

VersaStack Solution by Cisco and IBM with IBM DB2, IBM Spectrum Control, and IBM Spectrum Protect

Jon Tate

Vadi Bhatt

Sanjeev Naldurgkar

Filip Van Den Neucker

Asher Pemberton



Storage



International Technical Support Organization

**VersaStack Solution by Cisco and IBM with IBM DB2,
IBM Spectrum Control, and IBM Spectrum Protect**

February 2016

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

First Edition (February 2016)

This edition applies to the VersaStack software levels that are described in Chapter 3, “Software revisions and configuration guidelines” on page 11.

© 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	ix
Trademarks	x
IBM Redbooks promotions	xi
Preface	xiii
Authors	xiv
Now you can become a published author, too!	xv
Comments welcome	xvi
Stay connected to IBM Redbooks	xvi
Chapter 1. Introduction	1
1.1 Easy, efficient, and versatile	2
1.2 Evolving datacenter requirements	2
1.3 Holistic approach	3
1.4 Hardware options	3
1.5 Related information	5
Chapter 2. Architecture	7
Chapter 3. Software revisions and configuration guidelines	11
3.1 Software revisions	12
3.2 Configuration guidelines	13
Chapter 4. Planning for an IBM DB2 High Availability Server Failover Cluster implementation	19
4.1 Design considerations	20
4.1.1 Database workload	20
4.1.2 Server virtualization	20
4.1.3 Database availability	20
4.1.4 Quality of service and network segregation	20
4.1.5 Network availability	21
Chapter 5. Physical infrastructure	23
5.1 VersaStack cabling	24
5.2 Storage compatibility and interoperability	28
5.3 VersaStack system build process	29
Chapter 6. VersaStack Cisco Nexus 9000 Series Switches configuration	31
6.1 Cisco Nexus 9000 Series Switches network initial configuration setup	32
6.1.1 Configuring Cisco Nexus A	32
6.1.2 Configuring Cisco Nexus B	33
6.1.3 Enabling the Cisco Nexus 9000 Series Switch features and settings	34
6.1.4 Creating VLANs for VersaStack traffic	34
6.1.5 Configuring the Virtual Port Channel Domain	35
6.1.6 Configuring network interfaces for the vPC peer links	36
6.1.7 Configuring network interfaces to the Cisco UCS Fabric Interconnect	37
6.1.8 Linking in to an existing network infrastructure	40
Chapter 7. IBM Storwize V7000 storage configuration	41

7.1	Secure web access to the IBM Storwize V7000 service and management GUI.	42
7.2	IBM Storwize V7000 initial configuration setup.	42
Chapter 8. Cisco Unified Computing System configuration.		57
8.1	Performing the initial setup of Cisco UCS 6248 Fabric Interconnect for VersaStack environments	58
8.1.1	Cisco UCS 6248 A	58
8.1.2	Cisco UCS 6248 B	58
8.2	Cisco UCS for IBM Storwize V7000 storage system.	59
8.2.1	Logging in to Cisco UCS Manager	59
8.2.2	Upgrading Cisco UCS Manager software to Version 2.2(3d)	59
8.2.3	Adding a block of IP addresses for KVM access	59
8.2.4	Adding a block of IPv4 addresses for KVM access.	60
8.2.5	Synchronizing the Cisco UCS environment to NTP	61
8.2.6	Enabling the server and uplink ports.	63
8.2.7	Enabling Fibre Channel ports	66
8.2.8	Creating storage VSANs.	67
8.2.9	Configuring the FC storage ports	70
8.2.10	Configuring the VSAN for the FC storage ports	71
8.2.11	Creating WWNN pools	73
8.2.12	Creating WWPN pools	75
8.2.13	Creating vHBA templates for Fabric A and Fabric B.	78
8.2.14	Creating the storage connection policy for Fabric-A	80
8.2.15	Creating the Storage Connection Policy for Fabric-B	83
8.2.16	Acknowledging Cisco UCS chassis and FEX modules.	90
8.2.17	Creating uplink port channels to Cisco Nexus switches	90
8.2.18	Creating MAC address pools	93
8.2.19	Creating an UUID suffix pool.	95
8.2.20	Creating a server pool.	97
8.2.21	Creating virtual local area networks	98
8.2.22	Creating a host firmware package	100
8.2.23	Setting jumbo frames in Cisco UCS Fabric.	101
8.2.24	Creating a local disk configuration policy	103
8.2.25	Creating a network control policy for Cisco Discovery Protocol	104
8.2.26	Creating a power control policy.	105
8.2.27	Creating a server pool qualification policy (optional).	106
8.2.28	Creating a server BIOS policy	107
8.2.29	Creating a vNIC/vHBA placement policy for VM infrastructure hosts	109
8.2.30	Updating the default Maintenance Policy	110
8.2.31	Creating vNIC templates.	111
8.2.32	Creating boot policies	115
8.2.33	Creating service profile templates.	122
8.2.34	Creating service profiles	134
8.3	Backing up the Cisco UCS Manager configuration	137
Chapter 9. SAN boot		139
9.1	Adding hosts and mapping the boot volumes on the Storwize V7000 system	140
Chapter 10. VersaStack VMware ESXi 5.5 Update 2 SAN boot installation		145
10.1	The Cisco UCS 6200 Fabric Interconnect Cisco UCS Manager.	146
10.2	Setting up a VMware ESXi installation	147
10.2.1	ESXi hosts vm-host-infra-01 and vm-host-infra-02	147
10.3	Installing ESXi.	148
10.3.1	ESXi hosts vm-host-infra-01 and vm-host-infra-02	148

