to ensure that the required volumes have been mapped to them.



Virtual infrastructure configuration

This chapter provides detailed instructions for configuring VMware ESXi hosts and creating virtual machines in a VersaStack environment that is used for Oracle Database installation later.

This chapter includes the following sections:

- ▶ ESXi Configuration
- ▶ Create and modify virtual machines
- ▶ Considerations of installing Red Hat Enterprise Linux

7.1 ESXi Configuration

These sections show how to configure ESXi hosts in a VersaStack environment. Four ESXi hosts are available in this environment and, unless specified, the instructions in this section need to be executed on all four ESXi hosts.

7.1.1 Scan Disk Devices

After FlashSystem V9000 volumes are mapped to ESXi hosts, complete these steps to scan disk devices:

1. Log in using the Administrator@vsphere.local user from the vSphere Web Client.
2. Click **Hosts and Clusters** in the left pane, as shown in Figure 7-1.




Figure 7-1 VMware vCenter GUI

Note: Run all the actions with vSphere Web Client using Administrator@vsphere.local or an equivalent user.

- Click the cluster name **VersaStack_Oracle** to expand the ESXi host list, and click the first ESXi host **vm-host-oracle-01** in the inventory, as shown in Figure 7-2.




Figure 7-2 ESXi server

- Click **Manage** → **Storage** → **Storage Devices** to list current disk devices, and click **Rescan All Storage Adapters** (the second icon from left) to start a device scan on all storage adapters, as shown in Figure 7-3.




Figure 7-3 List storage devices

5. A window opens that prompts you to choose scan options. Select **Scan for new Storage Devices** and **Scan for new VMFS Volumes**, and click **OK** to continue, as shown in Figure 7-4.



Figure 7-4 Rescan storage options

6. After the rescan action is finished, the newly mapped volumes will show up in the list, as shown in Figure 7-5.




Figure 7-5 List storage devices

7.1.2 Mark FlashSystem V9000 Disks as Flash Disks

The newly mapped volumes cannot be recognized as Flash disks by VMware ESXi host automatically. For performance reasons, it is suggested that you mark these FlashSystem V9000 volumes as Flash disks.

To mark FlashSystem V9000 volumes as Flash disks, complete these procedures:

1. Select all the FlashSystem V9000 volumes that are listed as HDD in the list, and click **Mark Flash Disks** (the third icon from right) to mark these selected disks as Flash disks, as shown in Figure 7-6.

The screenshot shows a table titled "Storage Devices" with columns: Name, Type, Capacity, Operational State, Hardware Acceleration, Drive Type, and Transport. The first two rows are marked as "Flash" and the remaining five as "HDD". A tooltip points to the third icon from the right in the toolbar, stating "Marks the selected disks as flash disks."

Name	Type	Capacity	Operational State	Hardware Acceleration	Drive Type	Transport
IBM Fibre Channel Disk (naa.600507...)	disk	200 GB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	200 GB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	1 TB	Attached	Supported	HDD	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	HDD	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	HDD	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	HDD	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	HDD	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	100 GB	Attached	Supported	HDD	Fibre Channel

Figure 7-6 Mark selected disks as Flash disks

2. A warning window prompts you to confirm this action. Click **Yes** to continue, as shown in Figure 7-7.



Figure 7-7 Warning to confirm marking Flash disks

3. After the marking disks action is completed, all the FlashSystem V9000 volumes will show up as Flash disks in the list, as shown in Figure 7-8.

The screenshot shows the same "Storage Devices" table as in Figure 7-6, but now all seven rows have "Flash" listed in the Drive Type column. The "HDD" filter icon in the toolbar is no longer active.

Name	Type	Capacity	Operational State	Hardware Acceleration	Drive Type	Transport
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	40 GB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	2 TB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	100 GB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	1 TB	Attached	Supported	Flash	Fibre Channel
IBM Fibre Channel Disk (naa.600507...)	disk	200 GB	Attached	Supported	Flash	Fibre Channel

Figure 7-8 List storage devices

7.1.3 Modify Path Selection Policy

For each storage device, VMware ESXi host sets the Path Selection Policy (PSP) based on the defined claim rules, usually three PSPs are supported in ESXi, as shown in Table 7-1.

Table 7-1 ESXi Path Selection Policy

Policy	Meaning
Fixed	The host uses the designated preferred path, if it has been configured.
Most Recently Used	The host selects the path that it used most recently. When the path becomes unavailable, the host selects an alternative path.
Round Robin	The host uses an automatic path selection algorithm rotating through all active paths when connecting to storage subsystems.

For FlashSystem V9000 volumes, it is suggested to use **Round Robin** Path Selection Policy. To modify PSP for a FlashSystem V9000 volume, complete these steps:

1. Select one FlashSystem V9000 volume in storage device list. Navigate to Properties tab of Device Details pane.
2. The Path Selection Policy configured for this volume is **Most Recently Used**, as shown in Figure 7-9. To change the Path Selection Policy, click **Edit Multipathing** button.




Figure 7-9 Show disk device details

- From the dialog, it shows current Path Selection Policy, and all available paths and active paths to the disk, as shown in Figure 7-10.




Figure 7-10 List current Path Selection Policy

- Click the Path Selection Policy drop list, and select **Round Robin (VMware)** as the new Path Selection Policy. Click **OK** to continue, as shown in Figure 7-11.




Figure 7-11 Change disk Path Selection Policy

- The change is effective immediately and shown on the Properties tab, as shown in Figure 7-12.




Figure 7-12 Disk Properties tab

- Repeat these steps to change Path Selection Policy for all FlashSystem V9000 disks.

Tip: Changing Path Selection Policy through GUI can take a long time if there are a few FlashSystem V9000 volumes. In this case, running a command on ESXi host would be a quicker way.

Log in to the ESXi host using root user, and run the following command:

```
[root@VM-Host-Oracle-01:~] for i in `esxcli storage nmp device list | grep '^naa.60050768'`; do esxcli storage nmp device set --device $i --psp VMW_PSP_RR; done
```

7.1.4 Create datastore

According to our planning, three datastores need to be created in this environment as shown in Table 7-2.

Table 7-2 Datastore list

Datastore	Capacity	Usage
VM_OS	1 TB * 1	Boot disks of Linux virtual machines
OCR_REDO	100 GB * 1	Oracle RAC OCR and Redo log files
DATABASE	2 TB * 3	Oracle RAC database data files

To create datastores in the VMware ESXi host, complete these steps. It is only necessary to perform datastore creation actions on one ESXi host.

- From the vSphere Web Client, select **vm-host-oracle-01** in the inventory.
- Click **Related Objects** → **Datastores** to list current datastores.

3. Click **Create a New Datastore** (the first icon from left) to add a datastore, as shown in Figure 7-13.




Figure 7-13 List datastores

4. Select **VMFS** as the type of new datastore, as shown in Figure 7-14, and click **Next** to continue.




Figure 7-14 Choose new datastore type

- Input the name of new datastore, and select the disk device to be used to create a datastore, as shown in Figure 7-15. Click **Next** to continue.




Figure 7-15 Select disk device

- Select the **Use all available partitions** option, as shown in Figure 7-16. Click **Next** to continue.




Figure 7-16 Choose partition configuration

- From the datastore configuration summary window, as shown in Figure 7-17, click **Finish** to finish datastore configuration.




Figure 7-17 Datastore configuration summary

- After VMware ESXi host finishes datastore creation, the new datastore will be listed in Datastores tab, as shown in Figure 7-18.




Figure 7-18 List datastores

9. For the datastore, which is planned to have multiple disks, right-click the datastore name and select **Increase Datastore Capacity**, as shown in Figure 7-19.




Figure 7-19 Increase Datastore Capacity

10. Select the disk to be added into the datastore, as shown in Figure 7-20, and click **Next** to continue.




Figure 7-20 Select disk device

11. Select the **Use all available partitions** option, as shown in Figure 7-21, and click **Next** to continue.




Figure 7-21 Choose partition configuration

12. In the datastore configuration summary window, as shown in Figure 7-22, click **Finish** to complete the increasing datastore capacity process.




Figure 7-22 Increase Datastore Capacity summary

7.1.5 Configure NIC Teaming

NIC teaming is a feature that is provided in ESXi to increase network capacity for the virtual switch through teaming multiple Ethernet network adapters. It also provides failover if one of the adapters in the team goes down. It is suggested to configure NIC teaming on ESXi hosts in a VersaStack environment.

Complete these steps to configure NIC teaming on ESXi host from vSphere web client:

1. From the vSphere Web Client, select **vm-host-oracle-01** in the inventory.
2. Click **Manage** → **Networking** → **Virtual Switches**, then select **vSwitch0** from the list, as shown in Figure 7-23.




Figure 7-23 Virtual switches

3. Click **Manage the Physical Network Adapters** (the third icon from left) to manage the physical network adapters that are connected to **vSwitch0**.

4. Click **Add Adapters** (the first icon from left) to add adapters to virtual switch **vSwitch0**, as shown in Figure 7-24.




Figure 7-24 Manage physical network adapters for vSwitch0

5. Select network adapter **vmnic1** and select **Active Adapters** in **Failover order group**, as shown in Figure 7-25. Click **OK** to continue.




Figure 7-25 Add physical adapters to the Switch

- Confirm **vmnic1** is added in vSwitch0 active adapters, as shown in Figure 7-26. Click **OK** to continue.




Figure 7-26 Manage physical network adapters for vSwitch0

- The change is reflected on the virtual switch network diagram, as shown in Figure 7-27.




Figure 7-27 List virtual switches

VMware ESXi supports several load balancing policies. Make sure to use the default **Route based on originating virtual port** policy in ESXi hosts in VersaStack environment. To check the load balancing policy, click the **Edit Settings** icon (the fifth icon from left). A dialog is prompted as shown in Figure 7-28. Click **Teaming and Failover** → **Load Balancing** to check the current load balancing policy. Make sure that **Route based on originating virtual port** is used.




Figure 7-28 List load balancing policy

7.1.6 Enable Jumbo Frames

Enabling jumbo frames allow ESXi host to send larger frames out onto the physical network, which could reduce the CPU load caused by transferring data. To enable jumbo frames on ESXi virtual switches, complete these steps:

1. From the vSphere Web Client, select **vm-host-oracle-01** in the inventory.
2. Click **Manage** → **Networking** → **Virtual Switches**, select the **vSwitch0** from the list, as shown in Figure 7-27 on page 68.

3. Click **View Switch Settings** (the first icon from right) in Figure 7-27 on page 68, a window is displayed as shown in Figure 7-29. In this example, the current MTU of vSwitch0 is 1500.




Figure 7-29 vSwitch0 properties

4. Click **Edit Settings** (the fifth icon from left) in Figure 7-27 on page 68. A window is displayed as shown in Figure 7-30.




Figure 7-30 Edit settings for vSwitch0

5. Enter **9000** in the **MTU (Bytes)** box, as shown in Figure 7-31, and click **OK** to continue.




Figure 7-31 Edit settings for vSwitch0

7.1.7 Add Host Networking

This section shows how to add more host networkings in ESXi host. The planned host networkings are shown in Table 7-3.

Table 7-3 Planned host networking

Name	VLAN ID	Usage
Public	11	For Oracle RAC Database services
InterConnect	3175	For Oracle RAC nodes internal communication

To add host networking, complete these steps:

1. From the vSphere Web Client, select **vm-host-oracle-01** in the inventory.
2. Click **Manage** → **Networking** → **Virtual Switches**, select **vSwitch0** from the list, as shown in Figure 7-27 on page 68.

3. Click **Add Host Networking** (the first icon from left) in Figure 7-27 on page 68. In the window that is displayed, select **Virtual Machine Port Group for a Standard Switch**, as shown in Figure 7-32, and click **Next** to continue.




Figure 7-32 Select connection type

4. Select an existing standard switch, as shown in Figure 7-33, and click **Next** to continue.




Figure 7-33 Select target device

5. Input the planned Network label and VLAN ID, as shown in Figure 7-34. Click **Next** to continue.




Figure 7-34 Enter planned network label

6. Review the settings before the host networking is created and click **Finish** to submit, as shown in Figure 7-35.




Figure 7-35 Review settings

- After all the host networkings are created, they will be reflected on virtual switch network diagram, as shown in Figure 7-36.




Figure 7-36 List host networkings

7.2 Create and modify virtual machines

This section shows how to create virtual machines on an ESXi host, and make the necessary modifications to build the environment for installing the Oracle RAC Database.

The target of this section is to create four virtual machines as planned in Table 7-4.

Table 7-4 Planned virtual machine configuration

VM Name	ESXi Host	CPU	Memory	Ethernet Adapters	Disks
ITSO_VM1	vm-host-oracle-01	56 cores	240 GB	1 Public and 1 InterConnect	1 local, 8 shared
ITSO_VM2	vm-host-oracle-02	56 cores	240 GB	1 Public and 1 InterConnect	1 local, 8 shared
ITSO_VM3	vm-host-oracle-03	56 cores	240 GB	1 Public and 1 InterConnect	1 local, 8 shared
ITSO_VM4	vm-host-oracle-04	56 cores	240 GB	1 Public and 1 InterConnect	1 local, 8 shared

7.2.1 Create virtual machines

Complete these steps to create a virtual machine on one ESXi host from vSphere Web Client:

1. From the vSphere Web Client, select **vm-host-oracle-01** in the inventory, then right-click **vm-host-oracle-01**, and select **New Virtual Machine** → **New Virtual Machine** as shown in Figure 7-37.




Figure 7-37 New virtual machine

2. Select the creation type as **Create a new virtual machine**, as shown in Figure 7-38, and click **Next** to continue.




Figure 7-38 Select a creation type

3. Specify the name and location of the virtual machine, as shown in Figure 7-39, and click **Next** to continue.




Figure 7-39 Select virtual machine name and folder

4. Select the compute resource of the virtual machine, in this example it is shown on one ESXi host as in Figure 7-40, and click **Next** to continue.




Figure 7-40 Select compute resource

5. Select a Datastore for the boot disk of the virtual machine. The boot disk is then stored in the **VM_OS** datastore, as shown in Figure 7-41, and click **Next** to continue.




Figure 7-41 Select storage

6. For the virtual machine compatibility, keep the default setting as shown in Figure 7-42, and click **Next** to continue.




Figure 7-42 Select compatibility

7. Select a Guest Operating System, Choose **Linux** in the **Guest OS Family** menu, and **Red Hat Enterprise Linux 7 (64-bit)** as the **Guest OS Version**, as shown in Figure 7-43. Click **Next** to continue.




Figure 7-43 Select guest OS

8. Customize virtual machine hardware, input **CPU** and **Memory** and choose the **Network Adapter** to connect to Public Networking, as shown in Figure 7-44.




Figure 7-44 Customize hardware

9. Click the small arrow in front of **New Hard Disk** to expand the disk options, as shown in Figure 7-45. Select **Thick provision eager zeroed** for the **Disk Provisioning** type, and click **Next** to continue.




Figure 7-45 Select disk options

10. Review the virtual machine configuration in the summary tab, as shown in Figure 7-46. If the configuration is correct, click **Finish** to submit virtual machine creation.




Figure 7-46 Review virtual machine configuration

Repeat these steps to create four virtual machines according to the planning for this environment.

7.2.2 Add network adapter

Complete these steps to add one additional network adapter for the virtual machines that were created in 7.2.1, “Create virtual machines” on page 75:

1. From the vSphere Web Client, select **ITSO_VM1** in the inventory, right-click **ITSO_VM1**, and select **Edit Settings**, as shown in Figure 7-47.




Figure 7-47 Edit virtual machine settings

2. Click the Virtual Hardware tab, click **New Device** to expand the drop list, and select **Network** in the list, as shown in Figure 7-48, then click **Add** button to continue.




Figure 7-48 Edit virtual machine hardware

3. Select **InterConnect** as the host networking of the newly added network adapter, and select **Connect**, as shown in Figure 7-49. Click the **OK** button to add the network adapter.




Figure 7-49 Add a network adapter

7.2.3 Create shared disks for virtual machines

The plan is to build an environment for running Oracle RAC Database. The shared VMware virtual disks shown in Table 7-5 are planned on all four virtual machines.

Table 7-5 Disk planning for virtual machines

Disk Name	Datastore	Size (GB)	SCSI ID	Usage	Shared
OS	VM_OS	200	0:0	Linux OS	No
OCR1	OCR_REDO	5	1:0	Oracle OCR and Voting	Yes
OCR2		5	1:1	Oracle OCR and Voting	Yes
OCR3		5	1:2	Oracle OCR and Voting	Yes
REDO		80	1:3	Oracle Redo log file	Yes

Disk Name	Datastore	Size (GB)	SCSI ID	Usage	Shared
DATA1	DATABASE	1024	1:4	Oracle data files	Yes
DATA2		1024	1:5	Oracle data files	Yes
DATA3		1024	1:6	Oracle data files	Yes
DATA4		1024	1:8	Oracle data files	Yes

To create shared disks for virtual machines, complete these steps:

1. From the vSphere Web Client, select **ITSO_VM1** in the inventory, right-click **ITSO_VM1**, and select **Edit Settings** from pop-up menu.
2. Click the Virtual Hardware tab, click **New Device** to expand the drop list, select **SCSI Controller** in the list, as shown in Figure 7-50, then click **Add** to continue.




Figure 7-50 Add a SCSI Controller

3. Make sure that the type of the new SCSI controller is **VMware Paravirtual**, as shown in Figure 7-51.




Figure 7-51 Virtual machine hardware configuration

4. Click **New Device** to expand the drop list and select **New Hard Disk** in the list, as shown in Figure 7-52. Click **Add** to continue.




Figure 7-52 Add a new hard disk

5. Click the arrow in front of **New Hard Disk** to expand the settings, input the size of new disk, and select the location of the disk, as shown in Figure 7-53.




Figure 7-53 Select new disk location

6. Select the datastore where to store this virtual disk as shown in Figure 7-54, and click **OK** to continue.




Figure 7-54 Select datastore

7. Go back to disk settings, and set the following options, as shown in Figure 7-55 on page 89:
 - a. Disk Provisioning: **Thick provision eager zeroed**
 - b. Sharing: **Multi-writer**
 - c. Virtual Device Node: **SCSI controller 1**, refer to the plan for SCSI ID
 - d. Disk Mode: **Independent - Persistent**




Figure 7-55 Hard disk settings

8. Repeat these steps to create all eight shared disks for virtual machine **ITSO_VM1**.
9. From the vSphere Web Client, select **ITSO_VM2** in the inventory, right-click **ITSO_VM2**, and select **Edit Settings** from pop-up menu.
10. Click the Virtual Hardware tab, click **New Device** to expand the drop list, select **SCSI Controller** in the list, and click **Add** to add one additional SCSI controller.

11. Click **New Device** to expand the drop list, click **Existing Hard Disk** in the list, as shown in Figure 7-56, and click **Add** to continue.




Figure 7-56 Add an existing hard disk

12. Select the datastore, and find the file of the shared virtual disk (usually it is a vmdk file), as shown in Figure 7-57. Click **OK** to continue.




Figure 7-57 Select disk location

13. Click the arrow in front of **New Hard Disk** to expand the settings, and set the following options, as shown in Figure 7-58 on page 92:

- Sharing: **Multi-writer**
- Virtual Device Node: **SCSI controller 1**, refer to the plan for SCSI ID
- Disk Mode: **Independent - Persistent**

Note: The Thick provision eager zeroed virtual disk is shown as lazy zeroed type. Ignore the incorrect disk type that is shown in this example.




Figure 7-58 Hard disk settings

7.2.4 Install VMware Tools

VMware Tools is a suite of utilities provided by VMware. The utilities run on a virtual machine's guest operating system, and improve management of the virtual machine. It is suggested to install VMware Tools on virtual machines in a VersaStack environment.

To check whether VMware Tools is installed, log in to the vSphere Web Client and click the virtual machine's name in the inventory, Figure 7-59 shows a virtual machine with VMware Tools installed and running.




Figure 7-59 VMware Tools running

Figure 7-60 shows a virtual machine that does not have VMware Tools installed yet.




Figure 7-60 VMware Tools not running

To install VMware Tools for a virtual machine running Red Hat Linux, complete these steps:

1. Ensure that the virtual machine is powered on.
2. From the vSphere Web Client, select virtual machine name in the inventory, right-click virtual machine name, and select **Guest OS** → **Install VMware Tools** from pop-up menu, as shown in Figure 7-61.




Figure 7-61 Install VMware Tools

3. Click **Mount** in the window that opens to mount the VMware Tools disk image to the virtual machine, as shown in Figure 7-62.




Figure 7-62 Mount disk image of VMware Tools

4. Log in to Linux running on the virtual machine as the root user.

5. Create a mount point and mount cdrom in Linux, as shown in Example 7-1.

Example 7-1 Create mount point and mount cdrom

```
# mkdir /mnt/dvd
# mount /dev/sr0 /mnt/dvd
mount: /dev/sr0 is write-protected, mounting read-only
```

6. Copy the VMware Tools package file to a temporary local directory, and extract the package, as shown in Example 7-2.

Example 7-2 Copy and extract VMware tools package

```
# cp /mnt/dvd/VMwareTools-9.10.5-2981885.tar.gz /tmp
# cd /tmp/
# tar -zxvf VMwareTools-9.10.5-2981885.tar.gz
```

7. Change to the new directory called `vmware-tools-distrib`, then start the VMware Tools installation, as shown in Example 7-3.

Example 7-3 Install VMware tools

```
# cd /tmp/vmware-tools-distrib/
# ./vmware-install.pl
```

7.2.5 Disable VMware Tools Time Synchronization

Oracle RAC Database uses Cluster Time Synchronization Service (CTSS) service to synchronize time between Oracle RAC nodes, so it is suggested to disable VMware Tools time synchronization on virtual machines.

From the vSphere Web Client, select **ITSO_VM1** in the inventory, right-click **ITSO_VM1** and select **Edit Settings** from pop-up menu. Click VM Options tab, click the arrow in front of VM Tools to expand settings, as shown in Figure 7-63. Make sure that **Synchronize guest time with host** is cleared.




Figure 7-63 VMware Tools settings

7.2.6 Enable disk UUID

The disk Universally Unique Identifier (UUID) is the unique identifier of one disk. By enabling disk UUID on virtual machines, VMware will always present a consistent UUID to the virtual machines. It is useful to identify the name of shared disks on each virtual machine in the cluster.

Complete these steps to enable disk UUID in virtual machines:

1. Shut down **ITSO_VM1** either from the host or from the vSphere Client.
2. From the vSphere Web Client, select **ITSO_VM1** in the inventory, right-click **ITSO_VM1**, and select **Edit Settings** from the pop-up menu.

3. Click the VM Options tab, click the arrow in front of **Advanced** to expand settings, as shown in Figure 7-64, and click **Edit Configuration**.




Figure 7-64 Virtual machine options

- The current virtual machine configuration parameters are listed. Click **Add Row** to start adding one more parameter, as shown in Figure 7-65.

Configuration Parameters

Modify or add configuration parameters as needed for experimental features or as instructed by technical support. Entries cannot be removed.

Name	Value
nvrAm	ITSO_VM4.nvrAm
pciBridge0.present	TRUE
svga.present	TRUE
pciBridge4.present	TRUE
pciBridge4.virtualDev	pcieRootPort
pciBridge4.functions	8
pciBridge5.present	TRUE
pciBridge5.virtualDev	pcieRootPort
pciBridge5.functions	8
pciBridge6.present	TRUE
pciBridge6.virtualDev	pcieRootPort
pciBridge6.functions	8

Add Row

OK Cancel

Figure 7-65 List virtual machine configuration parameters

- Input the parameter name as `disk.EnableUUID`, and value is set to `true`, as shown in Figure 7-66. Click **OK** to continue.

Configuration Parameters

Modify or add configuration parameters as needed for experimental features or as instructed by technical support. Entries cannot be removed.

Name	Value
replay.filename	
scsi0:0.redo	
scsi1:0.redo	
scsi1:1.redo	
scsi1:2.redo	
scsi1:3.redo	
scsi1:4.redo	
scsi1:5.redo	
scsi1:6.redo	
scsi1:8.redo	
disk.EnableUUID	true

Add Row

OK Cancel

Figure 7-66 Add a configuration parameter

6. Power on the virtual machine.

Repeat the above steps to enable disk UUID on the rest of virtual machines.

To get the UUID of a virtual disk, first locate the file of the virtual disk on ESXi host. It is usually under the directory `/vmfs/volumes/datastore_name/virtual_machine_name`. Issue the command as shown in Example 7-4 with the root user.

Example 7-4 Inquiry virtual disk UUID

```
[root@VM-Host-Oracle-01:~] vmkfstools -J getuuid
/vmfs/volumes/OCR_RED0/ITS0_VM1/ITS0_VM1.vmdk
UUID is 60 00 C2 9d 36 01 a6 44-2f 61 2e f6 39 16 e9 cb
```

To list the UUIDs of all disks in Red Hat Linux running on virtual machine, go to `/dev/disks/by-id/` directory to list the link files whose names start with `wwn`, as shown in Example 7-5.

Example 7-5 List all disks UUIDs in Red Hat Linux

```
# cd /dev/disk/by-id/
# ls -l wwn*
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c29042e2c8f9c2c477aa529cdd43 ->
.././sdf
lrwxrwxrwx. 1 root root 9 May 27 23:16 wwn-0x6000c293cde59615ca88fcd00b3c974f ->
.././sda
lrwxrwxrwx. 1 root root 10 May 27 23:16
wwn-0x6000c293cde59615ca88fcd00b3c974f-part1 -> .././sda1
lrwxrwxrwx. 1 root root 10 May 27 23:16
wwn-0x6000c293cde59615ca88fcd00b3c974f-part2 -> .././sda2
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c2944d3a235c5c249145baebe634 ->
.././sdc
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c2968068b713c270ca0941fc70aa ->
.././sdg
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c29be3eec07dad92dc236dee7adf ->
.././sdd
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c29c0f2bb25dc26c365f6ca822e6 ->
.././sdh
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c29d3601a6442f612ef63916e9cb ->
.././sdb
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c29d55112026025f75304fd1ad12 ->
.././sdi
lrwxrwxrwx. 1 root root 9 May 27 21:16 wwn-0x6000c29d6ee47e293b0924813b485552 ->
.././sde
```

Use the following command to list UUID for a single disk, as shown in Example 7-6.

Example 7-6 Inquiry UUID for one disk

```
# sg_inq --id /dev/sda
VPD INQUIRY: Device Identification page
  Designation descriptor number 1, descriptor length: 20
    designator_type: NAA, code_set: Binary
    associated with the addressed logical unit
      NAA 6, IEEE Company_id: 0xc29
```

Vendor Specific Identifier: 0x3cde59615
Vendor Specific Identifier Extension: 0xca88fcd00b3c974f
[0x6000c293cde59615ca88fcd00b3c974f]

Note: The `sg_inq` command requires `sg3_utils` RPM package.

7.3 Considerations of installing Red Hat Enterprise Linux

The installation of Red Hat Enterprise Linux on virtual machines is beyond the intended scope of this book. However, there are two suggestions for the users who install Linux in this environment:

- ▶ Base Environment Selection during installation
- ▶ Swap size

7.3.1 Base Environment Selection during installation

Oracle Grid Infrastructure and RAC database installation requires a graphic interface. It is suggested to install Red Hat Enterprise Linux Server 7.2 x86_64 version, with **Server with GUI** base environment group selected, as shown in Figure 7-67. The Oracle installation described later is based on this OS version.




Figure 7-67 Red Hat Environment Groups

7.3.2 Swap size

Red Hat Linux will configure swap automatically during installation. However, it is suggest that you configure the swap size following the rules in Table 7-6, according to the Oracle Database installation guide.

Table 7-6 Swap size suggestion

Memory Size	Swap Size
Between 2 GB and 16 GB	Equal to the size of the RAM
More than 16 GB	16 GB



Oracle RAC installation

This chapter describes detailed instructions about how to install Oracle RAC 12c Release 1 (12.1.0.2.0) Enterprise Edition on virtual machines running on the VersaStack environment. It covers following tasks:

- ▶ Preparing for installing Oracle Grid Infrastructure
- ▶ Installing Oracle Grid Infrastructure
- ▶ Installing Oracle RAC database
- ▶ Creating ASM Disk Groups
- ▶ Creating Oracle RAC database

For more information, see *Oracle Grid Infrastructure Installation Guide 12c Release 1 (12.1) for Linux*, which is found at the following website:

<https://docs.oracle.com/database/121/CWLIN/E48914-18.pdf>

Also, see *Oracle Database Installation Guide 12c Release 1 (12.1) for Linux*, which is found at the following website:

<https://docs.oracle.com/database/121/LADBI/title.htm>

8.1 Oracle Cluster network and storage planning

This section shows the Ethernet network and storage planning of a sample 4-node Oracle RAC cluster.

8.1.1 Network planning

Oracle RAC has the following IP addresses requirements. All these IP addresses should be registered with relevant names in the DNS server before Oracle RAC installation.

- ▶ A public IP address for each node
- ▶ A virtual IP address for each node
- ▶ A private IP address for each node
- ▶ Three Single Client Access Name (SCAN) addresses for Oracle RAC cluster

The node IP addresses and names in this environment are detailed in Table 8-1.

Table 8-1 Node IP addresses

	Node1	Node2	Node3	Node4
Public name	itsovm1	itsovm2	itsovm3	itsovm4
Public IP address	192.168.161.71	192.168.161.72	192.168.161.73	192.168.161.74
Virtual name	itsovm1-vip	itsovm2-vip	itsovm3-vip	itsovm4-vip
Virtual IP address	192.168.161.81	192.168.161.82	192.168.161.83	192.168.161.84
Private name	itsovm1-priv	itsovm2-priv	itsovm3-priv	itsovm4-priv
Private IP address	10.0.0.1	10.0.0.2	10.0.0.3	10.0.0.4

The SCAN IP addresses and name in this environment are detailed in Table 8-2.

Table 8-2 SCAN IP addresses

SCAN name	SCAN IP
itso-cluster-scan	192.168.161.75
	192.168.161.76
	192.168.161.77

8.1.2 Storage planning

According to disk planning for virtual machines in Table 7-5 on page 83, eight shared disks are mapped to each node. The ASM disk groups are detailed in Table 8-3.

Table 8-3 Storage planning

ASM Disk Group	Disk Size(GB)	Disk UUID
OCR	5	6000c29d3601a6442f612ef63916e9cb
	5	6000c2944d3a235c5c249145baebe634
	5	6000c29be3eec07dad92dc236dee7adf

ASM Disk Group	Disk Size(GB)	Disk UUID
REDO	80	6000c29042e2c8f9c2c477aa529cdd43
DATA	1024	6000c29d6ee47e293b0924813b485552
	1024	6000c2968068b713c270ca0941fc70aa
	1024	6000c29c0f2bb25dc26c365f6ca822e6
	1024	6000c29d55112026025f75304fd1ad12

8.2 Oracle RAC installation preparation

This section describes the necessary preparation tasks on the Linux operating system before Oracle RAC installation, and includes the following areas:

- ▶ Hardware configuration check
- ▶ Kernel and Linux operating system check
- ▶ Linux operating system configuration
- ▶ Oracle ASMLib installation and configuration

Note:

1. Unless specified, run all the actions described in this section on each node of the Oracle RAC cluster.
2. Unless specified, run all the commands in this section using the root user.

8.2.1 Check hardware configuration

To install Oracle RAC 12c on Linux x86_64 platform, the following hardware requirements should be met:

- ▶ At least 4 GB of RAM for Oracle Grid Infrastructure and Oracle Database.
- ▶ At least one 1 GbE Ethernet interface card for public network communication (suggest having two cards bonded as one for redundancy), 10 GbE is preferred.
- ▶ At least one 1 GbE Ethernet interface card for private network communication (suggest having two cards bonded as one for redundancy), 10 GbE is preferred.
- ▶ Server should be started in runlevel 3 or runlevel 5.
- ▶ Server display cards provide at least 1024 x 768 display resolution.
- ▶ Server is connected to the network, contains a display monitor and DVD drive.
- ▶ At least 6.1 GB of disk space for an Oracle Database.
- ▶ At least 6.9 GB of disk space for an Oracle Grid Infrastructure.
- ▶ At least 1 GB allocated to /tmp.

To check hardware configuration on each node, complete these steps:

1. Use **grep MemTotal /proc/meminfo** command to check the system memory size, Example 8-1 shows that the total memory size of this virtual machine is 247458036 kB.

Example 8-1 Check memory size

```
# grep MemTotal /proc/meminfo
MemTotal:      247458036 kB
```

2. Use **ip link** or **ifconfig -a** command to check existing Ethernet interfaces, Example 8-2 shows that there are two Ethernet interfaces installed in this virtual machine.

Example 8-2 Check Ethernet interfaces

```
# ip link
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode
DEFAULT
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eno16782080: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
mode DEFAULT qlen 1000
    link/ether 00:50:56:b7:8b:75 brd ff:ff:ff:ff:ff:ff
3: eno33561344: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
mode DEFAULT qlen 1000
    link/ether 00:50:56:b7:de:f0 brd ff:ff:ff:ff:ff:ff
```

3. To check the current running speed of an Ethernet interface, use the **ethtool** command followed by the interface name, Example 8-3 shows this Ethernet is running at 10 Gbps (10000 Mb/s).

Example 8-3 Show Ethernet speed

```
# ethtool eno16782080
Settings for eno16782080:
  Supported ports: [ TP ]
  Supported link modes:   1000baseT/Full
                        10000baseT/Full

  Supported pause frame use: No
  Supports auto-negotiation: No
  Advertised link modes:  Not reported
  Advertised pause frame use: No
  Advertised auto-negotiation: No
  Speed: 10000Mb/s
  Duplex: Full
  Port: Twisted Pair
  PHYAD: 0
  Transceiver: internal
  Auto-negotiation: off
  MDI-X: Unknown
  Supports Wake-on: uag
  Wake-on: d
  Link detected: yes
```

4. Use `runlevel` command to check the current run level of operating system, Example 8-4 shows operating system is running at runlevel 5.

Example 8-4 Check Linux run level

```
# runlevel
N 5
```

Note: Runlevel 5 means Linux host boots into a graphical environment.

5. To check usable file system size, use `df -h` command in Linux.

8.2.2 Check Linux kernel and OS version

For Oracle running on Red Hat Enterprise Linux 7.x x86_64 platform, it requires Linux Kernel version 3.10.0-54.0.1.el7.x86_64 or later. Use the `uname -a` command to check current kernel version on Linux, as shown in Example 8-5.

Example 8-5 Check Linux kernel version

```
# uname -a
Linux localhost.localdomain 3.10.0-327.el7.x86_64 #1 SMP Thu Oct 29 17:29:29 EDT
2015 x86_64 x86_64 x86_64 GNU/Linux
```

Another alternative is to check Red Hat OS version. The kernel version of each Red Hat Linux release is published on the following Red Hat website:

<https://access.redhat.com/articles/3078>

To check the current OS version, use either of the following commands that are shown in Example 8-6. The OS version that was used in this environment is Red Hat Enterprise Linux Server 7.2.

Example 8-6 Check OS version

```
# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 7.2 (Maipo)

# lsb_release -a
LSB Version:::core-4.1-amd64:core-4.1-noarch
Distributor ID:RedHatEnterpriseServer
Description:Red Hat Enterprise Linux Server release 7.2 (Maipo)
Release:7.2
Codename:Maipo
```

Note: The `lsb_release` command requires the `redhat-lsb-core` RPM package to be installed.

8.2.3 Setup hostname

A host uses the default hostname `localhost` if it is not set during OS installation. A host deployed from VMware virtual machine template always has the same hostname with its template. Oracle RAC database requires that each node in the cluster has a unique name.

To set a hostname for a host, use the `hostnamectl set-hostname` command followed by the wanted hostname as shown in Example 8-7.

Example 8-7 Setup hostname

```
# hostnamectl set-hostname itsovm1
```

To verify the configured hostname, use the `hostnamectl` command as shown in Example 8-8.

Example 8-8 Listing hostname information

```
# hostnamectl
  Static hostname: itsovm1
            Icon name: computer-vm
            Chassis: vm
            Machine ID: d930cf0ba73c4d15b74d71f7d66ab66d
            Boot ID: 20bef504e6d242c08141d073aed3445a
  Virtualization: vmware
  Operating System: Red Hat Enterprise Linux Server 7.2 (Maipo)
            CPE OS Name: cpe:/o:redhat:enterprise_linux:7.2:GA:server
            Kernel: Linux 3.10.0-327.el7.x86_64
  Architecture: x86-64
```

After the hostname is set up, the next action is to update `/etc/hosts` file with a line containing the current IP address and hostname. A sample of the `/etc/hosts` file is shown in Example 8-9.

Example 8-9 A /etc/hosts sample

```
# more /etc/hosts
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
::1      localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.161.71 itsovm1
```

Note: Do not append hostname to the loopback address in the `/etc/hosts` file.

8.2.4 Configure IP address

IP address on the Ethernet interface cannot be configured during Red Hat OS installation, and to verify whether it is configured or not, use `ifconfig -a` command as shown in Example 8-10.

Example 8-10 Check IP addresses on Ethernet interfaces

```
# ifconfig -a
eno16782080: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.161.71 netmask 255.255.252.0 broadcast 192.168.163.255
    inet6 fe80::250:56ff:feb7:8b75 prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:b7:8b:75 txqueuelen 1000 (Ethernet)
    RX packets 3618801 bytes 5332420791 (4.9 GiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 171808 bytes 13282007 (12.6 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eno33561344: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    ether 00:50:56:b7:de:f0 txqueuelen 1000 (Ethernet)
```

```

RX packets 3 bytes 180 (180.0 B)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 0 bytes 0 (0.0 B)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

```

Example 8-10 shows Ethernet interface eno16782080 is configured with IP address 192.168.161.71, while Ethernet interface eno33561344 is not. Because the two Ethernet interfaces in this node are connected to a public and private network respectively, it is important to confirm that the correct IP address is configured on the correct Ethernet interface so that network communication is operational. To do so, complete the following steps:

1. Retrieve MAC addresses of public and private interfaces from VMware vSphere.

From the vSphere Web Client, select **ITSO_VM1** in the inventory, navigate to **VM Hardware** in **Summary** in right pane, and click **Network Adapter 1** and **Network Adapter 2** to show more information, as shown in Figure 8-1.




Figure 8-1 Check MAC Addresses of Ethernet interfaces

2. Determine the device names of public and private interfaces by comparing to MAC addresses from VMware and Linux.

By comparing the MAC addresses from vSphere Web Client with the output of command **ifconfig -a**, you can tell the device name of public Ethernet interface is eno16782080, and private Ethernet interface is eno33561344.

3. Verify or configure IP addresses for public and private interfaces.

According to 8.1.1, “Network planning” on page 104, the IP addresses of Ethernet adapters for this node is shown in Table 8-4.

Table 8-4 IP addresses for interfaces

Network	Device name	IP address
Public	eno16782080	192.168.161.71
Private	eno33561344	10.0.0.1

To verify whether IP address is configured on Ethernet interface or not, use the command **ifconfig -a**. If it is not configured, use a tool to modify the configuration file that is located in the `/etc/sysconfig/network-scripts` directory. The configuration file name should start with **ifcfg-**, and be followed by the device name of Ethernet interface. For example, **ifcfg-eno33561344** is the configuration file for Ethernet interface eno33561344.

Usually you must change the following parameters in the Ethernet interface configuration file.

BOOTPROTO	Change from dhcp to none
ONBOOT	Change from no to yes
NETMASK or PREFIX	Specify network mask
IPADDR	Specify IP address of interface

Example 8-11 shows a completed Ethernet interface configuration file.

Example 8-11 Ethernet interface configuration file

```
# cat /etc/sysconfig/network-scripts/ifcfg-eno33561344
TYPE=Ethernet
BOOTPROTO=none
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
NAME=eno33561344
UUID=493bbcbe-5c6d-4257-b627-cb817ff9bc13
DEVICE=eno33561344
ONBOOT=yes
IPADDR=10.0.0.1
PREFIX=24
IPV6_PEERDNS=yes
IPV6_PEERROUTES=yes
```

To activate network configuration with latest configuration file, run the **systemctl restart network.service** command.

Note: Clients can also use Red Hat graphic tool to configure IP address for Ethernet interface. To start this tool, run the **gnome-control-center** command in the Linux GUI, and click the **Network** icon.

8.2.5 Check swap size

Usually the swap size in Linux is related to the system memory size. Table 8-5 describes the suggested relationship between the system memory size and swap size.

Table 8-5 Swap size requirement

Memory Size	Swap Size
Between 2 GB and 16 GB	Equal to the size of the RAM
More than 16 GB	16 GB

The system memory size of the node in this environment is 240 GB. According to the above rule, a 16 GB swap size is required, Example 8-12 shows how to check swap size in Linux.

Example 8-12 Check swap size in Linux

```
# grep SwapTotal /proc/meminfo
SwapTotal:      16883708 kB
```

8.2.6 Enable jumbo frames support on private Ethernet interface

Jumbo frames are Ethernet frames with more than 1,500 bytes of payload (MTU). Typical jumbo frames can carry up to 9,000 bytes of payload. Usually enabling jumbo frames can improve network performance as well as reduce server usage and CPU cycles while transferring large files.

Jumbo Frames have these requirements on hardware and configuration:

1. Ethernet interfaces and switches are running in 10 Gb mode.
2. All network components from end to end support jumbo frames configuration.

Generally, enable jumbo frames support on the private Ethernet interface for Oracle RAC node. However, enabling jumbo frames support on the public Ethernet interface is optional, depending on customer network environment and server configuration.

Ethernet interfaces and switches in the VersaStack environment are all capable of running at 10 Gb mode, and it is possible to configure jumbo frames support on Ethernet switches.

Form information about configuring jumbo frames support on Ethernet interfaces in VMware ESXi, see 7.1.6, “Enable Jumbo Frames” on page 69.

To configure jumbo frames support on private Ethernet interface on Oracle RAC node, navigate to the `/etc/sysconfig/network-scripts` directory, open the private interface configuration file, and append one row that describes MTU, as shown in Example 8-13.

Example 8-13 Enable jumbo frames

```
# echo 'MTU=9000' >> /etc/sysconfig/network-scripts/ifcfg-eno33561344
```

After you restart the Linux network service, the private Ethernet interface should be able to support jumbo frames. Example 8-14 shows how to use the `ping` command to test jumbo frames in Linux.

Example 8-14 Jumbo Frames testing

```
# ping -s 8192 10.0.0.2
PING 10.0.0.2 (10.0.0.2) 8192(8220) bytes of data.
8200 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=0.609 ms
8200 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=0.326 ms
8200 bytes from 10.0.0.2: icmp_seq=3 ttl=64 time=0.346 ms
8200 bytes from 10.0.0.2: icmp_seq=4 ttl=64 time=0.348 ms
8200 bytes from 10.0.0.2: icmp_seq=5 ttl=64 time=0.308 ms
```

8.2.7 Configure DNS Server

Use DNS to resolve names that are used in Oracle RAC environment, especially SCAN. To configure DNS Servers for a node, modify the `/etc/resolv.conf` file to add the correct search and nameserver parameters. Example 8-15 shows a typical DNS configuration file on the Linux host.

Example 8-15 DNS configuration file

```
# cat /etc/resolv.conf
search versastack.local
nameserver 192.168.161.50
nameserver 192.168.161.51
```

According to our network planning, Oracle SCAN in DNS is required to be configured as one single name entry with three A records. After DNS is set up, make sure that all three A records are returned in `nslookup`. Example 8-16 shows how to test whether Oracle SCAN name resolution is correctly configured in DNS server.

Example 8-16 DNS lookup test

```
# nslookup
> itso-cluster-scan
Server: 192.168.161.50
Address: 192.168.161.50#53

Name: itso-cluster-scan.versastack.local
Address: 192.168.161.77
Name: itso-cluster-scan.versastack.local
Address: 192.168.161.76
Name: itso-cluster-scan.versastack.local
Address: 192.168.161.75
```

8.2.8 Synchronize Time with Time Server

Oracle RAC Cluster requires that the echo node in the cluster is configured with the same time zone, and the clock on each node is synced. There are two supported options for time synchronization:

- ▶ An operating system configured Network Time Protocol (NTP)
- ▶ Oracle Cluster Time Synchronization Service (CTSS)

To check current time and time zone information, run the `timedatectl` command as shown in Example 8-17.