10.4	Setting up management networking for ESXi hosts	150
10.4.1	ESXi Host vm-host-infra-01	150
10.4.2	ESXi Host vm-host-infra-02	152
10.5	vSphere setup	153
10.5.1	Downloading the VMware vSphere Client and vSphere Remote CLI	153
10.6	Setting up VMkernel ports and the virtual switch	154
10.6.1	ESXi Host vm-host-infra-01	154
10.7	Mapping the required VMFS Datastores	157
10.7.1	Mapping the VMFS Datastores to the first host	157
10.8	Storage I/O Control	159
10.9	VersaStack VMware vCenter 5.5 Update 2	159
10.9.1	Installation steps for a simple installation of vCenter Server 5.5	160
10.10	Setting up a vCenter Server	165
10.10.1	vCenter Server VM	165
10.11	Mapping the datastores on the IBM Storwize V7000 second host after enabling the cluster	167
10.12	Optional: Adding domain account permissions	167
Chapter 11. IBM DB2 High Availability server and failover cluster implementation		171
11.1	Creating virtual machines	172
11.2	Installing Windows Server 2012 R2	179
11.3	Preparing the virtual machines for clustering	182
11.3.1	Renaming and assigning IP addresses to network adapters	182
11.3.2	Enabling jumbo frames for CSV traffic	184
11.3.3	Configuring the network adapters binding order	185
11.3.4	Installing Windows updates and adding roles and features	186
11.3.5	Adding Raw Device Mapping disks to the first virtual machine node	187
11.3.6	Adding Raw Device Mapping disks to the second virtual machine node	190
11.3.7	Preparing the disks for cluster use	194
11.4	Installing a Windows Server Failover Cluster	196
11.4.1	Installing a DB2 Server Failover Cluster	203
11.5	Modifying the vSphere High Availability and Distributed Resource Scheduler settings for the Windows Server Failover Cluster virtual machines	204
11.5.1	Creating anti-affinity rules	204
11.5.2	Enabling strict enforcement of anti-affinity rules	206
11.5.3	Setting the Distributed Resource Scheduler automation level for clustered virtual machines	207
11.5.4	Using vSphere Distributed Resource Scheduler groups and VM-Host affinity rules with clustered virtual machines	209
11.6	Installing DB2	213
11.7	Installing DB2 Data Studio Client	230
11.8	Deploying the DB2 GSDB sample database	245
11.9	Configuring DB2 High Availability	248
11.9.1	Connecting to the DB2 instance with the Data Studio Client	250
11.9.2	Configuring the database transaction logging	252
11.9.3	Configuring the HADR setup	253
Chapter 12. IBM Spectrum Control integration		255
12.1	Spectrum Control overview	256
12.2	Storage hypervisor	256
12.3	IBM SmartCloud Virtual Storage Center component model	258
12.3.1	Storage management	258
12.3.2	Storage virtualization	261