Example 8-17 Check time and time zone

```
# timedatectl
   Local time: Mon 2016-05-30 13:18:58 EDT
   Universal time: Mon 2016-05-30 17:18:58 UTC
     RTC time: Mon 2016-05-30 19:19:02
   Time zone: America/New_York (EDT, -0400)
   NTP enabled: yes
NTP synchronized: no
   RTC in local TZ: no
     DST active: yes
Last DST change: DST began at
                  Sun 2016-03-13 01:59:59 EST
                  Sun 2016-03-13 03:00:00 EDT
Next DST change: DST ends (the clock jumps one hour backwards) at
                  Sun 2016-11-06 01:59:59 EDT
                  Sun 2016-11-06 01:00:00 EST
```

If the time zone is not correctly configured, run the `timedatectl set-timezone` command followed by the wanted time zone city, which can be gotten from the `timedatectl list-timezones` command as shown in Example 8-18.

Example 8-18 List timezones

```
# timedatectl list-timezones
Africa/Abidjan
Africa/Accra
Africa/Addis_Ababa
Africa/Algiers
...<truncated output>...
```

Before starting Oracle RAC installation, ensure that the clocks on all nodes are set to the same time. Therefore, a manual time synchronization with the time server is required. To start the synchronization process, run the `ntpdate` command followed by a legal time server provided by system administrator, as shown in Example 8-19.

Example 8-19 Synchronize time by using ntpdate

```
# ntpdate -v 192.168.160.254
26 May 11:34:37 ntpdate[26430]: ntpdate 4.2.6p5@1.2349-o Fri Oct 16 08:51:51 UTC
2015 (1)
26 May 11:34:56 ntpdate[26430]: step time server 192.168.160.254 offset 12.412580
sec
```

This environment uses the Oracle CTSS service to synchronize time between Oracle RAC nodes. In this case, disable the NTP service. First, check whether the `ntp` package is installed, as shown in Example 8-20.

Example 8-20 Query ntp package

```
# rpm -qa|grep ntp
ntpdate-4.2.6p5-22.e17.x86_64
fontpackages-filesystem-1.44-8.e17.noarch
```

```
python-ntplib-0.3.2-1.e17.noarch
ntp-4.2.6p5-22.e17.x86_64
```

If ntp package is not listed in the above output, no action is required. Otherwise, it is suggested to stop and disable ntp service, and remove the ntpd.pid file, as shown in Example 8-21.

Example 8-21 Stop and disable ntpd service

```
# systemctl status ntpd.service
# systemctl stop ntpd.service
# systemctl disable ntpd.service
# rm /var/run/ntpd.pid
```

8.2.9 Install required RPM packages

To install Oracle RAC Database 12c on Red Hat Enterprise Linux 7 (x86_64), the following RPM packages are required:

- ▶ compat-libcap1
- ▶ compat-libstdc++-33
- ▶ gcc
- ▶ gcc-c++
- ▶ glibc
- ▶ glibc-devel
- ▶ ksh
- ▶ libaio
- ▶ libaio-devel
- ▶ libgcc
- ▶ libstdc++
- ▶ libstdc++-devel
- ▶ libXi
- ▶ libXtst
- ▶ make
- ▶ sysstat

Tips:

- ▶ It is suggested to configure yum to install these required RPM packages because some of the packages that are listed have dependent packages.
- ▶ compat-libstdc++-33 package is in rhel-7-server-optional-rpms repository, which is not enabled by default.

To check whether these packages are installed, save this package list to a file named pkg.lst in Linux, and run the command shown in Example 8-22. A prompt of xxx is not installed is shown in the output if a required package is missing.

Example 8-22 Check required rpm packages

```
# for i in `cat pkg.lst`;do rpm -q $i;done
compat-libcap1-1.10-7.e17.x86_64
compat-libstdc++-33-3.2.3-72.e17.x86_64
gcc-4.8.5-4.e17.x86_64
gcc-c++-4.8.5-4.e17.x86_64
glibc-2.17-105.e17.x86_64
```

```
glibc-devel-2.17-105.e17.x86_64
ksh-20120801-22.e17_1.2.x86_64
libaio-0.3.109-13.e17.x86_64
libaio-devel-0.3.109-13.e17.x86_64
libgcc-4.8.5-4.e17.x86_64
libstdc++-4.8.5-4.e17.x86_64
libstdc++-devel-4.8.5-4.e17.x86_64
libXi-1.7.4-2.e17.x86_64
libXtst-1.2.2-2.1.e17.x86_64
make-3.82-21.e17.x86_64
sysstat-10.1.5-7.e17.x86_64
```

8.2.10 Create users and groups

For installing and administering Oracle Grid Infrastructure and Oracle RAC Database software, it is suggested to create different users and groups in Linux. The Oracle Grid Infrastructure software owner is the user `grid` and the Oracle Database software owner is the user `oracle` in this environment.

Example 8-23 shows the commands to create users, groups, and directories required for Oracle installation.

Example 8-23 Commands to create users, groups, and directories for Oracle installation

```
# /usr/sbin/groupadd -g 54321 oinstall
# /usr/sbin/groupadd -g 54322 dba
# /usr/sbin/groupadd -g 54323 oper
# /usr/sbin/groupadd -g 54324 backupdba
# /usr/sbin/groupadd -g 54325 dgdba
# /usr/sbin/groupadd -g 54326 kmdba
# /usr/sbin/groupadd -g 54327 asmdba
# /usr/sbin/groupadd -g 54328 asmoper
# /usr/sbin/groupadd -g 54329 asmadmin
# useradd -u 54421 -g oinstall -G dba,oper,backupdba,dgdba,kmdba,asmdba,asmoper
oracle
# useradd -u 54422 -g oinstall -G dba,asmadmin,asmdba,asmoper grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/12.1.0/grid
# mkdir -p /u01/app/oraInventory
# chown -R oracle:oinstall /u01
# chown oracle:oinstall /u01/app/oracle
# chown grid:oinstall /u01/app/grid
# chown grid:oinstall /u01/app/oraInventory
# chmod -R 775 /u01/
```

After the oracle user and grid user are created, use the `id` command to verify the group information of each user, as shown in Example 8-24.

Example 8-24 Verify user group information

```
# id grid
uid=54422(grid) gid=54321(oinstall)
groups=54321(oinstall),54322(dba),54327(asmdba),54328(asmoper),54329(asmadmin)
```

```
# id oracle
uid=54421(oracle) gid=54321(oinstall)
groups=54321(oinstall),54322(dba),54323(oper),54324(backupdba),54325(dgdba),54326(
kndba),54327(asmdba),54328(asmoper)
```

To allow a successful login of the oracle and grid user later, the password of each user has to be set, as shown in Example 8-25.

Example 8-25 Set up password for oracle and grid user

```
# passwd oracle
Changing password for user oracle.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

```
# passwd grid
Changing password for user grid.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

8.2.11 Configure firewall

Starting from Red Hat Enterprise Linux 7, a new dynamic firewall daemon `firewalld` is introduced. This daemon provides a dynamically managed firewall with support for network “zones” to assign a level of trust to a network and its associated connections and interfaces. For more information about `firewalld`, see Section 4.5 in the *Red Hat Enterprise Linux 7 Security Guide*, which is available at:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Security_Guide/sec-Using_Firewalls.html

A zone in `firewalld` defines the level of trust for network connections. There are nine zones predefined in `firewalld` in total. Only these zones are going to be modified in this environment:

- ▶ **public zone:** Default zone, for use in public areas. All connections from other computers are denied except for selected incoming connections.
- ▶ **trusted zone:** All network connections are accepted.

Example 8-26 shows how to list the current public and trusted zone configuration.

Example 8-26 List runtime public and trusted zone configuration

```
# firewall-cmd --zone=public --list-all
public (default, active)
  interfaces: eno16782080 eno33561344
  sources:
  services: dhcpv6-client ssh
  ports:
  masquerade: no
  forward-ports:
  icmp-blocks:
  rich rules:
```

```
# firewall-cmd --zone=trusted --list-all
trusted
  interfaces:
  sources:
  services:
  ports:
  masquerade: no
  forward-ports:
  icmp-blocks:
  rich rules:
```

To make sure that Oracle RAC Database Cluster is able to provide services for external clients, network connections from the following TCP ports should be accepted:

- ▶ Port 1521 - Oracle Transparent Network Substrate (TNS) Listener port
- ▶ Port 5500 - Oracle Enterprise Manager port

Example 8-27 shows how to add the port to the public zone.

Example 8-27 Add port to public zone

```
# firewall-cmd --zone=public --add-port=1521/tcp
success
# firewall-cmd --zone=public --add-port=5500/tcp
success
```

It is suggested to allow all network traffic on the private interconnect Ethernet interface. For Oracle RAC installation, it is better to trust service IP addresses of Oracle RAC nodes. Example 8-28 shows how to configure the trusted zone.

Example 8-28 Configure trusted zone

```
# firewall-cmd --zone=trusted --change-interface=en033561344
success
# firewall-cmd --zone=trusted --add-source=192.168.160.0/22
success
```

After an Ethernet interface is added to the trusted zone in firewall, the interface configuration file also needs to be updated to make sure that the configuration is persistent, as shown in Example 8-29.

Example 8-29 Modify interface configuration file

```
# echo 'ZONE=trusted' >> /etc/sysconfig/network-scripts/ifcfg-en033561344
```

Any changes that are made to `firewalld` need to be saved by using the command shown in Example 8-30 in case the configuration is lost after the firewall daemon reload.

Example 8-30 Save firewall configuration

```
# firewall-cmd --runtime-to-permanent
success
```

The final permanent firewall configuration is shown in Example 8-31.

Example 8-31 Show permanent firewall configuration

```
# firewall-cmd --permanent --zone=public --list-all
public (default, active)
  interfaces: eno16782080
  sources:
  services: dhcpv6-client ssh
  ports: 1521/tcp 5500/tcp
  masquerade: no
  forward-ports:
  icmp-blocks:
  rich rules:

# firewall-cmd --permanent --zone=trusted --list-all
trusted (active)
  interfaces: eno33561344
  sources: 192.168.160.0/22
  services:
  ports:
  masquerade: no
  forward-ports:
  icmp-blocks:
  rich rules:
```

Note: Linux Firewall can be also configured by using the graphical user interface tool **firewall-config**.

8.2.12 Configure SELinux

Security-Enhanced Linux (SELinux) is an implementation of a flexible mandatory access control architecture in the Linux operating system. It is suggested to run SELinux in enforcing mode. To check the status of SELinux, run the **sestatus** command as shown in Example 8-32.

Example 8-32 Check SELinux status

```
# sestatus
SELinux status:                enabled
SELinuxfs mount:              /sys/fs/selinux
SELinux root directory:       /etc/selinux
Loaded policy name:           targeted
Current mode:                 enforcing
Mode from config file:       enforcing
Policy MLS status:           enabled
Policy deny_unknown status:   allowed
Max kernel policy version:    28
```

There is a known issue that the `oracleasm` script might fail if SELinux is in enforcing mode. For more detailed information and instructions on how to work around the issue, see the following guide:

https://docs.oracle.com/cd/E52668_01/E53499/E53499.pdf

It is suggested to disable the SELinux policy module for oracleasm as shown in Example 8-33.

Example 8-33 Disable SELinux policy module for oracleasm

```
# semodule -d oracleasm
# semodule -l|grep oracleasm
oracleasm 1.0.0 Disabled
```

8.2.13 Set Kernel parameters

Many kernel parameters are required to be configured for Oracle installation. The descriptions of each kernel parameter and its suggested value are shown in Table 8-6.

Table 8-6 Kernel parameters

Kernel Parameter	Description	Suggested Value
kernel.sem	The kernel parameter sem is composed of four parameters: SEMMSL, SEMMNI, SEMMNS, SEMOPM.	250 32000 100 128
kernel.shmmax	This kernel parameter defines the maximum size in bytes of a single shared memory segment that a Linux process can allocate in its virtual address space.	Half of memory in bytes
kernel.shmmni	SHMMNI is the maximum number of shared memory segments.	4096
kernel.shmall	This parameter sets the total number of shared memory pages that can be used.	Memory size * 0.4 / Page size
fs.file-max	This parameter defines the maximum number of open file handles.	6815744
net.ipv4.ip_local_port_range	This parameter defines the range that local ports are allowed in for TCP and UDP traffic.	9000 65535
net.core.rmem_default	This parameter defines the default setting in bytes of the socket receive buffer.	262144
net.core.rmem_max	This parameter defines the maximum socket receive buffer size in bytes.	4194304
net.core.wmem_default	This parameter defines the default setting in bytes of the socket send buffer.	262144
net.core.wmem_max	This parameter defines the maximum socket send buffer size in bytes.	1048576
fs.aio-max-nr	This parameter sets the maximum number of concurrent asynchronous I/O requests.	1048576

To modify these kernel parameters, edit `/etc/sysctl.conf` to add the lines that are shown in Example 8-34.

Example 8-34 /etc/sysctl.conf

```
kernel.sem = 250 32000 100 128
kernel.shmmax = 126698514432
```

```

kernel.shmmni = 4096
kernel.shmall = 24745803
fs.file-max = 6815744
net.ipv4.ip_local_port_range = 9000 65535
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576
fs.aio-max-nr = 1048576

```

In order for the changes to take effect immediately, run the command shown in Example 8-35.

Example 8-35 Refresh with the new configuration

```
# sysctl -p /etc/sysctl.conf
```

8.2.14 Set user limits

For grid owner user and database owner user, Oracle requires that you set appropriate system resource limits, and the suggested ranges are shown in Table 8-7.

Table 8-7 Resource limit

Resource	Description	Soft Limit	Hard Limit
nofile	Maximum number of open file descriptors	>= 1024	>= 65536
nproc	Maximum number of processes	>= 2047	>= 16384
stack	Maximum stack size (KB)	> 10240KB	<= 32768KB
memlock	Maximum locked-in-memory address space (KB)	> = 90% Memory with HugePages	> = 90% Memory with HugePages

To check the current user soft and hard limits, log in to Linux with the username that you want, and run the command `ulimit -Sa` and `ulimit -Ha`.

To update the limits of oracle and grid user in this environment, open the `/etc/security/limits.conf` file and add content as shown in Example 8-36. The maximum number of processes limit (nproc) is not added to this file because the default limits setting meets the requirement.

Example 8-36 Oracle and grid user limit

```

oracle      soft  nofile      4096
oracle      hard  nofile      65536
oracle      soft  stack       10240
oracle      hard  stack       32768
oracle      soft  memlock     222712233
oracle      hard  memlock     222712233
grid        soft  nofile      4096
grid        hard  nofile      65536
grid        soft  stack       10240
grid        hard  stack       32768
grid        soft  memlock     222712233
grid        hard  memlock     222712233

```

The limits setting should be effective immediately after the user logs in again. Double-check the limits before Oracle installation. Example 8-37 shows the updated limits of the oracle user.

Example 8-37 Oracle user limits

```
[oracle@itsovm1 ~]$ ulimit -Sa
core file size          (blocks, -c) 0
data seg size          (kbytes, -d) unlimited
scheduling priority    (-e) 0
file size              (blocks, -f) unlimited
pending signals        (-i) 966511
max locked memory      (kbytes, -l) 222712233
max memory size        (kbytes, -m) unlimited
open files             (-n) 4096
pipe size              (512 bytes, -p) 8
POSIX message queues   (bytes, -q) 819200
real-time priority     (-r) 0
stack size             (kbytes, -s) 10240
cpu time               (seconds, -t) unlimited
max user processes     (-u) 4096
virtual memory         (kbytes, -v) unlimited
file locks             (-x) unlimited

[oracle@itsovm4 ~]$ ulimit -Ha
core file size          (blocks, -c) unlimited
data seg size          (kbytes, -d) unlimited
scheduling priority    (-e) 0
file size              (blocks, -f) unlimited
pending signals        (-i) 966511
max locked memory      (kbytes, -l) 222712233
max memory size        (kbytes, -m) unlimited
open files             (-n) 65536
pipe size              (512 bytes, -p) 8
POSIX message queues   (bytes, -q) 819200
real-time priority     (-r) 0
stack size             (kbytes, -s) 32768
cpu time               (seconds, -t) unlimited
max user processes     (-u) 966511
virtual memory         (kbytes, -v) unlimited
file locks             (-x) unlimited
```

8.2.15 Disable avahi-daemon service and Zero Configuration Networking

Avahi is a daemon that runs on Linux by default. It implements Apple's Zeroconf architecture (also known as "Rendezvous" or "Bonjour"). It might cause ora.cssd fails to join the cluster in certain situations. Therefore, Oracle suggests that you disable the avahi-daemon service. For more information, see the following Oracle document:

<https://support.oracle.com/epmos/faces/DocumentDisplay?parent=DOCUMENT&id=1501093.1>

To check current avahi-daemon status, run the `systemctl status avahi-daemon` command as shown in Example 8-38.

Example 8-38 Check avahi-daemon status

```
# systemctl status avahi-daemon
avahi-daemon.service - Avahi mDNS/DNS-SD Stack
  Loaded: loaded (/usr/lib/systemd/system/avahi-daemon.service; enabled; vendor
  preset: enabled)
  Active: active (running) since Fri 2016-05-27 11:33:38 EDT; 1h 2min ago
  Main PID: 1513 (avahi-daemon)
  Status: "avahi-daemon 0.6.31 starting up."
  CGroup: /system.slice/avahi-daemon.service
          ..1513 avahi-daemon: running [linux.local]
          ..1521 avahi-daemon: chroot helper
```

Example 8-39 shows how to disable the avahi-daemon persistently.

Example 8-39 Disable avahi-daemon service

```
# systemctl stop avahi-daemon
Warning: Stopping avahi-daemon.service, but it can still be activated by:
  avahi-daemon.socket

# systemctl disable avahi-daemon
Removed symlink /etc/systemd/system/multi-user.target.wants/avahi-daemon.service.
Removed symlink /etc/systemd/system/sockets.target.wants/avahi-daemon.socket.
Removed symlink /etc/systemd/system/dbus-org.freedesktop.Avahi.service.
```

The Zero Configuration Networking feature need to be disabled also, as it can cause communication issues between cluster member nodes. Update the `/etc/sysconfig/network` file with one line as shown in Example 8-40.

Example 8-40 Disable Zero Configuration Networking

```
# Created by anaconda
NOZEROCONF=yes
```

8.2.16 Disable RemoveIPC feature

By default in Red Hat Enterprise Linux 7.2, `systemd-logind` service removes all IPC objects owned by a user if this user fully logs out. This feature removes shared memory segments and semaphores for oracle and grid users if they are not connected any more, which crashes the Oracle ASM and database instances.

To disable this feature, update the `/etc/systemd/logind.conf` file as shown in Example 8-41.

Example 8-41 RemoveIPC configuration

```
# This file is part of systemd.
#
# systemd is free software; you can redistribute it and/or modify it
# under the terms of the GNU Lesser General Public License as published by
# the Free Software Foundation; either version 2.1 of the License, or
# (at your option) any later version.
#
```

```
# Entries in this file show the compile time defaults.
# You can change settings by editing this file.
# Defaults can be restored by simply deleting this file.
#
# See logind.conf(5) for details.
```

```
[Login]
#NAutoVTs=6
#ReserveVT=6
#KillUserProcesses=no
#KillOnlyUsers=
#KillExcludeUsers=root
#InhibitDelayMaxSec=5
#HandlePowerKey=poweroff
#HandleSuspendKey=suspend
#HandleHibernateKey=hibernate
#HandleLidSwitch=suspend
#HandleLidSwitchDocked=ignore
#PowerKeyIgnoreInhibited=no
#SuspendKeyIgnoreInhibited=no
#HibernateKeyIgnoreInhibited=no
#LidSwitchIgnoreInhibited=yes
#IdleAction=ignore
#IdleActionSec=30min
#RuntimeDirectorySize=10%
#RemoveIPC=yes
RemoveIPC=no
```

For more information about RemoveIPC, see the Oracle support website at:

<https://support.oracle.com/epmos/faces/DocumentDisplay?id=2081410.1>

It is also mentioned on Red Hat website at:

https://bugzilla.redhat.com/show_bug.cgi?id=1264533

8.2.17 Shared memory file system /dev/shm

Mounting tmpfs at shared memory file system /dev/shm is handled automatically by Red Hat Linux. This temporary file system size is always set to be half of the installed memory, so manual configuration in /etc/fstab is no longer necessary. However, Oracle installation program checks /etc/fstab for entry of mounting /dev/shm. It is suggested to add one entry in /etc/fstab as shown in Example 8-42.

Example 8-42 tmpfs in /etc/fstab

```
# for tmpfs
tmpfs                /dev/shm            tmpfs   rw,exec            0 0
```

8.2.18 Disable Transparent HugePages

Transparent HugePages (THP) allocates memory dynamically during run time. According to Oracle, it can cause memory allocation delays. To avoid performance issues of Oracle Database, Oracle suggests disabling THP on all Oracle Database servers.

To check whether THP is enabled, run the command shown in Example 8-43.

Example 8-43 Check Transparent HugePages status

```
# cat /sys/kernel/mm/transparent_hugepage/enabled
[always] madvise never
```

From the output in Example 8-43, [always] means that THP is enabled. To disable THP, update the /etc/default/grub file, and insert transparent_hugepage=never to the end of line starting with GRUB_CMDLINE_LINUX, as shown in Example 8-44.

Example 8-44 The /etc/default/grub file

```
# cat /etc/default/grub
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb
quiet transparent_hugepage=never"
GRUB_DISABLE_RECOVERY="true"
```

After the /etc/default/grub file is updated, run the following command to generate a new configuration file, and reboot the server to activate the new configuration, as shown in Example 8-45.