12.3.3	Application-aware data protection	264
12.4	IBM SmartCloud Virtual Storage Center features	265
12.4.1	Efficient by design	265
12.4.2	Self-optimizing	266
12.4.3	Cloud agility	267
12.5	IBM SmartCloud Virtual Storage Center interfaces	268
12.5.1	VMware	269
12.6	IBM SmartCloud Virtual Storage Center offerings	274
12.6.1	License model overview	275
12.6.2	VSC for Storwize Family license	275
12.7	VersaStack: Spectrum Control	276
12.7.1	Tivoli Productivity Center Virtual Storage Edition installation	276
12.7.2	Integrating the Storwize V7000 storage system with Spectrum Control	278
12.7.3	Monitoring and alerting	290
12.8	Advanced Analytics	324
12.8.1	Cloud Configuration	324
12.8.2	Provisioning	327
12.8.3	Integrating servers and virtual machines	342
12.8.4	Reporting for departments and applications	350
12.9	Resources	353
Chapter 13. IBM Spectrum Protect integration		355
13.1	IBM Spectrum Protect Suite for Unified Recovery overview	356
13.1.1	IBM Spectrum Software Defined Storage Suite	356
13.1.2	IBM Spectrum Protect Suite for Unified Recovery	358
13.1.3	Licensing metrics	365
13.2	IBM Spectrum Protect implementation	367
13.2.1	Architectural overview	367
13.2.2	Guest support for virtual machines and virtualization	370
13.2.3	Blueprints	371
13.2.4	Multi-site setup	373
13.2.5	Summary	378
13.3	Protecting the VMware infrastructure	379
13.3.1	Deploying IBM Spectrum Protect for Virtual Environments	379
13.3.2	Storwize V7000 FlashCopy mapping	381
13.3.3	Protecting VMware data	384
13.3.4	Summary	389
13.4	Protecting the DB2 HADR virtual machines	390
13.4.1	Using IBM Spectrum Protect to back up DB2 data	390
13.4.2	DB2 quiescing commands	391
13.5	Using IBM Spectrum Protect Advanced Protection and Recovery Technologies	392
13.5.1	Progressive incremental backups	392
13.5.2	Data deduplication	394
13.5.3	Node replication with automated failover	397
13.6	Monitoring and managing the IBM Spectrum Protect environment	399
13.6.1	Data Protection for VMware	400
13.6.2	IBM Spectrum Protect Operations Center	400
13.6.3	Reporting and monitoring for IBM Spectrum Protect	405
Chapter 14. General performance		413
14.1	IBM Easy Tier	414
14.2	Autotier	415
14.3	General performance metrics	418

14.3.1 B200 M4	418
14.3.2 VIC 1340	419
14.3.3 Storwize V7000 storage system	419
Chapter 15. General validation	421
15.1 Validation scenarios	422
15.2 Storwize V7000 failover validation	422
15.2.1 Unexpected Fibre Channel cable failure	423
15.2.2 Unexpected node failure	427
15.3 Cisco Nexus devices	432
15.3.1 vPC peer switch failure validation	432
15.4 Cisco UCS service profile	435
15.4.1 Service profile migration validation	436
Appendix A. Windows Active Directory and running configurations	441
Building Windows Active Directory Server virtual machines	442
Nexus 9000 running configuration	444
Nexus 9000 A running configuration	444
Nexus 9000 B running configuration	447
Related publications	451
IBM Redbooks	451
Other resources	451
Online resources	452
Help from IBM	453

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 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®	Informix®
Cognos®	IBM Elastic Storage™	Lotus®
DB2®	IBM FlashCore™	MicroLatency®
DB2 Connect™	IBM FlashSystem®	Netcool®
developerWorks®	IBM SmartCloud®	Passport Advantage®
Domino®	IBM Spectrum™	ProtecTIER®
DS4000®	IBM Spectrum Accelerate™	Real-time Compression™
DS5000™	IBM Spectrum Archive™	Redbooks®
DS6000™	IBM Spectrum Control™	Redbooks (logo)  ®
DS8000®	IBM Spectrum Protect™	Storwize®
Easy Tier®	IBM Spectrum Scale™	System Storage®
FlashCopy®	IBM Spectrum Storage™	Tivoli®
FlashSystem™	IBM Spectrum Virtualize™	XIV®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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.

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

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 require an IT infrastructure to be easy, efficient, and versatile.

The VersaStack solution by Cisco and IBM® can help you accelerate the deployment of your datacenters. 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 Unified Computing System (Cisco UCS) Integrated Infrastructure with the efficiency of the IBM Storwize® storage system. The Cisco UCS Integrated Infrastructure includes the Cisco UCS, Cisco Nexus and Cisco MDS switches, and Cisco UCS Director. The IBM Storwize V7000 storage system 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 (CVDs) to provide faster delivery of applications, greater IT efficiency, and less risk.

This IBM Redbooks® publication is aimed at experienced storage administrators that are tasked with deploying a VersaStack solution with IBM DB2® High Availability (DB2 HA), IBM Spectrum Protect™, and IBM Spectrum Control™.

Authors

This book was produced by a team of specialists from around the world working at the Cisco campus in San Jose, California.



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 29 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.



Vadi Bhatt has been a performance architect for the last four years with the solutions and benchmarking division of the Cisco UCS technical marketing team. Vadi leads a team in Bangalore, India that is primarily focused on developing solution guides and CVDs, and delivering industry-standard benchmark numbers on Cisco UCS. His areas of focus include Oracle Applications, Microsoft applications, and preferred practices for Cisco UCS infrastructure offerings, including VSPEX, Flexpod, Citrix Cloud Platform, and OpenStack. Vadi has over 18 years of experience in enterprise software development involving large-scale, distributed, and clustered systems. He has six patents to his credit in the area of enterprise relation database system architecture. He has extensive knowledge about distributed database systems, big data, and enterprise application design. Before his role at Cisco, Vadi was with Sybase Inc (acquired by SAP) as Technical Director in the performance engineering group, where he designed both OLTP (ASE) and DSS (Sybase IQ) systems for performance and scalability. Vadi holds a master degree in computer science and engineering from the Indian Institute of Technology, Mumbai.



Sanjeev Naldurgkar is a Technical Marketing Engineer with Cisco's Datacenter Group. He has 14 years of experience in information technology. His focus areas include Cisco UCS, Microsoft platforms, server virtualization, and storage technologies. Before joining Cisco, he was a Support Engineer at Microsoft Global Technical Support Center. Sanjeev holds a bachelor degree in electronics and communication engineering, and industry certifications from Microsoft and VMware.



Filip Van Den Neucker is a Senior Technical Consultant for IBM Systems Storage Software covering Belgium and Luxembourg. He joined IBM in 2011. His main areas of expertise are the IBM Spectrum Protect (formerly IBM Tivoli® Storage Manager) and IBM Spectrum Control (formerly IBM Tivoli Productivity Center) software products. Holding a master degree in linguistics, he combines his passion for languages and education with life-long learning and ongoing new technology exploration. Throughout his 19 years of experience in the IT industry, he held several positions in global support, development, and IT and datacenter management, which provides him with an in-depth background in IT systems, datacenter and communication infrastructures, virtualization technologies, storage hardware, storage management, and backup software. As he deems knowledge transfer to be important, he regularly organizes technical workshops and architecture sessions for IBM Business Partners, customers, and colleagues in the BeNeLux area. He also engages in local community projects advising about the IT infrastructure of local educational organizations and giving technology exploration lectures and workshops to youngsters.



Asher Pemberton works in the Manchester SAN Volume Controller test team. He has been working with IBM for two years as part of a global team developing, testing, and supporting IBM Storage products. Asher has been integral in system test verification of releases 7.2, 7.3, and 7.4 of SAN Volume Controller and IBM Storwize, and has been involved in the testing and development of new products in the Storwize product family from their inception to release to market. He recently trained a new team in Guadalajara, Mexico so that they can test the current SAN Volume Controller release. Before joining IBM, he received a master degree in physics from the University of Bristol.

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

Caela Dehaven, Chris O'Brien, Ruchi Jain, Vijay Durairaj
Cisco

Sally Neate, Paul Merrison, Matt Smith, Eric Stouffer, Ian Shave, Warren Hawkins, Rob Wallis
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

Cisco and IBM have a long history of working together to deliver technology that benefits their mutual clients. Expanding on this success, IBM and Cisco have announced VersaStack, an innovative validated design that brings together IBM Storwize storage systems 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 storage systems into the single pane of glass management environment that is provided by Cisco UCS Director, with future capabilities to deliver Application Centric Infrastructure (ACI) and Intercloud Fabric from Cisco and usage within the IBM Cloud Builder Professional Services offering from IBM Global Services.

VersaStack is backed by a series of Cisco Validated Designs (CVDs) and IBM Redbooks publications that are developed by Cisco and IBM, providing faster delivery of applications, greater reliability, and confidence for their customers and business partners.

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 require an IT infrastructure that is easy, efficient, and versatile. The VersaStack solution by Cisco and IBM helps accelerate datacenter infrastructure deployment, efficiently manage information and resources, and adapt to business change.

VersaStack includes the IBM Storwize V7000 storage system, which includes technologies that both complement and enhance virtual environments with built-in functions, such as IBM Data Virtualization, Real-time Compression, and Easy Tier, which 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 Unified Computing System (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 Storwize storage system.

VersaStack is backed by CVDs, for faster delivery of applications and increased IT efficiency with less risk. VersaStack is supported by a broad range of services from 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, which is 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 datacenter requirements

The datacenter industry is always evolving, and current trends make evolution more critical than ever. The datacenter has moved far beyond a simple repository for digital records, and way beyond just a vehicle for backup and restore.

Increasingly, the datacenter's 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 by using traditional computing tools.

Two other trends are making it easier to provision datacenter 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 datacenter 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 that is implied by SDN helps organizations manage the increased compute requirements of big data. Underpinning these trends are changes in hardware. Vendors are adapting specific datacenter components to address cloud, SDN, and big data requirements. IBM, for example, evolved its Storwize family of virtualized storage technologies specifically for software-defined environments (SDEs).

Cisco, meanwhile, developed Cisco Application Centric Infrastructure (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 datacenter. Previously, datacenter teams purchased computing, storage, and network building blocks separately and assembled, configured, and tested the various technologies with the hope everything worked 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

This approach is more than just integrating hardware and software. Both IBM and Cisco are fully aware of the requirements of the enterprise today. Therefore, it made 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 worldwide trust IBM and Cisco products as best in industry, and this partnership takes this quality to a new level.

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

1.4 Hardware options

All the screen captures and work in this document come from a Storwize V7000 Gen2 storage system, with a combination of SAS and SSD drives. VersaStack can also be used with Storwize V5000 and IBM FlashSystem™ V9000 storage systems (system validation upcoming).