Example 8-45 Generate new grub configuration file

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
Generating grub configuration file ...
Found linux image: /boot/vmlinuz-3.10.0-327.el7.x86_64
Found initrd image: /boot/initramfs-3.10.0-327.el7.x86_64.img
Found linux image: /boot/vmlinuz-0-rescue-de0320fd093243d8955006cc0e7efecb
Found initrd image: /boot/initramfs-0-rescue-de0320fd093243d8955006cc0e7efecb.img
done

# shutdown -r now
```

8.2.19 Setting disk I/O scheduler on Linux

Many disk I/O schedulers are available in Linux. The deadline I/O scheduler caps maximum latency per request and maintains a good disk throughput, which is best for disk-intensive database applications. Oracle suggests using Deadline scheduler for disks that are used in Oracle Database.

To check the current disk I/O scheduler on Linux, run the command shown in Example 8-46.

Example 8-46 Show Disk I/O Scheduler

```
# cat /sys/block/sd*/queue/scheduler
noop [deadline] cfq
```

```
noop [deadline] cfq
noop [deadline] cfq
noop [deadline] cfq
noop [deadline] cfq
```

Red Hat Enterprise Linux 7.x uses deadline as the default I/O scheduler for all SCSI devices. No additional steps that are required to select deadline IO scheduler for SAN devices in RHEL 7.x.

8.2.20 Create disk partitions

Because ASMLib supports disk partitions only, you must create partitions on disks to be managed by Oracle ASM.

It is important to identify the disk and align with storage planning mentioned in 8.1.2, “Storage planning” on page 104 before creating disk partitions. Example 8-47 shows how to list the size of disks.

Example 8-47 List size of disks

```
# fdisk -l|grep 'Disk /dev'
Disk /dev/sda: 214.7 GB, 214748364800 bytes, 419430400 sectors
Disk /dev/sdb: 5368 MB, 5368709120 bytes, 10485760 sectors
Disk /dev/sdc: 5368 MB, 5368709120 bytes, 10485760 sectors
Disk /dev/sdd: 5368 MB, 5368709120 bytes, 10485760 sectors
Disk /dev/sde: 85.9 GB, 85899345920 bytes, 167772160 sectors
Disk /dev/sdf: 1099.5 GB, 1099511627776 bytes, 2147483648 sectors
Disk /dev/sdg: 1099.5 GB, 1099511627776 bytes, 2147483648 sectors
Disk /dev/sdh: 1099.5 GB, 1099511627776 bytes, 2147483648 sectors
Disk /dev/sdi: 1099.5 GB, 1099511627776 bytes, 2147483648 sectors
Disk /dev/mapper/rhel-root: 196.5 GB, 196457005056 bytes, 383705088 sectors
Disk /dev/mapper/rhel-swap: 17.3 GB, 17288921088 bytes, 33767424 sectors
```

The most accurate way is to identify it through the disk UUID, which mentioned in 7.2.6, “Enable disk UUID” on page 96, and also the UUID of each shared disk in this environment is listed in Table 8-3 on page 104.

Complete these steps to create a partition that consumes all the available space:

1. Run **fdisk** followed by disk device name to start creating partition.
2. Enter **n** to create a partition.
3. Enter **p** to select the default primary partition type.
4. Enter **1** to create the first primary partition.
5. Press **Enter** to select the default starting sector.
6. Press **Enter** to select the default ending sector.
7. Enter **w** to write the partition table.

Example 8-48 shows the whole process of creating a partition on disk /dev/sdb.

Example 8-48 Create a disk partition

```
# fdisk /dev/sdb
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Device does not contain a recognized partition table
Building a new DOS disklabel with disk identifier 0xc0949480.

Command (m for help): n
Partition type:
   p   primary (0 primary, 0 extended, 4 free)
   e   extended
Select (default p): p
Partition number (1-4, default 1): 1
First sector (2048-10485759, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-10485759, default 10485759):
Using default value 10485759
Partition 1 of type Linux and of size 5 GiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Example 8-49 shows how to list the partitions that are configured on one disk.

Example 8-49 List disk partitions

```
# fdisk -l /dev/sdb

Disk /dev/sdb: 5368 MB, 5368709120 bytes, 10485760 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0xc0949480

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1            2048     10485759     5241856   83  Linux
```

Repeat these steps to create partitions on all disks to be managed by Oracle ASM. Be aware that it is only required to create disk partitions on one node only as all these disks are shared disks.

8.2.21 Install and Configure Oracle ASMLib

Oracle ASMLib consists of three RPM packages as shown in Table 8-8.

Table 8-8 Oracle ASMLib packages

Package Description	RPM Package Name	Source
ASMLib Kernel Driver	kmod-oracleasm	Red Hat CD
Userspace Library	oracleasm-lib	Oracle website
Driver Support Files	oracleasm-support	Oracle website

The later two packages are available from the Oracle website. The link below is for the Red Hat Enterprise Linux 7.x version:

<http://www.oracle.com/technetwork/server-storage/linux/asmlib/rhel7-2773795.html>

Use the commands shown in Example 8-50 to install Oracle ASMLib on each Linux node.

Example 8-50 Install Oracle ASMLib

```
# rpm -ivh kmod-oracleasm-2.0.8-15.el7.x86_64.rpm
# rpm -ivh oracleasm-lib-2.0.12-1.el7.x86_64.rpm
# rpm -ivh oracleasm-support-2.1.8-3.el7.x86_64.rpm
```

Start initial configuration as shown in Example 8-51. Set the user to own ASMLib driver interface to grid and group to own to asmadmin.

Example 8-51 Oracle ASMLib initial configuration

```
# /usr/sbin/oracleasm configure -i
Configuring the Oracle ASM library driver.
```

This will configure the on-boot properties of the Oracle ASM library driver. The following questions will determine whether the driver is loaded on boot and what permissions it will have. The current values will be shown in brackets ('[]'). Hitting <ENTER> without typing an answer will keep that current value. Ctrl-C will abort.

```
Default user to own the driver interface []: grid
Default group to own the driver interface []: asmadmin
Scan for Oracle ASM disks on boot (y/n) [y]: y
Writing Oracle ASM library driver configuration: done
```

Load the oracleasm driver module and mount the ASM driver file system manually, as shown in Example 8-52.

Example 8-52 Load ASMLib driver

```
# /usr/sbin/oracleasm init
Creating /dev/oracleasm mount point: /dev/oracleasm
Loading module "oracleasm": oracleasm
Configuring "oracleasm" to use device physical block size
Mounting ASMLib driver filesystem: /dev/oracleasm
```

To double check whether the *oracleasm* service is running and enabled, run the commands as shown in Example 8-53.

Example 8-53 Check oracleasm service

```
# systemctl status oracleasm.service
? oracleasm.service - Load oracleasm Modules
   Loaded: loaded (/usr/lib/systemd/system/oracleasm.service; enabled; vendor
   preset: disabled)
   Active: active (exited) since Fri 2016-05-27 12:27:16 EDT; 7s ago
   Process: 3087 ExecStart=/usr/sbin/service oracleasm start_sysctl (code=exited,
   status=0/SUCCESS)
   Main PID: 3087 (code=exited, status=0/SUCCESS)

# systemctl is-enabled oracleasm.service
enabled
```

To create an ASM disk, run the command **oracleasm createdisk** followed by the label of ASM disk, and then followed by the device name of the disk. Example 8-54 shows the process to create all ASM disks to be used in this Oracle RAC environment.

Example 8-54 Create ASM disks

```
# oracleasm createdisk OCR1 /dev/sdb1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk OCR2 /dev/sdc1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk OCR3 /dev/sdd1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk REDO /dev/sde1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk DATA1 /dev/sdf1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk DATA2 /dev/sdg1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk DATA3 /dev/sdh1
Writing disk header: done
Instantiating disk: done
# oracleasm createdisk DATA4 /dev/sdi1
Writing disk header: done
Instantiating disk: done
```

Note: It is only necessary to run ASM disks creation commands from one node, the other nodes are able to discover these ASM disks later.

To discover ASM disks on other nodes, run the following commands that are shown in Example 8-55.

Example 8-55 Discover and list ASM disks

```
# oracleasm scandisks
Reloading disk partitions: done
Cleaning any stale ASM disks...
Scanning system for ASM disks...
Instantiating disk "OCR1"
Instantiating disk "OCR2"
Instantiating disk "REDO"
Instantiating disk "OCR3"
Instantiating disk "DATA2"
Instantiating disk "DATA1"
Instantiating disk "DATA4"
Instantiating disk "DATA3"

# oracleasm listdisks
DATA1
DATA2
DATA3
DATA4
OCR1
OCR2
OCR3
REDO
```

8.3 Oracle Grid Infrastructure installation

This section shows how to install Oracle Grid Infrastructure, a prerequisite of installing Oracle RAC Database, which consists of Oracle Clusterware and ASM.

8.3.1 Get installation packages

Oracle Grid Infrastructure and Oracle Database 12c for Linux installation packages can be downloaded from Oracle website below, or they are available on Oracle Database 12c DVDs:

<http://www.oracle.com/technetwork/database/enterprise-edition/downloads/database12c-linux-download-2240591.html>

Copy the Oracle Grid Infrastructure installation packages to one node of the cluster, calculate the checksum of the files, and compare the values to those published on the Oracle website (Figure 8-2) to ensure file integrity as shown in Example 8-56.

Oracle Database 12c Release 1 (12.1.0.2.0) for Linux x86-64

📄 linuxamd64_12c_database_1of2.zip (1,673,544,724 bytes) (cksum - 839029806)
📄 linuxamd64_12c_database_2of2.zip (1,014,530,602 bytes) (cksum - 1187131466)

Directions

1. All files are in the .zip format. There is an unzip utility [here](#) if you need one.
2. Download and unzip both files to the same directory.
3. Installation guides and general Oracle Database 12c documentation are [here](#).

Oracle Database 12c Release 1 Grid Infrastructure (12.1.0.2.0) for Linux x86-64

📄 linuxamd64_12c_grid_1of2.zip (1,747,043,545 bytes) (cksum - 1194876808)
📄 linuxamd64_12c_grid_2of2.zip (646,972,897 bytes) (cksum - 2519919927)

Contains the Grid Infrastructure Software including Oracle Clusterware, Automated Storage Management (ASM), and ASM Cluster File System. Download and install prior to installing Oracle Real Application Clusters, Oracle Real Application Clusters One Node, or other application software in a Grid Environment

Figure 8-2 Checksums of Oracle Database packages

Example 8-56 Verify checksum of Oracle grid infrastructure packages

```
# cksum linuxamd64_12102_grid_1of2.zip
1194876808 1747043545 linuxamd64_12102_grid_1of2.zip
# cksum linuxamd64_12102_grid_2of2.zip
2519919927 646972897 linuxamd64_12102_grid_2of2.zip
# cksum linuxamd64_12102_database_1of2.zip
839029806 1673544724 linuxamd64_12102_database_1of2.zip
# cksum linuxamd64_12102_database_2of2.zip
1187131466 1014530602 linuxamd64_12102_database_2of2.zip
```

Extract the packages for installation as shown in Example 8-57.

Example 8-57 Extract Oracle installation packages

```
# unzip linuxamd64_12102_grid_1of2.zip
# unzip linuxamd64_12102_grid_2of2.zip
# unzip linuxamd64_12102_database_1of2.zip
# unzip linuxamd64_12102_database_2of2.zip
```

8.3.2 Install Oracle Cluster Verification Utility

Oracle provides a tool called Oracle Cluster Verification Utility (CVU) to perform cluster environment tests for Oracle RAC installation. It is bundled with the Oracle Grid Infrastructure software, and needs to be installed on each node before running Grid Infrastructure installation.

Example 8-58 shows the location of this package and how to install it.

Example 8-58 Install Oracle Cluster Verification Utility

```
# rpm -ivh grid/rpm/cvuqdisk-1.0.9-1.rpm
Preparing... ##### [100%]
Using default group oinstall to install package
Updating / installing...
 1:cvuqdisk-1.0.9-1 ##### [100%]
```

8.3.3 Oracle Grid Infrastructure installation

Oracle Grid Infrastructure and Oracle Database installation requires GUI access. It can be done either using Linux console or redirecting to X11 software on client machine. The installation steps in this section are from a virtual machine console provided by VMware vSphere. To launch a console for virtual machine **ITSO_VM1**, from the vSphere Web Client, select **ITSO_VM1** in the inventory, right-click **ITSO_VM1**, and select **Open Console** from pop-up menu, as shown in Figure 8-3.





Figure 8-3 Open virtual machine console

To install Oracle Grid Infrastructure from Linux console, complete these steps:

1. It is required to use grid owner user (grid in this environment) to start Grid Infrastructure installation. Log in with grid from the console, open a terminal, and start the installer by running **runInstaller** in the grid directory, as shown in Figure 8-4.

A terminal window titled 'grid@itsovm1:/software/grid' with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal shows the following commands and output:

```
[grid@itsovm1 Desktop]$ cd /software/grid
[grid@itsovm1 grid]$ ls
install    rpm          runInstaller stage
response  runcluvfy.sh sshsetup    welcome.html
[grid@itsovm1 grid]$ ./runInstaller
Starting Oracle Universal Installer...

Checking Temp space: must be greater than 415 MB.   Actual 174075 MB   Passed
Checking swap space: must be greater than 150 MB.   Actual 16487 MB   Passed
Checking monitor: must be configured to display at least 256 colors.   Actual 1
6777216   Passed
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2016-05-26_02
-33-29PM. Please wait ...[grid@itsovm1 grid]$
```

Figure 8-4 Run Oracle Grid Infrastructure installer

2. From the Oracle Grid Infrastructure installation options window, select **Install and Configure Oracle Grid Infrastructure for a Cluster** because it will be an Oracle RAC Database environment, as shown in Figure 8-5. Click **Next** to continue.




Figure 8-5 Grid Infrastructure installation options

3. From Oracle Grid Infrastructure Cluster Type option window, select **Configure a Standard Cluster**, as shown in Figure 8-6, and click **Next** to continue.




Figure 8-6 Grid cluster type options

4. For Oracle Grid Infrastructure Installation Type option, select **Advanced Installation**, as shown in Figure 8-7, and click **Next** to continue.




Figure 8-7 Grid installation type options

5. Oracle Grid Infrastructure supports multiple languages choices. **English** is selected by default. If additional language support is needed, select the language name from the left box, and click the **Arrow** button to add it to the list, as shown in Figure 8-8. Click **Next** to continue.




Figure 8-8 Grid language options

- From cluster name and SCAN setting window, enter the cluster name and SCAN name. Remember that the SCAN name here should be the one registered in DNS server. Click to clear **Configure GNS** because GNS is not used in this environment, as shown in Figure 8-9. Click **Next** to continue.




Figure 8-9 Cluster name and SCAN name

- The cluster node information window now shows node in this cluster including public node name and private node name, as shown in Figure 8-10. To add a node to the cluster, click the **Add** button.




Figure 8-10 Cluster node information

- Enter both the public hostname and private hostname of one node, as shown in Figure 8-11. Remember that these hostnames should be registered in the DNS server as planned. Click **OK** to add one node to the cluster.




Figure 8-11 Add one node to cluster

- After all the nodes are added to the cluster, click **SSH Connectivity**. Input the password of grid user, as shown in Figure 8-12, then click **Setup** to set up passwordless SSH connectivity between all cluster member nodes.




Figure 8-12 Setup passwordless SSH connectivity for grid user

- It takes a few seconds to set up SSH connectivity. After the process is complete, a prompt is displayed as shown in Figure 8-13. Click **OK** to continue.




Figure 8-13 SSH setup completion

11. Network interface usage configuration dialog lists all network interfaces found on cluster node. Make sure to select the correct usage for each interface as shown in Figure 8-14. Sometimes the interface is not shown due to the device name being not consistent on each node in the cluster. Make your decision based on subnet information because it will be consistent on all nodes. Click **Next** to continue if the choices are made.




Figure 8-14 Network interface usage

Note: Set the usage of any interface that is not used in Oracle environment to **Do Not Use**.

12. Oracle supports several ways to place Oracle Cluster Registry (OCR) and vote disk files, select **Use Standard ASM for storage**, as shown in Figure 8-15, and click **Next** to continue.




Figure 8-15 Storage options for OCR and voting disk files

13. Complete these steps to create an ASM disk group to store OCR and vote disks, as shown in Figure 8-16 on page 142:

- Disk group name** Specify the name of the ASM disk group here, such as **OCR** in this example.
- Redundancy** The number of vote disks depends on the Redundancy level chosen here: **EXTERNAL** creates one voting disk, **NORMAL** creates three voting disks, **HIGH** creates five voting disks. Select **NORMAL** because three voting disks are planned.
- Allocation Unit Size** An allocation unit is the fundamental unit of allocation within a disk group, considering CPU utilization and memory consumption. It is suggested to set allocation unit size to 4 MB.
- Add Disk** Select the names of the disks to be planned for creating the OCR disk group, usually the name of disk would be "ORCL:" followed by the label of the ASM disk. If the candidate list is empty, click Change Discovery Path button and change it to /dev/oracleasm/disks/* and retry.

For more information about ASM disk group, see the *Automatic Storage Management Administrator's Guide 12c Release 1* at:

<https://docs.oracle.com/database/121/OSTMG/E41058-11.pdf>

After the setting is finished, click **Next** to continue as shown in Figure 8-16.




Figure 8-16 Create ASM disk group for OCR and voting disk files

14. Oracle supports setting either the same or different passwords for SYS and ASMSNMP users. Select **Use same passwords for these accounts** and input the passwords, as shown in Figure 8-17. The SYS user is an ASM administration user with SYSASM privileges, and it is suggested to use a strong password. After the setting is finished, click **Next** to continue.




Figure 8-17 Set password for ASM users

15. Failure Isolation is a feature to isolate a failed node from the rest of the cluster to prevent data corruption. It can be done through Intelligent Management Platform Interface (IPMI) specification interface if it is available. Select **Do not use Intelligent Management Platform Interface specification (IPMI)** as shown in Figure 8-18, and click **Next** to continue if IPMI is not available.




Figure 8-18 Failure isolation options

16. Enterprise Manager Cloud Control is a tool that allows the database administrator to monitor and manage databases from a single console. Click to clear **Register with Enterprise Manager Cloud Control** as shown in Figure 8-19, because there is no OMS host in this environment. Click **Next** to continue.




Figure 8-19 Management options

17. To assign OS groups for Oracle ASM Management, select them according to Table 8-9, as shown in Figure 8-20 on page 146. Click **Next** to continue.

Table 8-9 ASM groups

Oracle ASM role	OS Group
Oracle ASM Administrator (OSASM) Group	asmadmin
Oracle ASM DBA (OSDBA for ASM) Group	asmdba
Oracle ASM Operator (OSOPER for ASM) Group	asmoper




Figure 8-20 Assign OS groups for ASM

18. Specify the locations for ORACLE_BASE and Grid Infrastructure software, as shown in Figure 8-21, and click **Next** to continue.




Figure 8-21 Specify installation location

19. Specify the location for installation metadata files, as shown in Figure 8-22, and click **Next** to continue.




Figure 8-22 Specify inventory directory

20. Two scripts are required to be run with root user on each node during Grid Infrastructure installation. The installation program offers options to run these scripts automatically, as shown in Figure 8-23. It is suggested to run these scripts manually. To do so, click to clear the **Automatically run configuration scripts**, and click **Next** to continue.




Figure 8-23 Root script execution configuration

21. The installation program performs prerequisite checks as shown in Figure 8-24. This process usually takes several minutes.




Figure 8-24 Perform prerequisite checks

22. The verification result shows as soon as the prerequisite checks are finished. Figure 8-25 shows that there is one warning on /dev/shm. This is a known issue and can be ignored, so select **Ignore All** and click **Next** to continue.

Note: The detailed warning information about /dev/shm, described as The size of in-memory file system mounted at /dev/shm is “24576” megabytes which does not match the size in /etc/fstab as “0” megabytes, is a known issue documented on the Oracle website at:

<https://support.oracle.com/epmos/faces/DocumentDisplay?id=1918620.1>




Figure 8-25 Prerequisite check result

23. A warning prompt is displayed as shown in Figure 8-26 to confirm the ignoring action. Click **Yes** to continue.




Figure 8-26 Confirm ignoring prerequisites

24. Figure 8-27 shows a summary of the installation settings. Click **Install** to continue.




Figure 8-27 Grid infrastructure installation summary

25. Grid Infrastructure installation begins. Check the progress bar for installation process as shown in Figure 8-28.




Figure 8-28 Installation in progress

26. During the installation, a window prompts you to run two scripts as root user, as shown in Figure 8-29.




Figure 8-29 Prompt to execute root scripts

27. Open a terminal window and login as root user on the first node in cluster, known as itsovm1 in this environment. Run the script /u01/app/oraInventory/orainstRoot.sh as shown in Example 8-59.

Example 8-59 Running /u01/app/oraInventory/orainstRoot.sh

```
# /u01/app/oraInventory/orainstRoot.sh
Changing permissions of /u01/app/oraInventory.
Adding read,write permissions for group.
Removing read,write,execute permissions for world.

Changing groupname of /u01/app/oraInventory to oinstall.
The execution of the script is complete.
```

28. Repeat the above step on the other three hosts in sequence.

29. Go back to the terminal on the first node in cluster, known as itsovm1 in the environment, and run the script /u01/app/12.1.0/grid/root.sh using the root user, as shown in Example 8-60.