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

The FlashSystem V9000 storage system 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. Using its flash-optimized design, the FlashSystem V9000 storage system 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 demanding extreme performance.

For more information about the FlashSystem V9000 storage system, see the following resources:

- ▶ <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 FlashSystem V9000 solution, the IBM FlashSystem® 900 storage system can go behind stand-alone SAN Volume Controller 2145-DH8 nodes, offering greater flexibility.

The FlashSystem 900 storage system 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 tier automatically I/O-intensive payloads to the FlashSystem 900 storage system.

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

This option is also available with the existing Storwize V7000 storage system. You simply can add the FlashSystem 900 storage system 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>

Table 1-1 shows a comparison of the SAN Volume Controller and Storwize nodes.

Table 1-1 A quick comparison of Storwize V5000, Storwize V7000, and SAN Volume Controller 2145-DH8 nodes

Component	Storwize V5000	Storwize V7000	SAN Volume Controller 2145-DH8
Standard Host Interface	6 Gb SAS, 1 Gb iSCSI, 8 Gb FC, or 10 Gb iSCSI/FCoE	1 Gb iSCSI	1 Gb iSCSI
Optional Host Interface	None	Two 8 Gb/16 Gb FC or 10 Gb iSCSI/FCoE)	Three (8 Gb/16 Gb FC or 10 Gb iSCSI/FCoE)
RAM (per node)	8 GB	32 or 64 GB	32 or 64 GB
Expansion Enclosures (per control enclosure)	Up to 19	Up to 20	Up to 2 (with 12 Gb SAS HIC)
Licensed Function Enforcement	Honor	Honor	Honor
IBM FlashCopy®	License (per enclosure)	License (per enclosure)	License (per TiB)
Remote Copy	License (per enclosure)	License (per enclosure)	License (per TiB)
Easy Tier	License (per enclosure)	License (per enclosure)	License (per TiB)
System Clustering	Yes: 2 control enclosures	Yes: 4 control enclosures	Yes: 4 control enclosures
General External Virtualization	License (per enclosure)	License (per enclosure)	License (per TiB)

Component	Storwize V5000	Storwize V7000	SAN Volume Controller 2145-DH8
Data Migration from external storage	Yes	Yes	Yes
Compression	No	License (per enclosure)	License (per TiB)
Compression Hardware	No	Yes, optional extra	Yes, optional extra
NAS	No	Yes, Storwize V7000 Unified	No

1.5 Related information

Here are links to resources that are related to VersaStack that might interest you:

- ▶ 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 and CISCO
<https://www.youtube.com/watch?v=E0W4gggN99o>
- ▶ 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 Easy Tier
<https://www.youtube.com/watch?v=J7Rr13fEv0U>
- ▶ Video: IBM and Cisco VersaStack - Flash Optimization and IBM Easy Tier
<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 and IBM
<https://www.youtube.com/watch?v=18mKR0skQ3o>

Architecture

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

The IBM DB2 on VersaStack design combines an IBM DB2 High Availability (HA) setup running on VersaStack with IBM Spectrum Control and IBM Spectrum Protect, as shown in Figure 2-1.