Example 8-60 Run /u01/app/12.1.0/grid/root.sh

```
# /u01/app/12.1.0/grid/root.sh
Performing root user operation.

The following environment variables are set as:
ORACLE_OWNER= grid
ORACLE_HOME= /u01/app/12.1.0/grid

Enter the full pathname of the local bin directory: [/usr/local/bin]:
```

```
Copying dbhome to /usr/local/bin ...
Copying oraenv to /usr/local/bin ...
Copying coraenv to /usr/local/bin ...
```

```
Creating /etc/oratab file...
```

```
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root script.
```

```
Now product-specific root actions will be performed.
```

```
Using configuration parameter file:
```

```
/u01/app/12.1.0/grid/crs/install/crsconfig_params
```

```
2016/05/27 14:43:53 CLSRSC-4001: Installing Oracle Trace File Analyzer (TFA)
Collector.
```

```
2016/05/27 14:44:16 CLSRSC-4002: Successfully installed Oracle Trace File
Analyzer (TFA) Collector.
```

```
2016/05/27 14:44:17 CLSRSC-363: User ignored prerequisites during installation
```

```
OLR initialization - successful
```

```
root wallet
root wallet cert
root cert export
peer wallet
profile reader wallet
pa wallet
peer wallet keys
pa wallet keys
peer cert request
pa cert request
peer cert
pa cert
peer root cert TP
profile reader root cert TP
pa root cert TP
peer pa cert TP
pa peer cert TP
profile reader pa cert TP
profile reader peer cert TP
peer user cert
pa user cert
```

```
2016/05/27 14:44:56 CLSRSC-330: Adding Clusterware entries to file
'oracle-ohasd.service'
```

```
CRS-4133: Oracle High Availability Services has been stopped.
CRS-4123: Oracle High Availability Services has been started.
CRS-4133: Oracle High Availability Services has been stopped.
CRS-4123: Oracle High Availability Services has been started.
CRS-2672: Attempting to start 'ora.evmd' on 'itsovm1'
CRS-2672: Attempting to start 'ora.mdnsd' on 'itsovm1'
CRS-2676: Start of 'ora.mdnsd' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.evmd' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.gpnpd' on 'itsovm1'
CRS-2676: Start of 'ora.gpnpd' on 'itsovm1' succeeded
```

```

CRS-2672: Attempting to start 'ora.cssdmonitor' on 'itsovm1'
CRS-2672: Attempting to start 'ora.gipcd' on 'itsovm1'
CRS-2676: Start of 'ora.cssdmonitor' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.gipcd' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.cssd' on 'itsovm1'
CRS-2672: Attempting to start 'ora.diskmon' on 'itsovm1'
CRS-2676: Start of 'ora.diskmon' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.cssd' on 'itsovm1' succeeded

```

ASM created and started successfully.

Disk Group OCR created successfully.

```

CRS-2672: Attempting to start 'ora.crf' on 'itsovm1'
CRS-2672: Attempting to start 'ora.storage' on 'itsovm1'
CRS-2676: Start of 'ora.storage' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.crf' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.crsd' on 'itsovm1'
CRS-2676: Start of 'ora.crsd' on 'itsovm1' succeeded
CRS-4256: Updating the profile
Successful addition of voting disk 7cd47f01f1e84f82bfef846752f4bab2.
Successful addition of voting disk ab468b2ebcca4ffbbf793400a8301cec.
Successful addition of voting disk 7316d6b073a44f6dbf03323dcbb983ef.
Successfully replaced voting disk group with +OCR.

```

```

CRS-4256: Updating the profile
CRS-4266: Voting file(s) successfully replaced
## STATE File Universal Id File Name Disk group
--  -----
 1. ONLINE 7cd47f01f1e84f82bfef846752f4bab2 (ORCL:OCR1) [OCR]
 2. ONLINE ab468b2ebcca4ffbbf793400a8301cec (ORCL:OCR2) [OCR]
 3. ONLINE 7316d6b073a44f6dbf03323dcbb983ef (ORCL:OCR3) [OCR]

```

Located 3 voting disk(s).

```

CRS-2791: Starting shutdown of Oracle High Availability Services-managed
resources on 'itsovm1'
CRS-2673: Attempting to stop 'ora.crsd' on 'itsovm1'
CRS-2677: Stop of 'ora.crsd' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.storage' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.mdnsd' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.gpnpd' on 'itsovm1'
CRS-2677: Stop of 'ora.storage' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.crf' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.ctssd' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.evmd' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.asm' on 'itsovm1'
CRS-2677: Stop of 'ora.mdnsd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.gpnpd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.crf' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.evmd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.ctssd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.asm' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.cluster_interconnect.haip' on 'itsovm1'
CRS-2677: Stop of 'ora.cluster_interconnect.haip' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.cssd' on 'itsovm1'
CRS-2677: Stop of 'ora.cssd' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.gipcd' on 'itsovm1'

```

```

CRS-2677: Stop of 'ora.gipcd' on 'itsovm1' succeeded
CRS-2793: Shutdown of Oracle High Availability Services-managed resources on
'itsovm1' has completed
CRS-4133: Oracle High Availability Services has been stopped.
CRS-4123: Starting Oracle High Availability Services-managed resources
CRS-2672: Attempting to start 'ora.mdnsd' on 'itsovm1'
CRS-2672: Attempting to start 'ora.evmd' on 'itsovm1'
CRS-2676: Start of 'ora.mdnsd' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.evmd' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.gpnpd' on 'itsovm1'
CRS-2676: Start of 'ora.gpnpd' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.gipcd' on 'itsovm1'
CRS-2676: Start of 'ora.gipcd' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'itsovm1'
CRS-2676: Start of 'ora.cssdmonitor' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.cssd' on 'itsovm1'
CRS-2672: Attempting to start 'ora.diskmon' on 'itsovm1'
CRS-2676: Start of 'ora.diskmon' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.cssd' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.cluster_interconnect.haip' on 'itsovm1'
CRS-2672: Attempting to start 'ora.ctssd' on 'itsovm1'
CRS-2676: Start of 'ora.ctssd' on 'itsovm1' succeeded
CRS-2676: Start of 'ora.cluster_interconnect.haip' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.asm' on 'itsovm1'
CRS-2676: Start of 'ora.asm' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.storage' on 'itsovm1'
CRS-2676: Start of 'ora.storage' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.crf' on 'itsovm1'
CRS-2676: Start of 'ora.crf' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.crsd' on 'itsovm1'
CRS-2676: Start of 'ora.crsd' on 'itsovm1' succeeded
CRS-6023: Starting Oracle Cluster Ready Services-managed resources
CRS-6017: Processing resource auto-start for servers: itsovm1
CRS-6016: Resource auto-start has completed for server itsovm1
CRS-6024: Completed start of Oracle Cluster Ready Services-managed resources
CRS-4123: Oracle High Availability Services has been started.
2016/05/27 14:50:28 CLSRSC-343: Successfully started Oracle Clusterware stack

CRS-2672: Attempting to start 'ora.asm' on 'itsovm1'
CRS-2676: Start of 'ora.asm' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.OCR.dg' on 'itsovm1'
CRS-2676: Start of 'ora.OCR.dg' on 'itsovm1' succeeded
2016/05/27 14:51:44 CLSRSC-325: Configure Oracle Grid Infrastructure for a
Cluster ... succeeded

```

30. Repeat the above step on the other three hosts in sequence.

31. After the scripts complete, click **OK** to continue. The installation program finishes the rest of procedures and then shows a successful installation window as shown in Figure 8-30.





Figure 8-30 Completion of Grid infrastructure installation

8.4 Oracle RAC Database installation

This section introduces how to install Oracle RAC Database software on four nodes. Complete these steps to start Oracle RAC Database installation:

1. It is required to use Oracle database owner user (also known as. oracle in this environment) to start the Oracle Database installation. Log in with oracle from console, open a terminal, and start the installer by running **runInstaller** in database directory, as shown in Figure 8-31.

A terminal window titled 'oracle@itsovm1:/software/database' with a menu bar (File, Edit, View, Search, Terminal, Help). The terminal shows the following commands and output:

```
[oracle@itsovm1 Desktop]$ cd /software/database/
[oracle@itsovm1 database]$ ls
install response rpm runInstaller sshsetup stage welcome.html
[oracle@itsovm1 database]$ ./runInstaller
Starting Oracle Universal Installer...

Checking Temp space: must be greater than 500 MB.   Actual 167145 MB   Passed
Checking swap space: must be greater than 150 MB.   Actual 16487 MB   Passed
Checking monitor: must be configured to display at least 256 colors.   Actual 1
6777216   Passed
Preparing to launch Oracle Universal Installer from /tmp/OraInstall2016-05-26_05
-29-34PM. Please wait ...
```

Figure 8-31 Run Oracle database installer

- Oracle supports sending security updates if an Oracle Support email and password is provided. Click to clear **I wish to receive security updates via My Oracle Support** if it is not necessary, as shown in Figure 8-32. Click **Next** to continue.




Figure 8-32 Configure security updates

- A window is displayed appears to confirm the action, as shown in Figure 8-33. Click **Yes** to continue.



Figure 8-33 Confirm not to provide email

4. Select **Install database software only** and click **Next** to continue, as shown in Figure 8-34.




Figure 8-34 Oracle database installation options

5. Select **Create Real Application Clusters database installation** and click **Next** to continue, as shown in Figure 8-35.




Figure 8-35 Oracle database type options

6. Select the four nodes in the list, click the **SSH Connectivity** button, and input the password of the Oracle user. Then click the **Setup** button to set up passwordless SSH connectivity between all cluster member nodes, as shown in Figure 8-36.




Figure 8-36 Config nodes and passwordless SSH connectivity for oracle user

7. A dialog will be prompted after the passwordless SSH connectivity setup is completed, as shown in Figure 8-37.




Figure 8-37 Passwordless SSH connectivity setup completion

8. Oracle database supports multiple languages, **English** is selected by default. If additional language support is needed, select the language name from the left box, and click the **Arrow** button to add it to the list, as shown in Figure 8-38, then click **Next** to continue.




Figure 8-38 Oracle database language options

9. There are two Oracle database editions. Make sure that **Enterprise Edition** is selected, as shown in Figure 8-39, and click **Next** to continue.




Figure 8-39 Oracle database edition options

10. Specify the locations for ORACLE_BASE and Oracle Database software (Oracle Home), as shown in Figure 8-40, and click **Next** to continue.



Figure 8-40 Oracle database installation location

11. Assign OS groups for Oracle Database Management according to the criteria in Table 8-10, as shown in Figure 8-41 on page 168, and click **Next** to continue.

Table 8-10 Oracle Database groups

Oracle Database Role	OS Group
Database Administrator (OSDBA) group	dba
Database Operator (OSOPER) group	oper
Database Backup and Recovery (OSBACKUPDBA) group	backupdba
Data Guard administrative (OSDGDBA) group	dgdba
Encryption Key Management administrative (OSKMDBA) group	kmdba




Figure 8-41 Oracle database OS groups

12. The installation program performs prerequisite checks as shown in Figure 8-42. This process usually takes several minutes.




Figure 8-42 Perform prerequisite checks

13. The verification result shows as soon as the prerequisite checks are finished, Figure 8-43 shows that there is one warning on /dev/shm. This is a known issue, and can be ignored. Select **Ignore All** and click **Next** to continue.




Figure 8-43 Prerequisite check result

14. Figure 8-44 shows a summary of the installation settings. Click **Install** to continue.




Figure 8-44 Oracle database installation summary

15. Oracle Database installation begins. Check the progress bar for installation process as shown in Figure 8-45.




Figure 8-45 Oracle database installation in progress

16. During the installation, a window will open prompting you to run a script as root user, as shown in Figure 8-46.




Figure 8-46 Prompt to execute root script

17. Open a terminal window and login as root user on the first node in cluster, known as itsovm1 in this environment. Run the `/u01/app/oracle/product/12.1.0/dbhome/root.sh` script as shown in Example 8-61.

Example 8-61 Execute /u01/app/oracle/product/12.1.0/dbhome/root.sh

```
# /u01/app/oracle/product/12.1.0/dbhome/root.sh
Performing root user operation.
```

```
The following environment variables are set as:
ORACLE_OWNER= oracle
ORACLE_HOME= /u01/app/oracle/product/12.1.0/dbhome
```

```
Enter the full pathname of the local bin directory: [/usr/local/bin]:
The contents of "dbhome" have not changed. No need to overwrite.
The contents of "oraenv" have not changed. No need to overwrite.
The contents of "coraenv" have not changed. No need to overwrite.
```

```
Entries will be added to the /etc/oratab file as needed by
Database Configuration Assistant when a database is created
Finished running generic part of root script.
Now product-specific root actions will be performed.
```

18. Repeat the above step on the other three hosts in sequence.

19. After you finish running the script, click **OK** to continue. The installation program finishes the rest of procedures and shows a successful installation window as shown in Figure 8-47.




Figure 8-47 Oracle RAC database installation completion

8.5 Create Oracle ASM disk groups

This section describes how to create Oracle ASM disk groups using a graphic tool called ASM Configuration Assistant. ASM disk groups are used to store Oracle Database files. To create ASM disk groups, complete these steps:

1. To launch ASM Configuration Assistant, log in with `grid` from the console, open a terminal, and issue the command `/u01/app/12.1.0/grid/bin/asmca`.
2. ASM Configuration Assistant lists existing ASM disk groups information, as shown in Figure 8-48. Click **Create** button to create an ASM disk group.




Figure 8-48 ASM configuration assistant

3. Click **Show Advanced Options**, as shown in Figure 8-49, and complete these steps:
 - a. Input new ASM disk group name, such as REDO in this environment.
 - b. Select the redundancy level as **External (None)** because it is not necessary to keep multiple data copies.
 - c. Check the disks to be used to create REDO disk group from the list, for example ORCL:REDO in this environment.
 - d. Select **Allocation Unit size** as **4 MB**.
 - e. Click **OK** to create the ASM disk group.

Create Disk Group

Disk Group Name

Redundancy
 Redundancy is achieved by storing multiple copies of the data on different failure groups. Normal redundancy needs disks from at least two different failure groups, and high redundancy from at least three different failure groups.

High
 Normal
 External (None)

Select Member Disks
 Show Eligible
 Show All

Quorum failure groups are used to store voting files in extended clusters and do not contain any user data. They require ASM compatibility of 11.2 or higher.

<input type="checkbox"/>	Disk Path	Header Status	Disk Name	Size (MB)	Quorum
<input type="checkbox"/>	ORCL:DATA1	PROVISIONED		1048575	<input type="checkbox"/>
<input type="checkbox"/>	ORCL:DATA2	PROVISIONED		1048575	<input type="checkbox"/>
<input type="checkbox"/>	ORCL:DATA3	PROVISIONED		1048575	<input type="checkbox"/>
<input type="checkbox"/>	ORCL:DATA4	PROVISIONED		1048575	<input type="checkbox"/>
<input checked="" type="checkbox"/>	ORCL:REDO	PROVISIONED		81919	<input type="checkbox"/>

Note: If you do not see the disks which you believe are available, check the Disk Discovery Path and read/write permissions on the disks. The Disk Discovery Path limits set of disks considered for discovery.

Disk Discovery Path: <default>

Disk Group Attributes
 An allocation unit (AU) is the fundamental unit in which contiguous disk space is allocated to ASM files. ASM file extent size is a multiple of AUs. The AU size cannot be modified later.

Allocation Unit Size (MB)

Specify minimum software versions for ASM, Database and ASM volumes that this disk group need to be compatible with.

ASM Compatibility

Database Compatibility

ADVM Compatibility

Refer Oracle Automatic Storage Management Administrator's Guide for more details on the Compatibility matrix.

Figure 8-49 Create ASM disk group

- Repeat the above steps to create more ASM disk groups. After you are finished, all the ASM disk groups are listed in ASM Configuration Assistant, as shown in Figure 8-50.




Figure 8-50 List ASM disk groups

8.6 Oracle RAC Database creation

This section describes how to create a sample Oracle RAC database using a tool called Database Configuration Assistant.

8.6.1 Oracle RAC Database creation

To create an Oracle RAC Database, complete these steps:

1. To launch Database Configuration Assistant, log in with `oracle` from console, open a terminal, and issue the command `/u01/app/oracle/product/12.1.0/dbhome/bin/dbca`.
2. In the Database Operation window, select **Create Database**, as shown in Figure 8-51, and click **Next** to continue.




Figure 8-51 Create Oracle database

3. Select **Advanced Mode** to allow more flexible configuration, as shown in Figure 8-52, and click **Next** to continue.




Figure 8-52 Creation mode options

4. Select **Database Type** as **Oracle Real Application Clusters (RAC) database**, **Configuration Type** as **Admin-Managed**, select **General Purpose for Transaction Processing template**, as shown in Figure 8-53, and click **Next** to continue.




Figure 8-53 Database templates

5. Enter the **Global Database Name** and **SID Prefix**, and clear **Create as Container Database**, as shown in Figure 8-54, and click **Next** to continue.

Database Configuration Assistant – Create Database – Step 4 of 15

ORACLE 12c DATABASE

Database Identification

Database Operation
Creation Mode
Database Template
Database Identification
Database Placement
Management Options
Database Credentials
Storage Locations
Database Options
Initialization Parameters
Creation Options
Prerequisite Checks
Summary
Progress Page
Finish

Provide the identifier information required to access the database uniquely. An Oracle database is uniquely identified by a Global Database Name, typically of the form "name.domain". A database is referenced by an Oracle instance on each cluster database node. Specify a prefix to be used to name the cluster database instances.

Global Database Name: sample

SID Prefix: sample

Create As Container Database

Creates a database container for consolidating multiple databases into a single database and enables database virtualization. A container database (CDB) can have zero or more pluggable databases (PDB).

Create an Empty Container Database

Create a Container Database with one or more PDBs

Number of PDBs: 1

PDB Name:

Help < Back Next > Finish Cancel

Figure 8-54 Database identification

6. Only the node where DBCA is run from is chosen to join Oracle RAC Cluster by default. Select the other three nodes in the left pane, and click the **Right Arrow** to add them to Oracle RAC Cluster, as shown in Figure 8-55.




Figure 8-55 Select Oracle RAC nodes

7. Verify whether the list of nodes is correct, as shown in Figure 8-56, and click **Next** to continue.




Figure 8-56 Select all 4 nodes in cluster

8. Clear **Run Cluster Verification Utility (CVU) Checks Periodically**, **Configure Enterprise Manager (EM) Database Express**, and **Register with Enterprise Manager (EM) Cloud Control**, as shown in Figure 8-57, and click **Next** to continue.




Figure 8-57 Oracle database management options

- Set password for SYS and SYSTEM users. Select **Use the same administrative Password for All accounts**, and input the password, as shown in Figure 8-58. The SYS user is a database administration user with SYSDBA privileges. It is suggested to use a strong password. After the setting is finished, click **Next** to continue.




Figure 8-58 Set password for Oracle database

10. In the storage locations window, as shown in Figure 8-59, complete these steps:
 - a. Select **Automatic Storage Management (ASM)** for **Database files Storage Type**.
 - b. Select **Use Common Location for All Database Files**.
 - c. Input the name of the ASM disk group storing database files, such as +DATA in this environment.
 - d. Select **Use Oracle-Managed Files**.




Figure 8-59 Oracle database storage locations

- e. Click **Multiplex Redo Logs and Control Files** to open the window shown in Figure 8-60.




Figure 8-60 Specify redo log files location

- f. Input the name of ASM disk group storage redo logs and control files, such as +REDO in this environment. Click **OK** to continue.
- g. Clear **Specify Fast Recovery Area**.
- h. Clear **Enable Archiving**.
- i. Click **Next** to continue.

11. Figure 8-61 shows the database options. Keep the default settings, and click **Next** to continue.




Figure 8-61 Oracle database options

12. From the initialization parameters window, select **Typical Settings** for memory and clear **Use Automatic Memory Management** because AMM is not compatible with HugePages, as shown in Figure 8-62. Click **Next** to continue.




Figure 8-62 Oracle database initialization parameters

13. Select **Create Database**, and clear **Generate Database Create Scripts**, as shown in Figure 8-63, then click **Customize Storage Locations** button to modify redo log size.




Figure 8-63 Oracle database creation options

14. In the Customize Storage window, it lists control files, datafiles, and Redo Log groups, as shown in Figure 8-64.




Figure 8-64 Customize storage

15. Click **Redo Logs Group 1** to show the current Redo log file size in the right pane, as shown in Figure 8-65. Change it to **1024 MB** and click the **Apply** button.




Figure 8-65 Customize redo log files

16. Repeat the above steps to change the file size for all eight Redo Log files, as shown in Figure 8-66, and click **OK** to continue.




Figure 8-66 List redo log files size

17. The Database Configuration Assistant program performs prerequisite checks as shown in Figure 8-67. This process usually takes several minutes.




Figure 8-67 Perform prerequisite checks

18. Figure 8-68 shows a summary of creating database settings. Click **Finish** to continue.




Figure 8-68 Oracle database creation summary

19. Figure 8-69 shows the progress of database creation. The process usually takes about 10 minutes.




Figure 8-69 Oracle database creation in progress

20. After the database creation is completed, a window like that shown in Figure 8-70 is shown. Only SYS and SYSTEM users are unlocked by default, and to unlock other users or change password for other users, click **Password Management**.




Figure 8-70 Oracle database creation completion

21. Lock/unlock database users, or change the password for users in the Password Management window, as shown in Figure 8-71. Click **OK** to finish database creation.




Figure 8-71 Oracle database password management

8.6.2 Update grid and oracle user profile

Oracle RAC Database is a database with multiple instances from different nodes, which means the Oracle instance ID on each node is different. It is suggested to configure Oracle instance ID for the grid and oracle user in its profile on each node.

The grid user's Oracle instance ID in this environment is shown in Table 8-11.

Table 8-11 grid user Oracle instance ID

Node	Instance ID
itsovm1	+ASM1
itsovm2	+ASM2
itsovm3	+ASM3
itsovm4	+ASM4

The oracle user's Oracle instance ID in this environment is shown in Table 8-12.

Table 8-12 oracle user Oracle instance ID

Node	Instance ID
itsovm1	sample1
itsovm2	sample2
itsovm3	sample3
itsovm4	sample4

To configure grid user profile on itsovm1, log in to the node with the grid user, use a text editor to open the .bash_profile file, and append the lines shown in Example 8-62.

Example 8-62 The grid user profile

```
export ORACLE_BASE=/u01/app/grid
export ORACLE_HOME=/u01/app/12.1.0/grid
export ORACLE_SID=+ASM1
export PATH=$ORACLE_HOME/bin:$PATH
umask 022
```

Repeat the above step on the other three nodes. Remember to change the Oracle instance ID.

To configure oracle user profile on itsovm1, log in to the node with the oracle user, use a text editor to open the .bash_profile file, and append the lines shown in Example 8-63.

Example 8-63 The oracle user profile

```
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=/u01/app/oracle/product/12.1.0/dbhome
export ORACLE_SID=sample1
export PATH=$ORACLE_HOME/bin:$PATH
umask 022
```

Repeat the above step on the other three nodes. Remember to change the Oracle instance ID.

8.6.3 HugePages Configuration

Usually the default memory page size is 4 KB in Linux. HugePages has been available since Linux Kernel 2.6. You can improve system performance by reducing the amount of system resources that are required to access page table entries by using HugePages in Linux.

To check whether HugePages is configured in Linux, run the command shown in Example 8-64.

Example 8-64 Show HugePages number

```
# grep -i HugePages /proc/meminfo
AnonHugePages:      0 kB
HugePages_Total:    0
HugePages_Free:     0
HugePages_Rsvd:     0
HugePages_Surp:     0
Hugepagesize:       2048 kB
```

HugePages_Total is the number of HugePages that are configured on the host. It is not configured by default. To configure HugePages on Oracle RAC node, complete these steps:

1. Make sure that the Oracle database is up and running. Check whether oracle and grid user are shown in the output of the `ipcs -m` command, as shown in Example 8-65.

Example 8-65 Show shared memory segments

```
# ipcs -m

----- Shared Memory Segments -----
key          shmid      owner      perms      bytes      nattch     status
0x00000000  221380608  grid       640        4096       0
0x00000000  221413377  grid       640        4096       0
0xf675ea7c  221446146  grid       640        24576     33
0x00000000  275349507  oracle     640        8388608   295
0x00000000  275382276  oracle     640        75967234048 295
0x00000000  275415045  oracle     640        262144000 295
0xc1c7c898  275447814  oracle     640        24576     295
0x00000000  278560775  grid       640        4194304   39
0x00000000  278593544  grid       640        780140544 39
0x00000000  278626313  grid       640        6291456   39
0x381a5a60  278659082  grid       640        16384     39
0x00000000  1368817675 root       600        524288    2      dest
0x00000000  1368850444 root       600        4194304   2      dest
0x00000000  1369014285 root       600        4194304   2      dest
0x00000000  1370161167 root       600        2097152   2      dest
```

2. Log in to one node with root user, create a script named `hugepages.sh` with the content shown in Example 8-66.

Example 8-66 Script to calculate Hugepages number

```
#!/bin/bash
HPG_SZ=`grep Hugepagesize /proc/meminfo | awk {'print $2'}`
NUM_PG=1
for SEG_BYTES in `ipcs -m | awk {'print $5'} | grep "[0-9][0-9]*"`
do
    MIN_PG=`echo "$SEG_BYTES/($HPG_SZ*1024)" | bc -q`
    if [ $MIN_PG -gt 0 ]; then
        NUM_PG=`echo "$NUM_PG+$MIN_PG+1" | bc -q`
    fi
done
echo $NUM_PG
```

3. Run the script `hugepages.sh` and document the number shown, for example 36745 in this environment.
4. Append a line with content `vm.nr_hugepages = 36745` to `/etc/sysctl.conf` file, as shown in Example 8-67.

Example 8-67 Modify /etc/sysctl.conf file

```
echo "vm.nr_hugepages = 36745" >> /etc/sysctl.conf
```

5. Repeat the above step on each node in the Oracle RAC cluster.
6. Stop all the database instances and reboot the system.

8.7 Oracle RAC Database Management

This section introduces some Oracle RAC Database management operations, including the following topics:

- ▶ Enable archive log mode
- ▶ Add redo log group
- ▶ Multiplex Oracle control files
- ▶ Update processes and sessions parameters
- ▶ Startup and shutdown of Oracle CRS and Database
- ▶ OCR and voting disk management

8.7.1 Enable archive log mode

Oracle database can run in two log modes: No Archive Mode or Archive Mode. The major difference is that the Oracle archiver background processes (ARCn) archive redo logs in Archive mode. By default, the database is created in No Archive Mode. However, it is suggested to enable database in Archive Mode if it is in production.

For more information about No Archive Mode and Archive Mode, see the *Oracle Database Administrator's Guide*.

To enable Oracle RAC database in Archive Mode, complete these steps:

1. Connect to database with SYS user, and check the current database log mode, Example 8-68 shows database is in No Archive Mode.

Example 8-68 Inquiry database archive log mode

```
[oracle@itsovml ~]$ sqlplus /nolog
```

```
SQL*Plus: Release 12.1.0.2.0 Production on Sat May 28 16:11:14 2016
```

```
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
SQL> connect / as sysdba;  
Connected.
```

```
SQL> archive log list;  
Database log mode          No Archive Mode  
Automatic archival        Disabled  
Archive destination        /u01/app/oracle/product/12.1.0/dbhome/dbs/arch  
Oldest online log sequence 397  
Current log sequence       398
```

```
SQL> select inst_id,instance_name,version,archiver,status from gv$instance;
```

INST_ID	INSTANCE_NAME	VERSION	ARCHIVE STATUS
1	sample1	12.1.0.2.0	STOPPED OPEN
4	sample4	12.1.0.2.0	STOPPED OPEN
3	sample3	12.1.0.2.0	STOPPED OPEN
2	sample2	12.1.0.2.0	STOPPED OPEN

2. Check the current log archive destination, log archive destination state, and log archive format, as shown in Example 8-69.

Example 8-69 Show archive log destination parameter

```
SQL> show parameter log_archive_dest_1
```

NAME	TYPE	VALUE
log_archive_dest_1	string	
log_archive_dest_10	string	
log_archive_dest_11	string	

```
SQL> show parameter log_archive_dest_state_1
```

NAME	TYPE	VALUE
log_archive_dest_state_1	string	enable
log_archive_dest_state_10	string	enable
log_archive_dest_state_11	string	enable

```
SQL> show parameter log_archive_format;
```

NAME	TYPE	VALUE
log_archive_format	string	%t_%s_%r.dbf

3. Modify the log archive destination to the +DATA ASM disk group, as shown in Example 8-70.

Example 8-70 Modify archive log destination

```
SQL> alter system set log_archive_dest_1='LOCATION=+DATA/sample/archivelog';
```

```
System altered.
```

4. Stop the database using oracle user, as shown in Example 8-71.

Example 8-71 Stop Oracle database

```
[oracle@itsovm1 ~]$ srvctl stop listener  
[oracle@itsovm1 ~]$ srvctl stop database -d sample
```

5. Start the Oracle database to mount state with single instance, as shown in Example 8-72.

Example 8-72 Start one Oracle database instance to mount mode

```
[oracle@itsovm1 ~]$ srvctl start instance -d sample -i sample1 -o mount
```

6. Connect to the database with SYS user, and run the command to enable archive mode for the database, as shown in Example 8-73.

Example 8-73 Enable archivelog mode

```
[oracle@itsovm1 ~]$ sqlplus /nolog
```

```
SQL*Plus: Release 12.1.0.2.0 Production on Sat May 28 16:29:58 2016
```

```
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
SQL> conn / as sysdba;
Connected.
SQL> alter database archive log;
Database altered.
```

7. Start the Oracle database to open stage, and check the current archive log status, as shown in Example 8-74.

Example 8-74 Inquiry archive log mode

```
SQL> alter database open;
Database altered.

SQL> archive log list;
Database log mode          Archive Mode
Automatic archival        Enabled
Archive destination       +DATA/sample/archivelog
Oldest online log sequence 397
Next log sequence to archive 398
Current log sequence      398
```

8. Issue a manual **switch logfile** command, as shown in Example 8-75.

Example 8-75 Trigger manual log file switch

```
SQL> alter system switch logfile;

System altered.
```

9. Shut down the database manually and startup all instances, as shown in Example 8-76.

Example 8-76 Restart Oracle database

```
SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> exit

[oracle@itsovm1 ~]$ srvctl start database -d sample
```

10. Switch to grid user and run **asmcmd** to check whether Redo log is archived during manual switching logfile, as shown in Example 8-77.

Example 8-77 List archived log in ASM

```
[grid@itsovm1 ~]$ asmcmd -p
ASMCMDS [+] > ls +DATA/sample/archivelog
1_398_912970625.dbf
2016_05_28/
```

8.7.2 Add redo log group

This Oracle RAC environment has four database instances. Each database instance has its own redo thread. By default each thread has 2 online redo log groups. It is suggested to add one more redo log group to each thread. To do so, complete these steps:

1. Log in as SYS user, run the following SQL to get current redo log information, as shown in Example 8-78.

Example 8-78 Inquiry redo log files

```
SQL> select a.group#, b.thread#, b.bytes, b.status, a.member from v$logfile a,
v$log b where a.group# = b.group#;
```

GROUP#	THREAD#	BYTES	STATUS	MEMBER
2	1	1073741824	CURRENT	+REDO/SAMPLE/ONLINELOG/group_2.258.912970627
1	1	1073741824	INACTIVE	+REDO/SAMPLE/ONLINELOG/group_1.257.912970625
5	3	1073741824	INACTIVE	+REDO/SAMPLE/ONLINELOG/group_5.259.912970667
6	3	1073741824	CURRENT	+REDO/SAMPLE/ONLINELOG/group_6.260.912970667
3	2	1073741824	INACTIVE	+REDO/SAMPLE/ONLINELOG/group_3.261.912970669
4	2	1073741824	CURRENT	+REDO/SAMPLE/ONLINELOG/group_4.262.912970669
7	4	1073741824	CURRENT	+REDO/SAMPLE/ONLINELOG/group_7.263.912970671
8	4	1073741824	INACTIVE	+REDO/SAMPLE/ONLINELOG/group_8.264.912970671

8 rows selected.

2. Add one more redo log file for each thread, as shown in Example 8-79.

Example 8-79 Add redo log files

```
SQL> alter database add logfile thread 1 group 9 size 1024M;
Database altered.
```

```
SQL> alter database add logfile thread 2 group 10 size 1024M;
Database altered.
```

```
SQL> alter database add logfile thread 3 group 11 size 1024M;
Database altered.
```

```
SQL> alter database add logfile thread 4 group 12 size 1024M;
Database altered.
```

3. Make an inquiry of the latest redo log file information. The newly added redo log groups are in an unused state, as shown in Example 8-80.

Example 8-80 Inquiry redo log files

```
SQL> select a.group#, b.thread#, b.bytes, b.status, a.member from v$logfile a,
v$log b where a.group# = b.group#;
```

GROUP#	THREAD#	BYTES	STATUS	MEMBER
2	1	1073741824	CURRENT	+REDO/SAMPLE/ONLINELOG/group_2.258.912970627
1	1	1073741824	INACTIVE	+REDO/SAMPLE/ONLINELOG/group_1.257.912970625
5	3	1073741824	INACTIVE	+REDO/SAMPLE/ONLINELOG/group_5.259.912970667

```

6      3 1073741824 CURRENT +REDO/SAMPLE/ONLINELOG/group_6.260.912970667
3      2 1073741824 INACTIVE +REDO/SAMPLE/ONLINELOG/group_3.261.912970669
4      2 1073741824 CURRENT +REDO/SAMPLE/ONLINELOG/group_4.262.912970669
7      4 1073741824 CURRENT +REDO/SAMPLE/ONLINELOG/group_7.263.912970671
8      4 1073741824 INACTIVE +REDO/SAMPLE/ONLINELOG/group_8.264.912970671
9      1 1073741824 UNUSED +REDO/SAMPLE/ONLINELOG/group_9.265.913751035
10     2 1073741824 UNUSED +REDO/SAMPLE/ONLINELOG/group_10.266.913751057
11     3 1073741824 UNUSED +REDO/SAMPLE/ONLINELOG/group_11.267.913751089
12     4 1073741824 UNUSED +REDO/SAMPLE/ONLINELOG/group_12.268.913751107

```

12 rows selected.

8.7.3 Multiplex Oracle control files

Control file is an important file in the Oracle database. The default Oracle installation has only one control file. If the only control file is damaged due to a disk failure, the database is shut down. For this reason, each Oracle Database should have at least two control files, each stored on a different disk. The following steps show how to multiplex Oracle control files on different ASM disk groups:

1. Log in as SYS user to query current control file information, as shown in Example 8-81.

Example 8-81 Inquiry control files

```
SQL> select name from v$controlfile;
```

```
NAME
```

```
-----
+REDO/SAMPLE/CONTROLFILE/current.256.912970623
```

2. Shut down the Oracle database to multiplex control files, as shown in Example 8-82.

Example 8-82 Stop Oracle database

```
[oracle@itsovm1 ~]$ srvctl stop listener
[oracle@itsovm1 ~]$ srvctl stop database -d sample
```

3. Use RMAN to copy the current control file to other ASM disk groups, as shown in Example 8-83.

Example 8-83 Copy control files

```
[oracle@itsovm1 ~]$ rman target / nocatalog
```

```
Recovery Manager: Release 12.1.0.2.0 - Production on Sun May 29 12:32:51 2016
```

```
Copyright (c) 1982, 2014, Oracle and/or its affiliates. All rights reserved.
```

```
connected to target database (not started)
```

```
RMAN> startup nomount;
```

```
Oracle instance started
```

```
Total System Global Area 76235669504 bytes
```

```
Fixed Size 7652568 bytes
```

```
Variable Size          14227082024 bytes
Database Buffers      61740154880 bytes
Redo Buffers          260780032 bytes
```

```
RMAN> restore controlfile to '+OCR' from
'+REDO/SAMPLE/CONTROLFILE/current.256.912970623';
```

```
Starting restore at 29-MAY-16
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=197 instance=sample1 device type=DISK
```

```
channel ORA_DISK_1: copied control file copy
Finished restore at 29-MAY-16
```

```
RMAN> restore controlfile to '+DATA' from
'+REDO/SAMPLE/CONTROLFILE/current.256.912970623';
```

```
Starting restore at 29-MAY-16
using channel ORA_DISK_1
```

```
channel ORA_DISK_1: copied control file copy
Finished restore at 29-MAY-16
```

-
4. Log in to the host using the grid user, and run `asmcmd` to check the file names copied by RMAN, as shown in Example 8-84.

Example 8-84 List names of new control files

```
[grid@itsovm1 ~]$ asmcmd
ASMCMD> ls -l +DATA/sample/controlfile/
Type          Redund Striped Time                Sys Name
CONTROLFILE  UNPROT FINE    MAY 29 12:00:00  Y   current.308.913120473
ASMCMD> ls -l +OCR/sample/controlfile/
Type          Redund Striped Time                Sys Name
CONTROLFILE  HIGH   FINE    MAY 29 12:00:00  Y   current.276.913120449
```

5. Log in to the Oracle database as the SYS user to modify the `control_files` parameter, as shown in Example 8-85.

Example 8-85 Modify control_files parameter

```
SQL> conn / as sysdba;
Connected.
SQL> select status from v$instance;
```

```
STATUS
-----
STARTED
```

```
SQL> alter system set
control_files='+REDO/SAMPLE/CONTROLFILE/current.256.912970623','+DATA/sample/co
ntrolfile/current.308.913120473','+OCR/sample/controlfile/current.276.913120449
' scope=spfile;
```

```
System altered.
```

6. Shut down the Oracle database and start it up to normal state, as shown in Example 8-86.

Example 8-86 Restart Oracle database

```
SQL> shutdown immediate;
ORA-01507: database not mounted
ORACLE instance shut down.
SQL> exit
Disconnected from Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 -
64bit Production
With the Partitioning, Real Application Clusters, Automatic Storage Management,
OLAP,
Advanced Analytics and Real Application Testing options
```

```
[oracle@itsovm1 ~]$ srvctl start database -d sample
```

7. Log in to the Oracle database as the SYS user to see the latest control files information, as shown in Example 8-87.

Example 8-87 Inquiry control files

```
SQL> select name from v$controlfile;
```

```
NAME
```

```
-----
+REDO/SAMPLE/CONTROLFILE/current.256.912970623
+DATA/sample/controlfile/current.308.913120473
+OCR/sample/controlfile/current.276.913120449
```

8.7.4 Update processes and sessions parameters

Oracle parameter processes specify the maximum number of operating system user processes that can simultaneously connect to Oracle. Parameter sessions specifies the maximum number of sessions that can be created in the system. For a heavy load Oracle database, update these two parameters to a large number to avoid ORA-12516 or ORA-12520 errors.

To check the current settings of processes and sessions parameter, use the SYS user to log in and run the command shown in Example 8-88.

Example 8-88 Show processes and sessions parameters

```
SQL> conn / as sysdba;
Connected.
SQL> show parameter processes;
```

```
NAME                                TYPE VALUE
-----
aq_tm_processes                      integer 1
db_writer_processes                  integer 7
gcs_server_processes                 integer 3
global_txn_processes                 integer 1
job_queue_processes                  integer 112
log_archive_max_processes             integer 4
processes                             integer 300
```

```
SQL> show parameter sessions;
```

NAME	TYPE	VALUE
java_max_sessionspace_size	integer	0
java_soft_sessionspace_limit	integer	0
license_max_sessions	integer	0
license_sessions_warning	integer	0
sessions	integer	300
shared_server_sessions	integer	

Usually the value of sessions is approximately (1.5 * processes) + 22. Example 8-89 shows adjusting processes to 1000 and sessions to 1522.

Example 8-89 Modify processes and sessions parameters

```
SQL> alter system set processes=1000 scope=spfile;
System altered.
```

```
SQL> alter system set sessions=1522 scope=spfile;
System altered.
```

The new values will be effective until the Oracle Database's next startup. Reboot the Oracle Database to apply the changes as shown in Example 8-90.

Example 8-90 Restart Oracle database

```
[oracle@itsovm1 ~]$ srvctl stop database -d sample
[oracle@itsovm1 ~]$ srvctl start database -d sample
```

8.7.5 Startup and shutdown of Oracle CRS and Database

To shut down Oracle CRS and Database, run the command shown in Example 8-91 with the root user.

Example 8-91 Stop Oracle CRS and database

```
# /u01/app/12.1.0/grid/bin/crsctl stop crs
CRS-2791: Starting shutdown of Oracle High Availability Services-managed resources
on 'itsovm1'
CRS-2673: Attempting to stop 'ora.crsd' on 'itsovm1'
CRS-2790: Starting shutdown of Cluster Ready Services-managed resources on
'itsovm1'
CRS-2673: Attempting to stop 'ora.oc4j' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.mgtdb' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.sample.db' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.LISTENER.lsnr' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.cvu' on 'itsovm1'
CRS-2677: Stop of 'ora.cvu' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.cvu' on 'itsovm3'
CRS-2676: Start of 'ora.cvu' on 'itsovm3' succeeded
CRS-2677: Stop of 'ora.LISTENER.lsnr' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.itsovm1.vip' on 'itsovm1'
CRS-2677: Stop of 'ora.itsovm1.vip' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.itsovm1.vip' on 'itsovm2'
```

CRS-2677: Stop of 'ora.sample.db' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.DATA.dg' on 'itsovm1'
CRS-2677: Stop of 'ora.DATA.dg' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.mgmtdb' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.MGMTLSNR' on 'itsovm1'
CRS-2677: Stop of 'ora.MGMTLSNR' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.MGMTLSNR' on 'itsovm2'
CRS-2676: Start of 'ora.itsovm1.vip' on 'itsovm2' succeeded
CRS-2677: Stop of 'ora.oc4j' on 'itsovm1' succeeded
CRS-2672: Attempting to start 'ora.oc4j' on 'itsovm4'
CRS-2676: Start of 'ora.MGMTLSNR' on 'itsovm2' succeeded
CRS-2672: Attempting to start 'ora.mgmtdb' on 'itsovm2'
CRS-2676: Start of 'ora.oc4j' on 'itsovm4' succeeded
CRS-2676: Start of 'ora.mgmtdb' on 'itsovm2' succeeded
CRS-2673: Attempting to stop 'ora.OCR.dg' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.REDO.dg' on 'itsovm1'
CRS-2677: Stop of 'ora.OCR.dg' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.REDO.dg' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.asm' on 'itsovm1'
CRS-2677: Stop of 'ora.asm' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.ons' on 'itsovm1'
CRS-2677: Stop of 'ora.ons' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.net1.network' on 'itsovm1'
CRS-2677: Stop of 'ora.net1.network' on 'itsovm1' succeeded
CRS-2792: Shutdown of Cluster Ready Services-managed resources on 'itsovm1' has completed
CRS-2677: Stop of 'ora.crsd' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.storage' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.mdnsd' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.gpnpd' on 'itsovm1'
CRS-2677: Stop of 'ora.storage' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.crf' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.ctssd' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.evmd' on 'itsovm1'
CRS-2673: Attempting to stop 'ora.asm' on 'itsovm1'
CRS-2677: Stop of 'ora.gpnpd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.mdnsd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.crf' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.ctssd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.evmd' on 'itsovm1' succeeded
CRS-2677: Stop of 'ora.asm' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.cluster_interconnect.haip' on 'itsovm1'
CRS-2677: Stop of 'ora.cluster_interconnect.haip' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.cssd' on 'itsovm1'
CRS-2677: Stop of 'ora.cssd' on 'itsovm1' succeeded
CRS-2673: Attempting to stop 'ora.gipcd' on 'itsovm1'
CRS-2677: Stop of 'ora.gipcd' on 'itsovm1' succeeded
CRS-2793: Shutdown of Oracle High Availability Services-managed resources on 'itsovm1' has completed
CRS-4133: Oracle High Availability Services has been stopped.

To start Oracle CRS and Database, run the command shown in Example 8-92 with the root user.

Example 8-92 Start Oracle CRS and database

```
# /u01/app/12.1.0/grid/bin/crsctl start crs
```

To list Oracle CRS resource status, or check CRS startup progress, run the command shown in Example 8-93 with the grid user.