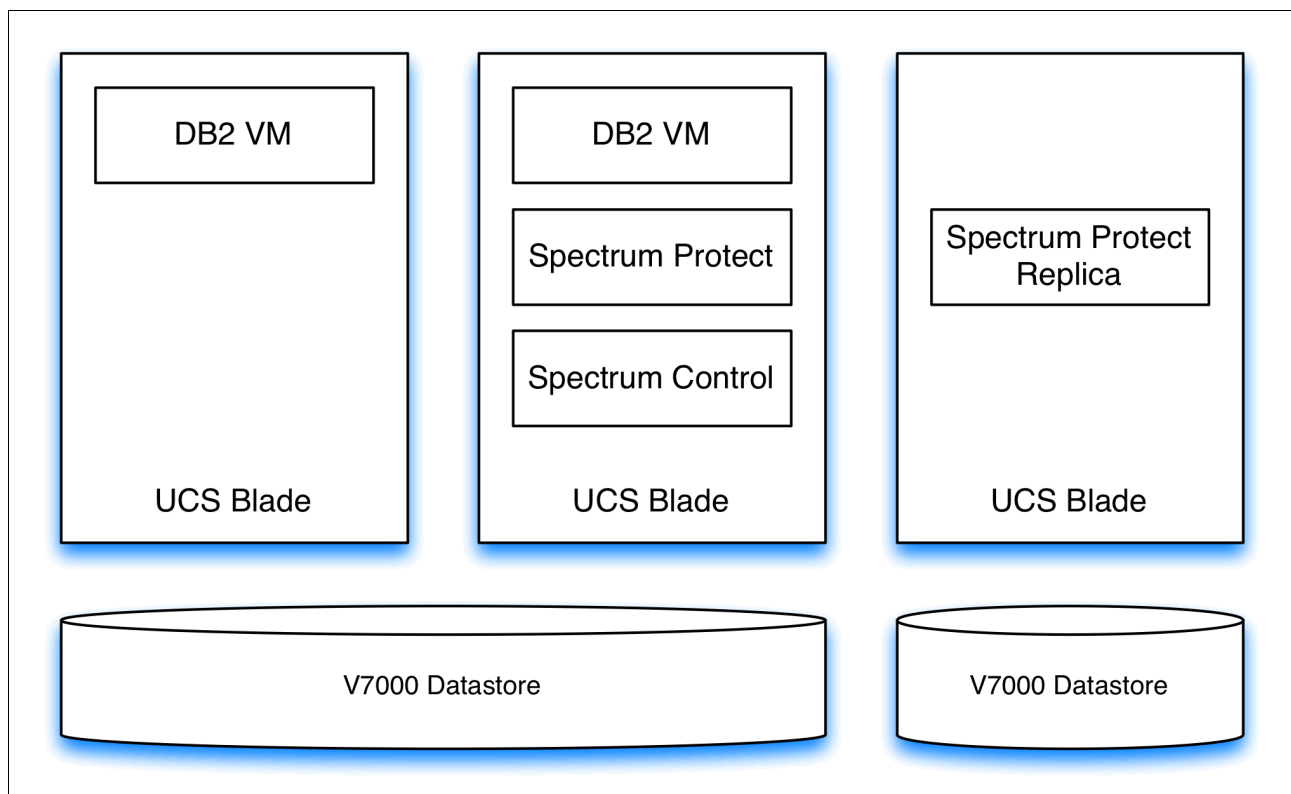


Figure 2-1 DB2 on VersaStack with IBM Spectrum Protect and Spectrum Control

It uses the flexibility of the Cisco Fabric Interconnect to operate in FC Switching Mode. This deployment model eliminates the need for a separate Fibre Channel switch to help reduce deployment costs. UCS Manager SAN Connectivity Policies are used to help automate SAN zoning for the administrator.

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

Specifically, VersaStack is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions.

VMware vSphere built on VersaStack includes a Storwize V7000 storage system, Cisco networking, the Cisco UCS, Cisco 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 datacenter rack or be deployed according to a customer's datacenter design. Port density enables the networking components to accommodate multiple configurations.

One benefit of the VersaStack architecture is its ability to meet any customer's capacity or performance needs in a cost-effective manner. A storage system that can serve multiple protocols across a single interface allows for customer choice and investment protection because it is a wire-once architecture.

This architecture references relevant criteria pertaining to resiliency, cost benefit, and ease of deployment of all components, including a Storwize V7000 storage system.

The architecture for this solution, which is shown in Figure 2-2 on page 9, uses two sets of hardware resources:

- ▶ Common Infrastructure services on redundant and self-contained hardware
- ▶ VersaStack PoD with DB2 HA, Spectrum Control, and IBM Spectrum Protect