Example 8-93 Show CRS resources states

```
[grid@itsovm1 ~]$ crsctl stat res -t
```

Name	Target	State	Server	State details

Local Resources				

ora.DATA.dg	ONLINE	ONLINE	itsovm1	STABLE
	ONLINE	ONLINE	itsovm2	STABLE
	ONLINE	ONLINE	itsovm3	STABLE
	ONLINE	ONLINE	itsovm4	STABLE
ora.LISTENER.lsnr	ONLINE	ONLINE	itsovm1	STABLE
	ONLINE	ONLINE	itsovm2	STABLE
	ONLINE	ONLINE	itsovm3	STABLE
	ONLINE	ONLINE	itsovm4	STABLE
ora.OCR.dg	ONLINE	ONLINE	itsovm1	STABLE
	ONLINE	ONLINE	itsovm2	STABLE
	ONLINE	ONLINE	itsovm3	STABLE
	ONLINE	ONLINE	itsovm4	STABLE
ora.REDO.dg	ONLINE	ONLINE	itsovm1	STABLE
	ONLINE	ONLINE	itsovm2	STABLE
	ONLINE	ONLINE	itsovm3	STABLE
	ONLINE	ONLINE	itsovm4	STABLE
ora.asm	ONLINE	ONLINE	itsovm1	Started,STABLE
	ONLINE	ONLINE	itsovm2	Started,STABLE
	ONLINE	ONLINE	itsovm3	Started,STABLE
	ONLINE	ONLINE	itsovm4	Started,STABLE
ora.net1.network	ONLINE	ONLINE	itsovm1	STABLE
	ONLINE	ONLINE	itsovm2	STABLE
	ONLINE	ONLINE	itsovm3	STABLE
	ONLINE	ONLINE	itsovm4	STABLE
ora.ons	ONLINE	ONLINE	itsovm1	STABLE
	ONLINE	ONLINE	itsovm2	STABLE
	ONLINE	ONLINE	itsovm3	STABLE
	ONLINE	ONLINE	itsovm4	STABLE

Cluster Resources				

ora.LISTENER_SCAN1.lsnr	1	ONLINE	ONLINE	itsovm1	STABLE
ora.LISTENER_SCAN2.lsnr	1	ONLINE	ONLINE	itsovm2	STABLE
ora.LISTENER_SCAN3.lsnr	1	ONLINE	ONLINE	itsovm4	STABLE
ora.MGMTLSNR	1	ONLINE	ONLINE	itsovm2	169.254.142.90 10.0.0.2,STABLE
ora.cvu	1	ONLINE	ONLINE	itsovm2	STABLE
ora.itsovm1.vip	1	ONLINE	ONLINE	itsovm1	STABLE
ora.itsovm2.vip	1	ONLINE	ONLINE	itsovm2	STABLE
ora.itsovm3.vip	1	ONLINE	ONLINE	itsovm3	STABLE
ora.itsovm4.vip	1	ONLINE	ONLINE	itsovm4	STABLE
ora.mgmtdb	1	ONLINE	ONLINE	itsovm2	Open,STABLE
ora.oc4j	1	ONLINE	ONLINE	itsovm2	STABLE
ora.sample.db	1	ONLINE	ONLINE	itsovm1	Open,STABLE
	2	ONLINE	ONLINE	itsovm2	Open,STABLE
	3	ONLINE	ONLINE	itsovm3	Open,STABLE
	4	ONLINE	ONLINE	itsovm4	Open,STABLE
ora.scan1.vip	1	ONLINE	ONLINE	itsovm1	STABLE
ora.scan2.vip	1	ONLINE	ONLINE	itsovm2	STABLE
ora.scan3.vip	1	ONLINE	ONLINE	itsovm4	STABLE

It is also possible to start or stop Oracle database only, and leave CRS running. First log in using the oracle user to get the database name, as shown in Example 8-94.

Example 8-94 Get database name

```
[oracle@itsovm1 ~]$ srvctl config database
sample
```

To stop a specific Oracle database, run the following command using the oracle user, as shown in Example 8-95.

Example 8-95 Stop Oracle database

```
[oracle@itsovm1 ~]$ srvctl stop database -d sample
```

To start a specific Oracle database, run the command shown in Example 8-96 using the oracle user.

Example 8-96 Start Oracle database

```
[oracle@itsovm1 ~]$ srvctl start database -d sample
```

And it is also possible to start or stop one Oracle database instance, and leave the rest of instances running. First log in using the oracle user to get the instance names as shown in Example 8-97.

Example 8-97 Show database configuration

```
[oracle@itsovm1 ~]$ srvctl config database -d sample -a
Database unique name: sample
Database name: sample
Oracle home: /u01/app/oracle/product/12.1.0/dbhome
Oracle user: oracle
Spfile: +DATA/SAMPLE/PARAMETERFILE/spfile.265.912970673
Password file: +DATA/SAMPLE/PASSWORD/pwdsample.256.912970459
Domain:
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: DATA,REDO,OCR
Mount point paths:
Services:
Type: RAC
Start concurrency:
Stop concurrency:
Database is enabled
Database is individually enabled on nodes:
Database is individually disabled on nodes:
OSDBA group: dba
OSOPER group: oper
Database instances: sample1,sample2,sample3,sample4
Configured nodes: itsovm1,itsovm2,itsovm3,itsovm4
Database is administrator managed
```

To stop a specific Oracle database instance, run the command shown in Example 8-98 using the oracle user.

Example 8-98 Stop a database instance

```
[oracle@itsovm1 ~]$ srvctl stop instance -d sample -i sample1
```

To start a specific Oracle database instance, run the command shown in Example 8-99 using the oracle user.

Example 8-99 Start a database instance

```
[oracle@itsovm1 ~]$ srvctl start instance -d sample -i sample1
```

8.7.6 OCR and voting disk management

To verify the integrity of OCR, run the command shown in Example 8-100 using the grid user.

Example 8-100 Perform OCR check

```
[grid@itsovm1 ~]$ ocrcheck
Status of Oracle Cluster Registry is as follows :
  Version                :                4
  Total space (kbytes)   :            409568
  Used space (kbytes)    :              1768
  Available space (kbytes) :          407800
  ID                     :    2057472910
  Device/File Name      :                +OCR
                        Device/File integrity check succeeded
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured
                        Device/File not configured

Cluster registry integrity check succeeded

Logical corruption check bypassed due to non-privileged user
```

To list current voting disks, run the command shown in Example 8-101 using the grid user.

Example 8-101 List voting disks

```
[grid@itsovm1 ~]$ crsctl query css votedisk
## STATE      File Universal Id                File Name Disk group
--  -
  1. ONLINE   aa4865c2ae924f4cbf95862ad4985605 (/dev/oracleasm/disks/OCR1) [OCR]
  2. ONLINE   95b6d848bb6f4f4fbff7471f0d803159 (/dev/oracleasm/disks/OCR2) [OCR]
  3. ONLINE   a0527f228c404f9fbfb5ef3642fa67e4 (/dev/oracleasm/disks/OCR3) [OCR]
Located 3 voting disk(s).
```



Using IBM Spectrum Protect

This chapter covers the steps necessary to use IBM Spectrum Protect to ensure that you have a current backup of the Oracle RAC database.

It is beyond the scope of this book to show how to install IBM Spectrum Protect. This chapter provide information, where applicable, that shows where to obtain the code, documentation, and scripts necessary to install IBM Spectrum Protect. Contact your IBM representative if you need any assistance in preparing for, installing, or configuring IBM Spectrum Protect.

This chapter includes the following sections:

- ▶ Environment overview
- ▶ IBM Spectrum Protect server deployment
- ▶ DP for Oracle: Install and configure IBM Spectrum Protect on Oracle server node

9.1 Environment overview

This section details the IBM Spectrum Protect server components that were deployed on the example Oracle DB server on VersaStack environment.

IBM Spectrum Protect includes these core components:

- ▶ IBM Spectrum Protect backup server
- ▶ Spectrum Operations Center
- ▶ IBM Spectrum Protect Data Protection for Databases: Oracle

9.1.1 IBM Spectrum Protect Backup Server

IBM Spectrum Protect is a highly scalable backup solution that can be deployed on multiple hardware and software platforms. For a list of Tivoli Storage Manager Supported Operating Systems, refer to the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21243309#Server%20Table>

Within the Oracle on VersaStack setup, we deployed IBM Spectrum Protect V7.1.6 on Red Hat Enterprise Linux Server (RHEL) V7.2. This example uses a host name of TSM716.

Linux x86_64 Server Requirements and Support for running IBM Spectrum Protect on Linux x86_64 can be found at the following website:

http://www.ibm.com/support/docview.wss?rs=663&context=SSGSG7&q1=ServerRequirements&uid=swg21204361&loc=en_US&cs=utf-8&lang=en

9.1.2 IBM Spectrum Protect Blueprints

The blueprint consists of a document, or “cookbook”, that describes the three reference architectures in detail, including IBM hardware model numbers and configuration requirements. It also includes scripts to speed up the installation and configuration, increasing time-to-value. The storage preparation script automates preparation of the file systems that are used by the Tivoli Storage Manager server. The blueprint configuration script verifies that the hardware configuration meets the blueprint specifications, validates kernel settings on Linux systems, and verifies the configuration of required file systems before running the standard Tivoli Storage Manager server installation. The script also configures the Tivoli Storage Manager server using best practices and performs these tasks:

- ▶ Creates a DB2 instance
- ▶ Defines data deduplication storage pools with optimal performance settings
- ▶ Defines administrative maintenance tasks that are optimized for data deduplication scalability
- ▶ Defines Tivoli Storage Manager database backup to disk
- ▶ Creates a dsmserv.opt file with best practice option overrides
- ▶ Creates policy domains for database, mail and file servers with management classes for 30, 60, and 120-day retention
- ▶ Defines backup schedules for all client types that can be easily selected when deploying the wanted client workloads

The workload simulation script runs simulated Tivoli Storage Manager database and storage pool workloads. It provides performance measurements that can be used to compare as a reference against those measured on the blueprint configuration.

More information can be found at the following website:

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/IBM%20Spectrum%20Protect%20Blueprints>

9.1.3 IBM Spectrum Protect Operations Center

The IBM Spectrum Protect Operations Center is a light-weight management application that offers the daily dashboard and management interface for the IBM Spectrum Protect servers. It can be deployed on the same system hosting the primary IBM Spectrum Protect server or on, for example, the VM that also hosts the IBM Tivoli Monitoring for IBM Spectrum Protect server.

Tivoli Storage Manager Operations Center Software and Hardware Requirements can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21653418>

The example deployment installed the Operations Center on the IBM Spectrum Protect server.

9.1.4 IBM Spectrum Protect Data Protection for Databases: Oracle

IBM Spectrum Protect for Databases (formerly IBM Tivoli Storage Manager for Databases) helps protect Oracle and Microsoft SQL data no matter where it is stored. You can continue running primary applications on your database servers while they back up, and restore data to and from auxiliary storage by using automated tasks, utilities, and interfaces. This software performs online, consistent, and centralized backups to help you avoid downtime, protect vital enterprise data, and minimize operational costs.

IBM Spectrum Protect Data Protection for Databases: Oracle on UNIX, IBM AIX®, and Linux hardware and software requirements and installation prerequisites can be found at the following website:

http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.3/db.orc/r_dporc_inst_prereq.html

9.2 IBM Spectrum Protect server deployment

The IBM Spectrum Protect server deployment was performed in two major steps. First, the code was installed, and then the ISP server was configured.

Download the relevant software through IBM Passport Advantage® and Passport Advantage Express available here:

<https://www.ibm.com/software/passportadvantage/>

If you have problems accessing this, contact your IBM sales representative.

9.2.1 Code installation: Step 1

Two packages were installed on the example Red Hat system. The first package was the Client code, referred to as Backup/Archive Client (BAC). The package installed by using rpm was:

7.1.6.0-TIV-TSMBAC-LinuxX86.tar

For detailed installation instructions, see *Installing the Tivoli Storage Manager backup-archive clients* in IBM Knowledge Center available at:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.6/client/c_inst.html

The second package was the Server code. The package installed using the installation script (install.sh) was:

7.1.6.000-TIV-TSMSRV-Linuxx86_64.bin

For detailed installation instructions, see *Installing and upgrading the server* in IBM Knowledge Center available at:

http://www.ibm.com/support/knowledgecenter/en/SSGSG7_7.1.6/srv.common/t_installing_srv.html?view=kc

Both packages were installed by using the instructions in that documentation.

These links are updated as changes are made to the code, and contain any relevant information.

9.2.2 Server configuration using automated Blueprint script: Step 2

This step contains two substeps: Linux disk mapping and configuring the ISP server using the IBM Blueprint automation script.

The example disk was configured to support the TSM server, DB2 space – database, log, and archive and database backup disk space and disk space for the backup data (tsmstg01-07). Our disk configuration is shown in Figure 9-1.

```
[root@TSM716 ~]# df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/rhel-root      50G       8.1G   42G  17% /
devtmpfs                  126G       0    126G   0% /dev
tmpfs                     126G     160K   126G   1% /dev/shm
tmpfs                     126G     51M   126G   1% /run
tmpfs                     126G       0    126G   0% /sys/fs/cgroup
/dev/sdbel                494M     158M   337M  32% /boot
/dev/mapper/rhel-home     58G       11G    47G  19% /home
tmpfs                     26G       0     26G   0% /run/user/0
tmpfs                     26G     40K     26G   1% /run/user/1000
/dev/mapper/tsmdb512-tsmdb512a 512G     785M   501G   1% /tsmdb512a
/dev/mapper/tsmdb512-tsmdb512b 512G     785M   501G   1% /tsmdb512b
/dev/mapper/tsmdb512-tsmdb512c 512G     785M   501G   1% /tsmdb512c
/dev/mapper/tsmdb512-tsmdb512d 512G     785M   501G   1% /tsmdb512d
/dev/mapper/tsmarch-tsmarch01 512G     6.4G   496G   2% /tsmarch01
/dev/mapper/tsmdbbk-tsmdbbk01 504G     1.5G   477G   1% /tsmdbbk01
/dev/mapper/tsmdbbk-tsmdbbk02 504G     1.7G   477G   1% /tsmdbbk02
/dev/mapper/tsmdbbk-tsmdbbk03 504G     2.1G   477G   1% /tsmdbbk03
/dev/mapper/tsmdbbk-tsmdbbk04 504G     1.3G   477G   1% /tsmdbbk04
/dev/mapper/tsmstg-tsmstg01 3.4T     221G   3.0T   7% /tsmstg01
/dev/mapper/tsmstg-tsmstg02 1.5T     197G   1.2T  14% /tsmstg02
/dev/mapper/tsmstg-tsmstg03 2.5T     212G   2.2T   9% /tsmstg03
/dev/mapper/tsmstg-tsmstg04 3.4T     211G   3.1T   7% /tsmstg04
/dev/mapper/tsmstg-tsmstg05 3.9T     202G   3.6T   6% /tsmstg05
/dev/mapper/tsmstg-tsmstg06 504G     201G   278G  42% /tsmstg06
/dev/mapper/tsmstg-tsmstg07 504G     183G   296G  39% /tsmstg07
/dev/mapper/tsmdb512-tsmlog00 256G     129G   123G  52% /tsmact01
[root@TSM716 ~]#
```

Figure 9-1 Mappings of the disk configuration

ISP server configuration

To download the documentation and the scripts, navigate to *Links to the version 2.3 blueprint downloads*, available at:

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/IBM%20Spectrum%20Protect%20Blueprints>

The server was configured by using the Blueprint Perl script (TSMserverconfig.pl) using the response file for a small server as defined in the IBM Blueprint documentation (*Blueprint for Linux x86*) `srv_blueprint_xlinux_v23.pdf`.

Table 19 on page 45 of that document details the parameters and values in the response file.

Figure 9-2 contains the contents of the response file that is used for this deployment. We followed the directions contained in the Blueprint documentation to achieve this configuration.

```
[root@TSM716 tsmconfig]# cat responsefilesmall.txt
serverscale s
db2user tsminst1
db2userpw Object00
db2group tsmsrvrs
db2userhomedir /home/tsminst1
instdirmountpoint /home/tsminst1/tsminst1
dbdirpaths /tsmdb512a,/tsmdb512b,/tsmdb512c,/tsmdb512d
tsmstgpaths /tsmstg01,/tsmstg02,/tsmstg03,/tsmstg04,/tsmstg05,/tsmstg06,/tsmstg07
actlogpath /tsmact01
archlogpath /tsmarch01
dbbackdirpaths /tsmdbbk01,/tsmdbbk02,/tsmdbbk03,/tsmdbbk04
backupstarttime 22:00
tsmsysadminid admin1
tsmsysadminpw Object00
tcpport 1500
servername TSM716
serverpassword Object00
[root@TSM716 tsmconfig]#
```

Figure 9-2 Response file

The configuration was started with the following command, which was entered on the Linux CLI:

```
Perl TSMserverconfig.pl responsefilesmall.txt
```

Disk full error: During the TSM server configuration, an issue was encountered between IBM DB2, the TSM implementation, and the 4096 allocation unit size (segment size) disks in the environment. During initial formatting of the DB2 database used by TSM to track backup data, the formatting process failed due to the use of 4096 block size disks for the IBM DB2 database files (data volumes and log volume).

The following message was found in the IBM DB2 logs:

```
FUNCTION: DB2 UDB, oper system services, sqlseekwrite64, probe:20
MESSAGE : ZRC=0x850F000C=-2062614516=SQL0_DISK "Disk full."
```

This problem is called *Database creation is failing with "Disk full" error during DMS tablespace creation on a storage subsystem with a sector size other than 512 bytes*. It is documented here:

<http://www.ibm.com/support/docview.wss?uid=swg21449214>

Our solution was to acquire four disks that use a 512 segment size and allocated the DB2 mounts to these four disks.

The five DB2 database mount points were mapped to the four disks with a 512 segment size as follows:

```
/dev/mapper/tsmdb512-tsmdb512a /tsmdb512a          ext4      defaults    0 0
/dev/mapper/tsmdb512-tsmdb512b /tsmdb512b          ext4      defaults    0 0
/dev/mapper/tsmdb512-tsmdb512c /tsmdb512c          ext4      defaults    0 0
/dev/mapper/tsmdb512-tsmdb512d /tsmdb512d          ext4      defaults    0 0
/dev/mapper/tsmdb512-tsmlog00 /tsmact01           ext4      defaults    0 0
```

ISP server initial testing

As a quick test after the successful completion of the TSM server configuration Perl script, start a CLI session by using the `dsmadm` command found in `/opt/tivoli/tsm/client/ba/bin` as shown in Figure 9-3.

```
[root@TSM716 bin]# pwd
/opt/tivoli/tsm/client/ba/bin
[root@TSM716 bin]# dsmadm
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 7, Release 1, Level 6.0
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

Enter your user id:  admin1

Enter your password:

Session established with server TSM716: Linux/x86_64
  Server Version 7, Release 1, Level 6.0
  Server date/time: 07/14/2016 08:25:37  Last access: 07/14/2016 08:16:20

tsm: TSM716>
```

Figure 9-3 The `dsmadm` command

9.3 DP for Oracle: Install and configure IBM Spectrum Protect on Oracle server node

This section covers the installation and configuration of IBM Spectrum Protect on the Oracle server node.

9.3.1 BAC Installation

Install the same BAC code package on the Oracle server as you did on the TSM server because the platforms are the same. This package is `7.1.6.0-TIV-TSMBAC-LinuxX86.tar`, per the product documentation.

For detailed installation instructions, see *Installing the Tivoli Storage Manager backup-archive clients* in IBM Knowledge Center available at:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.6/client/c_inst.html

BAC Configuration

After installing the Client code, create a `dsm.sys` and `dsm.opt` file to reflect your TSM implementation.

Figure 9-4 shows the example BAC dsm.sys file.

```
/opt/tivoli/tsm/client/ba/bin
[root@itsovml bin]# cat dsm.sys
*****
* Tivoli Storage Manager                                     *
*                                                         *
* Sample Client System Options file for UNIX (dsm.sys.smp) *
*****

* This file contains the minimum options required to get started
* using TSM. Copy dsm.sys.smp to dsm.sys. In the dsm.sys file,
* enter the appropriate values for each option listed below and
* remove the leading asterisk (*) for each one.

* If your client node communicates with multiple TSM servers, be
* sure to add a stanza, beginning with the SERVERNAME option, for
* each additional server.

*****

SErvername TSM716
  COMMMethod      TCPip
  TCPPort         1500
  TCPServeraddress 192.168.161.45
  nodename        ITSOVML
  passwordaccess  generate
```

Figure 9-4 Example BAC dsm.sys file

BAC dsm.opt is shown in Figure 9-5.

```
[root@itsovml bin64]# cat dsm.opt
SErvername tdpo
[root@itsovml bin64]# cd ../../ba/bin
[root@itsovml bin]# cat dsm.opt
*****
* Tivoli Storage Manager                                     *
*                                                         *
* Sample Client User Options file for UNIX (dsm.opt.smp) *
*****

* This file contains an option you can use to specify the TSM
* server to contact if more than one is defined in your client
* system options file (dsm.sys). Copy dsm.opt.smp to dsm.opt.
* If you enter a server name for the option below, remove the
* leading asterisk (*).

*****

* SErvername          A server name defined in the dsm.sys file
SErvername           TSM716
```

Figure 9-5 Example BAC dsm.opt file

BAC initial test

To test the installation and configuration of the base client code, perform a simple file level test backup as shown in Figure 9-6.

```
[root@itsovm1 bin]# pwd
/opt/tivoli/tsm/client/ba/bin
[root@itsovm1 bin]# dsmc incr /opt/tivoli/tsm/client/ba/bin*.*
IBM Tivoli Storage Manager
Command Line Backup-Archive Client Interface
Client Version 7, Release 1, Level 6.0
Client date/time: 07/14/2016 08:17:15
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

Node Name: ITSOVM1
Session established with server TSM716: Linux/x86_64
Server Version 7, Release 1, Level 6.0
Server date/time: 07/14/2016 08:16:50 Last access: 07/14/2016 07:25:10

Incremental backup of volume '/opt/tivoli/tsm/client/ba/bin*.*'
Successful incremental backup of '/opt/tivoli/tsm/client/ba/bin*.*'

Total number of objects inspected:          15
Total number of objects backed up:          0
Total number of objects updated:           0
Total number of objects rebound:          0
Total number of objects deleted:           0
Total number of objects expired:           0
Total number of objects failed:            0
Total number of objects encrypted:         0
Total number of objects grew:              0
Total number of retries:                   0
Total number of bytes inspected:           70.43 KB
Total number of bytes transferred:         0 B
Data transfer time:                         0.00 sec
Network data transfer rate:                 0.00 KB/sec
Aggregate data transfer rate:               0.00 KB/sec
Objects compressed by:                     0%
Total data reduction ratio:                 100.00%
Elapsed processing time:                    00:00:01
```

Figure 9-6 Initial test

9.3.2 Install and configure Data Protection (DP) for Databases (TDPO)

Now that you have a functioning TSM environment, you need to include the existing Oracle server into the environment.

The goal is to use TDPO to send a backup to the TSM server. After you install TDPO and configure your environment, an RMAN backup will, with the addition of just one parameter, automatically forward the backup to TSM for safe-keeping.

The example Oracle server environment has these characteristics:

- ▶ Oracle server information
 - Hostname: itsovm1
 - Ipaddr: 192.168.161.71
 - Oracle version: 12c – 12.1.0

► Oracle database information:

```
ORACLE_SID=sample1
ORACLE_BASE=/u01/app/oracle
ORACLE_HOME=/u01/app/oracle/product/12.1.0/dbhome
```

DP for Oracle Install package

Use the DP for Oracle product documentation available at:

http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.3/db.orc/t_protect_dpdborc.html?view=embed

Specifically use this file that clients and Business Partners can obtain by using Passport Advantage:

ISP_DB_716_DP_ORACLE_LINUX_MP_ML.tar.gz

Note: DP for Oracle was not changed in v7.1.6 and at the time of writing the most current version information is v7.1.3.

Install this package by following the product documentation referenced previously.

DP for Oracle configuration

The installation creates a new directory (oracle) under /opt/tivoli/tsm/client/

To configure the TDPO, make changes to files in the TSM client API and Oracle directories as shown in Example 9-1.

Example 9-1 Directories (api and oracle)

```
[root@itsovm1 client]# pwd
/opt/tivoli/tsm/client
[root@itsovm1 client]# ls -l
total 12
drwxr-xr-x. 4 root bin 4096 Jul 12 09:29 api
drwxr-xr-x. 4 root bin 4096 Jul 12 09:30 ba
drwxrwxr-x. 6 root root 4096 Jul 12 14:07 oracle
```

9.3.3 TDPO.opt

There is a bin64 subdirectory under oracle that contains a sample tdpo.opt (tdpo.opt.smp). Modify this file for your environment as shown Figure 9-7.

```
[root@itsovm1 bin64]# pwd
/opt/tivoli/tsm/client/oracle/bin64
[root@itsovm1 bin64]# cat tdpo.opt
*****
* IBM Tivoli Storage Manager for Databases
*   Data Protection for Oracle
*
* Sample tdpo.opt for the LinuxAMD64 Data Protection for Oracle
*****

DSMI_ORC_CONFIG      /opt/tivoli/tsm/client/oracle/bin64/dsm.opt
*DSMI_LOG            /opt/tivoli/tsm/client/oracle/bin64
DSMI_LOG             /home/oracle

*TDPO_FS             /adsmorc
*TDPO_NODE           <hostname>
*TDPO_OWNER          <username>
*TDPO_PSWDPATH       /opt/tivoli/tsm/client/oracle/bin64
*TDPO_PSWDPATH       /home/oracle

*TDPO_DATE_FMT       1
*TDPO_NUM_FMT        1
*TDPO_TIME_FMT       1

*TDPO_MGMT_CLASS_2   mgmtclass2
*TDPO_MGMT_CLASS_3   mgmtclass3
*TDPO_MGMT_CLASS_4   mgmtclass4
```

Figure 9-7 Example TDPO.opt file

9.3.4 TDPO dsm.opt

Create a dsm.opt following the product documentation in the same directory as shown in Figure 9-8.

```
[root@itsovm1 bin64]# pwd
/opt/tivoli/tsm/client/oracle/bin64
[root@itsovm1 bin64]# cat dsm.opt
SErvername tdpo
```

Figure 9-8 Example TDPO dsm.opt file

9.3.5 BAC API dsm.sys

The final configuration step is to update the client API dsm.sys file to conform to your TSM implementation.

Placing the `SErvername tdpo` stanza before `SErvername TSM716` stanza resolves issues with the `tdpofconf password` command later in the configuration. The completed `api/bin64/dsm.sys` file is shown in Figure 9-9.

```
[root@itsovm1 bin64]# pwd
/opt/tivoli/tsm/client/api/bin64
[root@itsovm1 bin64]# cat dsm.sys
*****
* Tivoli Storage Manager                                     *
*                                                         *
* Sample Client System Options file for UNIX (dsm.sys.smp) *
*****

* This file contains the minimum options required to get started
* using TSM. Copy dsm.sys.smp to dsm.sys. In the dsm.sys file,
* enter the appropriate values for each option listed below and
* remove the leading asterisk (*) for each one.

* If your client node communicates with multiple TSM servers, be
* sure to add a stanza, beginning with the SERVERNAME option, for
* each additional server.

*****

SErvername tdpo
  COMMethod      TCPip
  TCPPort        1500
  TCPServeraddress 192.168.161.45
  passwordaccess generate
  nodename       itsovm1_ora
  passworddir    /home/oracle

SErvername TSM716
  COMMethod      TCPip
  TCPPort        1500
  TCPServeraddress 192.168.161.45
  passwordaccess generate
  nodename       itsovm1
```

Figure 9-9 Completed BAC API `dsm.sys` file

Set the TSM node password for TDPO by running the `tdpofconf` command with `/opt/tivoli/tsm/client/oracle/bin64/tdpofconf` as the oracle user.

Set the TSM node password for TDPO to enable authenticated communications between TDPO process on the Oracle server and the TSM server.

Create a node on the TSM server named `itsovm1_ora` with the password `Object00` before you issue the `tdpofconf` command.

The command `tdpofconf password` creates an encrypted `TSM.PWD` password file that allows this authentication as shown in Example 9-2.

Example 9-2 Password file

```
[oracle@itsovm1 ~]$ whoami
Oracle
[oracle@itsovm1 ~]$ env | grep ORA
```

```

ORACLE_SID=sample1
ORACLE_BASE=/u01/app/oracle
ORACLE_HOME=/u01/app/oracle/product/12.1.0/dbhome
[oracle@itsovm1 bin64]$ pwd
/opt/tivoli/tsm/client/oracle/bin64
[oracle@itsovm1 ~]$ ./tdpoconf password
IBM Tivoli Storage Manager for Databases:
Data Protection for Oracle
Version 7, Release 1, Level 3.0
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.
*****
*   IBM Tivoli Storage Manager for Databases Utility
*   Password file initialization/update program
*****
Please enter current password:
Please enter new password:
Please reenter new password for verification:
ANU0260I Password successfully changed.

```

The TSM.PWD file is now updated/created in the directory indicated in the configuration and the command `ls -l` shown in Example 9-3 shows where this file was created.

Example 9-3 Created pwd file

```

[oracle@itsovm1 bin64]$ ls -l /home/oracle/TSM.PWD
-rw----- . 1 oracle oinstall 132 Jul 14 09:11 /home/oracle/TSM.PWD

```

The TSM server has two nodes defined. The first (ITSOVM1) is used for standard file backup and restore, and the second (ITSOVM1_ORA) is used for Oracle backup and restore operations. Figure 9-10 shows the `q node` command listing these two nodes.

```

tsm: TSM716>q node
Session established with server TSM716: Linux/x86_64
  Server Version 7, Release 1, Level 6.0
  Server date/time: 07/14/2016 08:59:39  Last access: 07/14/2016 08:46:22

```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
ITSOVM1	Linux x86-64	STANDARD	<1	1	No
ITSOVM1_ORA	TDPO Linux8- 6-64	STANDARD	1	2	No

```

tsm: TSM716>

```

Figure 9-10 Query node command

9.3.6 Use RMAN to back up the Oracle Database using TDPO

The use of TDPO requires minor changes to the processes, procedures, and scripts developed by the Oracle DBAs to protect their systems. Just one parameter is added to direct the backup or restore to use TSM. TDPO is fully supported with or without the use to an RMAN catalog database.

This is the only operational change that is needed to use TDPO.

An example RMAN backup of the sample database is shown in Example 9-4 using the `rman` command. This command is usually used in a script. Note the addition of the TDPO parameter pointer to TDPO configuration that was established earlier.

Example 9-4 RMAN backup

```
[oracle@itsovm1 ~]$ rman target /
RMAN> run
2> {
3> allocate channel t1 type 'sbt_tape' parms
4> 'ENV=(TDPO_OPTFILE=/opt/tivoli/tsm/client/oracle/bin64/tdpo.opt)';
5> backup filesperset 5
6> format 'df_%t_%s_%p'
7> (database);
8> }
```

```
using target database control file instead of recovery catalog
allocated channel: t1
channel t1: SID=2812 instance=sample1 device type=SBT_TAPE
channel t1: Data Protection for Oracle: version 7.1.3.0
```

```
Starting backup at 13-JUL-16
channel t1: starting full datafile backup set
channel t1: specifying datafile(s) in backup set
input datafile file number=00008 name=+DATA/sample/soe.dbf
input datafile file number=00001 name=+DATA/SAMPLE/DATAFILE/system.258.912970523
input datafile file number=00006 name=+DATA/SAMPLE/DATAFILE/users.259.912970569
channel t1: starting piece 1 at 13-JUL-16
channel t1: finished piece 1 at 13-JUL-16
piece handle=df_917091597_4_1 tag=TAG20160713T113957 comment=API Version 2.0,MMS
Version 7.1.3.0
channel t1: backup set complete, elapsed time: 07:49:07
channel t1: starting full datafile backup set
channel t1: specifying datafile(s) in backup set
input datafile file number=00004 name=+DATA/SAMPLE/DATAFILE/undotbs1.260.9129705
69
input datafile file number=00003 name=+DATA/SAMPLE/DATAFILE/sysaux.257.912970481
input datafile file number=00002 name=+DATA/SAMPLE/DATAFILE/undotbs3.263.9129706
43
input datafile file number=00005 name=+DATA/SAMPLE/DATAFILE/undotbs2.262.9129706
43
input datafile file number=00007 name=+DATA/SAMPLE/DATAFILE/undotbs4.264.9129706
43
channel t1: starting piece 1 at 13-JUL-16
channel t1: finished piece 1 at 13-JUL-16
piece handle=df_917119744_5_1 tag=TAG20160713T113957 comment=API Version 2.0,MMS
Version 7.1.3.0
channel t1: backup set complete, elapsed time: 00:01:55
channel t1: starting full datafile backup set
channel t1: specifying datafile(s) in backup set
including current control file in backup set
including current SPFILE in backup set
channel t1: starting piece 1 at 13-JUL-16
channel t1: finished piece 1 at 13-JUL-16
```

```
piecehandle=df_917119859_6_1 tag=TAG20160713T113957 comment=API Version 2.0,MMS
Version 7.1.3.0
channel t1: backup set complete, elapsed time: 00:00:01
Finished backup at 13-JUL-16
released channel: t1
```

9.3.7 IBM Spectrum Protect activity log for Oracle database backup

Example 9-5 shows the TSM activity log from the third full backup (see Example 9-4 on page 226) of the SAMPLE database. Note that the total database size is nearly 1.6 TB.

Example 9-5 TSM activity log

```
07/13/2016 19:28:32 ANR0951I Session 242 for node ITS0VM1_ORA processed 1
files by using inline data deduplication or compression,
or both. The number of original bytes was
1,580,024,201,216. Inline data deduplication reduced the
data by 636,189,798 bytes and inline compression reduced
the data by 828,340,493,675 bytes. (SESSION: 242)
```

9.3.8 View of backup data from IBM Spectrum Protect Operations Center

Using TSM as the backup repository (Container pool) that has inline data deduplication and compression causes a 54% reduction in the amount to data stored. Figure 9-11 shows the deduppool.




Figure 9-11 The deduppool

9.3.9 View of Oracle server node from IBM Spectrum Protect Operations Center

Figure 9-12 shows the TSM file space for the example ITSOVM1_ORA Oracle server node. It shows the amount of data and files that are currently held for this node.




Figure 9-12 Oracle server node

9.3.10 TSM CLI query occ command after backups

Example 9-6 shows the TSM CLI `query occ` command after backups. This is another way of displaying the information that is shown in Figure 9-12.

Example 9-6 Output of query occ command

```
tsm: TSM716>q occ
```

Node Name	Type	Filespace Name	FSID	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
ITSOVM1	Bkup	/	1	DEDUPPOOL	49	-	3.93
ITSOVM1	Bkup	/run	2	DEDUPPOOL	784	-	152.37
ITSOVM1_ORA	Bkup	/adsmorc	1	DEDUPPOOL	459	-	4,546,479.0
							0

Figure 9-13 displays the summary information for the Oracle server TSM node, ITSOVM1_ORA, using the IBM Spectrum Protect Operations Center. This summary shows the recent backup activity for the Oracle database.




Figure 9-13 Summary information for ITSOVM1_ORA

The following examples show configuration details for the IBM Spectrum Protect server, including this information:

- ▶ General OS and processor details
- ▶ File system table (fstab)

Example 9-7 shows the example TSM716 server OS information.

Example 9-7 TSM716 server OS information

```

1nmonq14gqqqqq[H for help]qqqHostname=itsovm1qqqqqqRefresh= 1secs qq08:48.22qqqqqqqqqqqqqqqqqqqq
x Linux and Processor Details qqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqq
x Linux: Linux version 3.10.0-327.el7.x86_64 (mockbuild@x86-034.build.eng.bos.redhat.com) x
x Build: (gcc version 4.8.3 20140911 (Red Hat 4.8.3-9) (GCC) ) x
x Release : 3.10.0-327.el7.x86_64 x
x Version : #1 SMP Thu Oct 29 17:29:29 EDT 2015 x
x cpuinfo: model name : Intel(R) Xeon(R) CPU E5-2660 v4 @ 2.00GHz x
x cpuinfo: vendor_id : GenuineIntel x
x cpuinfo: microcode : 0xb000014 x
x cpuinfo: cpuid level : 13 x
x # of CPUs: 56 x
x Machine : x86_64 x
x Nodename : itsovm1 x
x /etc/*ease[1]: NAME="Red Hat Enterprise Linux Server" x
x /etc/*ease[2]: VERSION="7.2 (Maipo)" x
x /etc/*ease[3]: ID="rhel" x

```

Figure 9-14 shows the file system table.

```
[root@TSM716 ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Thu Jun 30 08:17:22 2016
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/rhel-root / ext4 defaults 0 0
UUID=6edaf44e-7c08-44ee-a23e-27375f0c6368 /boot ext4 defaults 0 0
/dev/mapper/rhel-home /home ext4 defaults 0 0
/dev/mapper/rhel-swap swap swap defaults 0 0
/dev/mapper/tsmdb512-tsmdb512a /tsmdb512a ext4 defaults 0 0
/dev/mapper/tsmdb512-tsmdb512b /tsmdb512b ext4 defaults 0 0
/dev/mapper/tsmdb512-tsmdb512c /tsmdb512c ext4 defaults 0 0
/dev/mapper/tsmdb512-tsmdb512d /tsmdb512d ext4 defaults 0 0
/dev/mapper/tsmdb512-tsmlog00 /tsmact01 ext4 defaults 0 0
/dev/mapper/tsmarch-tsmarch01 /tsmarch01 ext4 defaults 0 0
/dev/mapper/tsmdbbk-tsmdbbk01 /tsmdbbk01 ext4 defaults 0 0
/dev/mapper/tsmdbbk-tsmdbbk02 /tsmdbbk02 ext4 defaults 0 0
/dev/mapper/tsmdbbk-tsmdbbk03 /tsmdbbk03 ext4 defaults 0 0
/dev/mapper/tsmdbbk-tsmdbbk04 /tsmdbbk04 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg01 /tsmstg01 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg02 /tsmstg02 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg03 /tsmstg03 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg04 /tsmstg04 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg05 /tsmstg05 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg06 /tsmstg06 ext4 defaults 0 0
/dev/mapper/tsmstg-tsmstg07 /tsmstg07 ext4 defaults 0 0
[root@TSM716 ~]# █
```

Figure 9-14 File system table

Figure 9-15 shows the RMAN output of the `list backup` command (truncated for brevity).

```

RMAN> list backup;

using target database control file instead of recovery catalog

List of Backup Sets
=====

BS Key Type LV Size Device Type Elapsed Time Completion Time
-----
1 Full 1.44T SBT_TAPE 07:42:03 12-JUL-16
BP Key: 1 Status: AVAILABLE Compressed: NO Tag: TAG20160712T155101
Handle: df_917020262_1_1 Media: 0
List of Datafiles in backup set 1
File LV Type Ckp SCN Ckp Time Name
-----
1 Full 60336183 12-JUL-16 +DATA/SAMPLE/DATAFILE/system.258.912970523
6 Full 60336183 12-JUL-16 +DATA/SAMPLE/DATAFILE/users.259.912970569
8 Full 60336183 12-JUL-16 +DATA/sample/soe.dbf

BS Key Type LV Size Device Type Elapsed Time Completion Time
-----
2 Full 8.44G SBT_TAPE 00:02:52 12-JUL-16
BP Key: 2 Status: AVAILABLE Compressed: NO Tag: TAG20160712T155101
Handle: df_917047989_2_1 Media: 0
List of Datafiles in backup set 2
File LV Type Ckp SCN Ckp Time Name
-----
2 Full 60414385 12-JUL-16 +DATA/SAMPLE/DATAFILE/undotbs3.263.912970643
3 Full 60414385 12-JUL-16 +DATA/SAMPLE/DATAFILE/sysaux.257.912970481
4 Full 60414385 12-JUL-16 +DATA/SAMPLE/DATAFILE/undotbs1.260.912970569
5 Full 60414385 12-JUL-16 +DATA/SAMPLE/DATAFILE/undotbs2.262.912970643
7 Full 60414385 12-JUL-16 +DATA/SAMPLE/DATAFILE/undotbs4.264.912970643

BS Key Type LV Size Device Type Elapsed Time Completion Time
-----
3 Full 18.75M SBT_TAPE 00:00:02 12-JUL-16
BP Key: 3 Status: AVAILABLE Compressed: NO Tag: TAG20160712T155101
Handle: df_917048184_3_1 Media: 0
SPFILE Included: Modification time: 07-JUN-16
SPFILE db_unique_name: SAMPLE
Control File Included: Ckp SCN: 60414703 Ckp time: 12-JUL-16

```

Figure 9-15 RMAN list backup

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM FlashSystem V9000 in a VersaStack Environment*, REDP-5264
- ▶ *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933
- ▶ *Implementing the IBM Storwize V7000 and IBM Spectrum Virtualize V7.6*, SG24-7938
- ▶ *Introducing and Implementing IBM FlashSystem V9000*, SG24-8273
- ▶ *VersaStack Solution by Cisco and IBM with IBM DB2, IBM Spectrum Control, and IBM Spectrum Protect*, SG24-8302
- ▶ *VersaStack Solution by Cisco and IBM with SQL, Spectrum Control, and Spectrum Protect*, SG24-8301

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other resources

These publications are also relevant as further information sources:

- ▶ *IBM System Storage Open Software Family SAN Volume Controller: CIM Agent Developers Reference*, SC26-7545
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Command-Line Interface User's Guide*, SC26-7544
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Configuration Guide*, SC26-7543
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Host Attachment Guide*, SC26-7563
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Installation Guide*, SC26-7541
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Planning Guide*, GA22-1052
- ▶ *IBM System Storage Open Software Family SAN Volume Controller: Service Guide*, SC26-7542

- ▶ *IBM System Storage SAN Volume Controller - Software Installation and Configuration Guide, SC23-6628*
- ▶ *IBM System Storage SAN Volume Controller V6.2.0 - Software Installation and Configuration Guide, GC27-2286*

Online resources

These websites are also relevant as further information sources:

- ▶ VersaStack Designs (links to PDF download page)
<http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/versastack-designs.html>
- ▶ VersaStack Solution - Cisco
<http://www.cisco.com/c/en/us/solutions/data-center-virtualization/versastack-solution-cisco-ibm/index.html>
- ▶ VersaStack Solution by Cisco and IBM
http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=PM&subtype=SP&htmlfid=TS03159USEN&appname=TAB_2_1_Appname
- ▶ Video: Client value of VersaStack
<https://www.youtube.com/watch?v=dvDG6UHMEuQ>
- ▶ Video: Growth Opportunities with VersaStack Solution
<https://www.youtube.com/watch?v=h32TsA2smLk>
- ▶ Video: High-Level Business Value of VersaStack from IBM and CISCO
<https://www.youtube.com/watch?v=E0W4ggyN99o>
- ▶ Video: IBM and Cisco VersaStack - Compression
<https://www.youtube.com/watch?v=xDbk4ddXzL0>
- ▶ Video: IBM and Cisco VersaStack - Data Virtualization
<https://www.youtube.com/watch?v=N-rNcokXzf0>
- ▶ Video: IBM and Cisco VersaStack - Flash Optimization and IBM Easy Tier
<https://www.youtube.com/watch?v=J7Rr13fEv0U>
- ▶ Video: IBM and Cisco VersaStack - Introduction
<https://www.youtube.com/watch?v=mkg1fkpAKII>
- ▶ Video: IBM and Cisco VersaStack - Turbo Compression
https://www.youtube.com/watch?v=PR_Uir1mxXE
- ▶ Video: New VersaStack Solution by Cisco and IBM
<https://www.youtube.com/watch?v=HHtgEABDYts>
- ▶ Video: Take 5 - VersaStack by Cisco and IBM
<https://www.youtube.com/watch?v=18mKR0skQ3o>
- ▶ Video: Talking VersaStack with Your Customers
<https://www.youtube.com/watch?v=UHANwo51ie0>

- ▶ Oracle Real Application Clusters - Overview
<https://www.oracle.com/database/real-application-clusters/index.html>
- ▶ Oracle Real Application Clusters - Resources
<https://www.oracle.com/database/real-application-clusters/resources.html>
- ▶ Oracle Database 12c
<http://www.oracle.com/technetwork/database/enterprise-edition/overview/index.html>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

VersaStack Solution by Cisco and IBM with Oracle RAC, IBM FlashSystem V9000, and IBM Spectrum Protect

SG24-8364-00

ISBN 0738442038



(0.5" spine)

0.475" x 0.873"

250 <-> 459 pages



SG24-8364-00

ISBN 0738442038

Printed in U.S.A.

Get connected