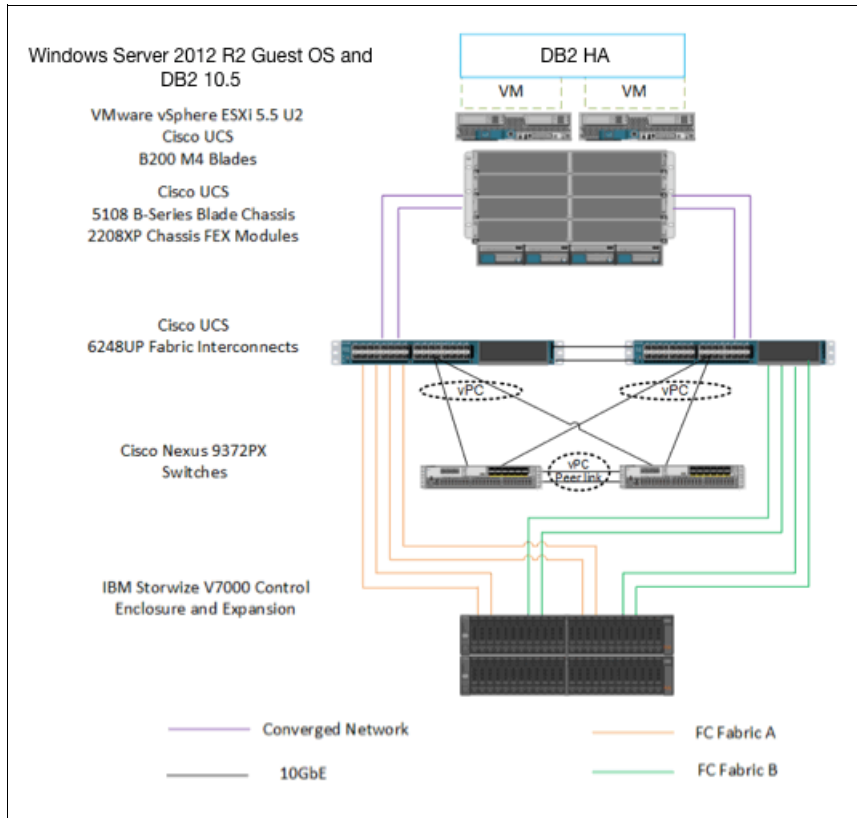


Figure 2-2 DB2 HA built on VersaStack

The common infrastructure services include Active Directory, DNS, DHCP, vCenter, Nexus 1000v virtual supervisor module (VSM), and any other shared service. These components are considered core infrastructure because they provide necessary datacenter-wide services for where the VersaStack PoD is. Because these services are integral to the deployment and operation of the platform, there is a need to adhere to preferred practices in their design and implementation, which includes such features as high-availability, appropriate RAID setup, and performance and scalability considerations because such services might need to be extended to multiple PoDs. At a customer's site, depending on whether this is a new datacenter, there might not be a need to build this infrastructure piece.

Figure 2-2 illustrates IBM DB2 HA built on VersaStack components and the network connections for a configuration with a Storwize V7000 storage system. This Fabric Interconnect direct attached design allows a connection to the Storwize V7000 storage controllers without using separate Fibre Channel switches.

VersaStack uses the Cisco Nexus 9396, and Cisco UCS B-Series with the Cisco UCS virtual interface card (VIC) and the IBM Storwize V7000 storage controllers connected in a highly available design that uses Cisco Virtual Port Channels (vPCs). This infrastructure is deployed to provide FC-booted hosts with block-level access to shared storage datastores.

The reference hardware configuration includes the following items:

- ▶ Two Cisco Nexus 9396 or 9372 switches
- ▶ Two Cisco UCS 6248UP Fabric Interconnects
- ▶ 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 additional fabric extenders and blade server chassis
- ▶ One IBM Storwize V7000 storage system, which contains a V7000 control enclosure and V7000 expansion enclosure. Support for up to 504 small form-factor (SFF) disks of any capacity.
- ▶ Support for up to a total of four V7000 control enclosures, up to 80 Storwize V7000 expansion enclosures, and up to 1056 SFF or large form-factor (LFF) disks of any capacity.

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, additional 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, DB2 HA deployment, and Spectrum Control and IBM Spectrum Protect overviews.

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

http://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/UCS_CVDs/Versastack_design.pdf



Software revisions and configuration guidelines

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

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	2.2(3c)	Embedded management
	Cisco UCS 5108 Blade Server Chassis	N/A	Software runs on FI
	Cisco UCS B 200 M4	2.2(3c)	Software bundle release
	Cisco ENIC	2.1.2.59	Ethernet driver for Cisco VIC
	Cisco FNIC	1.6.0.12	FCoE driver for Cisco VIC
Network	Cisco Nexus 9000 c9372PX	6.1(2)I3(3a)	Operating system version
Storage	IBM Storwize V7000	7.5.0.0	Software version

Layer	Device	Version/Release	Details
Software	Cisco UCS hosts	VMware vSphere ESXi 5.5u2	Operating system version
	Microsoft SQL Server	Microsoft SQL Server 2008 R2	Built-in server for vCenter
	VMware vCenter	5.5u2	Software version
	Windows Server	Windows Server 2012 R2	Operating system version
	IBM DB2	IBM DB2 10.5	Operating system version
	IBM Spectrum Control (IBM SmartCloud® Virtual Storage Center)	5.2.6	Software version
	IBM Spectrum Protect for Virtual Environments	7.1.2	Software version
	IBM Spectrum Protect	7.1.1	Software version
	IBM Tivoli Monitoring for Spectrum Protect	7.1	Software version
	IBM Spectrum Protect Snapshot	4.1.2	Software version

3.2 Configuration guidelines

This document provides details about configuring a fully redundant, highly available VersaStack unit with a Storwize V7000 storage system. Therefore, reference is made at each step to the component being configured as either 01 or 02. For example, node01 and node02 are used to identify the two IBM storage controllers that are provisioned with this document, and Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured.

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

Finally, to indicate that you should include information pertinent to your environment in a given 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 Network port

```
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	1	All infrastructure management in this VLAN
vMotion	30	VMware vMotion traffic
WinClus	40	Windows Cluster heartbeat traffic
WinCSV	50	Windows Cluster shared volume traffic
Backup	60	Backup traffic for storage

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	Customer host name
Active Directory (contains DHCP and DNS)	
vCenter Server	

Table 3-4 lists the configuration variables that are used throughout this book. This table can be completed based on the specific site variables and 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	

Variable	Description	Customer value
<<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	
<<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	

Variable	Description	Customer value
<<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_winclus_vlan_id>>	Windows Cluster heartbeat traffic	
<<var_wincsv_vlan_id>>	Windows Cluster shared volume traffic	
<<var_backup_vlan_id>>	Backup traffic for storage	
<<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_infra_01_ip>>	VMware ESXi host 01 in-band management IP	
<<var_vmhost_infra_01_2nd_ip>>	VMware ESXi host 01 secondary in-band management IP	

Variable	Description	Customer value
<<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.		



Planning for an IBM DB2 High Availability Server Failover Cluster implementation

This chapter describes some of the considerations and assumptions that are followed during the design of the DB2 High Availability (HA) installation.

4.1 Design considerations

The goal of this chapter is to come up with a simple and efficient DB2 Server database design that is suited for a VersaStack solution. The major design considerations of the recommended 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 driving 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. There are certain OLTP workloads that 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, you can use SSDs or flash memory to store temporary database (tempdb) files.

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 DB2 Server transactional workloads. Each of the virtual machines (VMs) hosting the DB2 Server 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

The configuration ensures database instance level availability by using IBM DB2 High Availability that runs on Microsoft Server Clustering technology. 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, in this configuration for DB2, VMs use built-in DB2 capabilities to provide the high availability in an active-passive setup. Alternatively, the Microsoft Failover Cluster can be used for failover functions. On the DB2 VMs, anti-affinity rules are set to prevent VMs from migrating under the HA/DRS feature, which ensures that VMs are not placed on the same ESXi, 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.



Physical infrastructure

This chapter describes the physical infrastructure that was implemented and used in this book.

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.

The tables in this section contain details for the prescribed and supported configuration of an IBM Storwize V7000 storage system running Version 7.4.0.

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 result in changes to the deployment procedures that follow because specific port locations are mentioned.

It is possible to order Storwize V7000 storage systems in a different configuration from what is presented in the tables in this section. Before starting, be sure that the configuration matches the descriptions in the tables and diagrams in this section.

Figure 5-1 and Figure 5-2 on page 25 show cabling diagrams for a VersaStack configuration that uses the Cisco Nexus 9300 and Storwize V7000 storage system. For SAS cabling information, the Storwize V7000 control enclosure and expansion enclosure should be connected according to the cabling guide found at the following website:

http://www.ibm.com/support/knowledgecenter/ST3FR7_7.4.0/com.ibm.storwize.v7000.740.doc/v3500_qisascables_b4jtyu.html?cp=ST3FR7%2F1-3-0-1-3

Figure 5-1 shows the VersaStack block-only cable diagram.

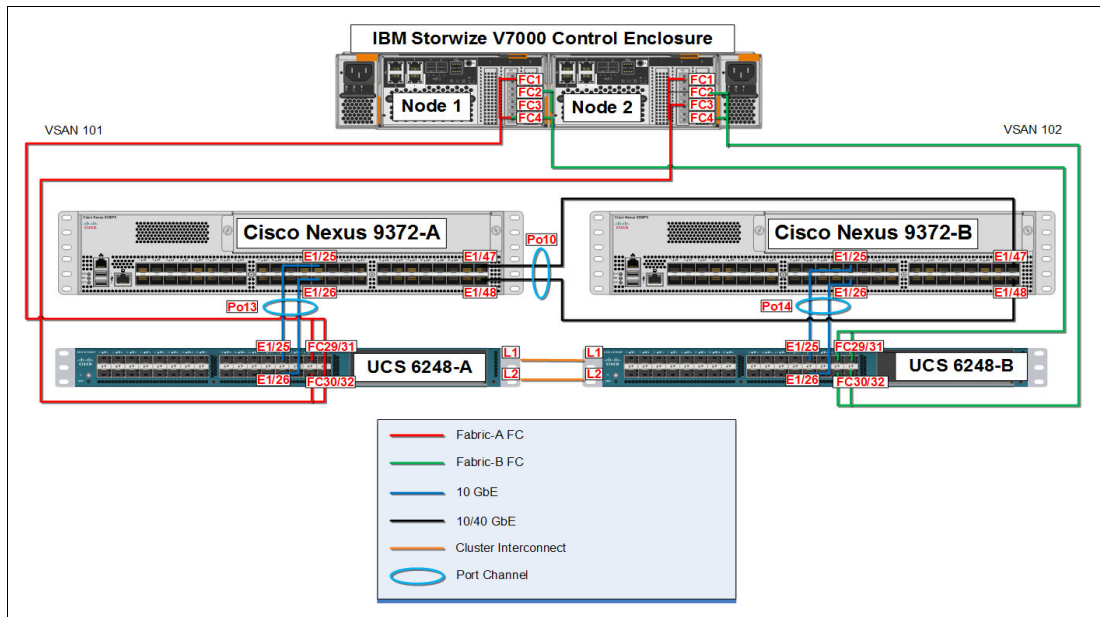


Figure 5-1 VersaStack block-only cable diagram

Figure 5-2 on page 25 shows the VersaStack management cable diagram.

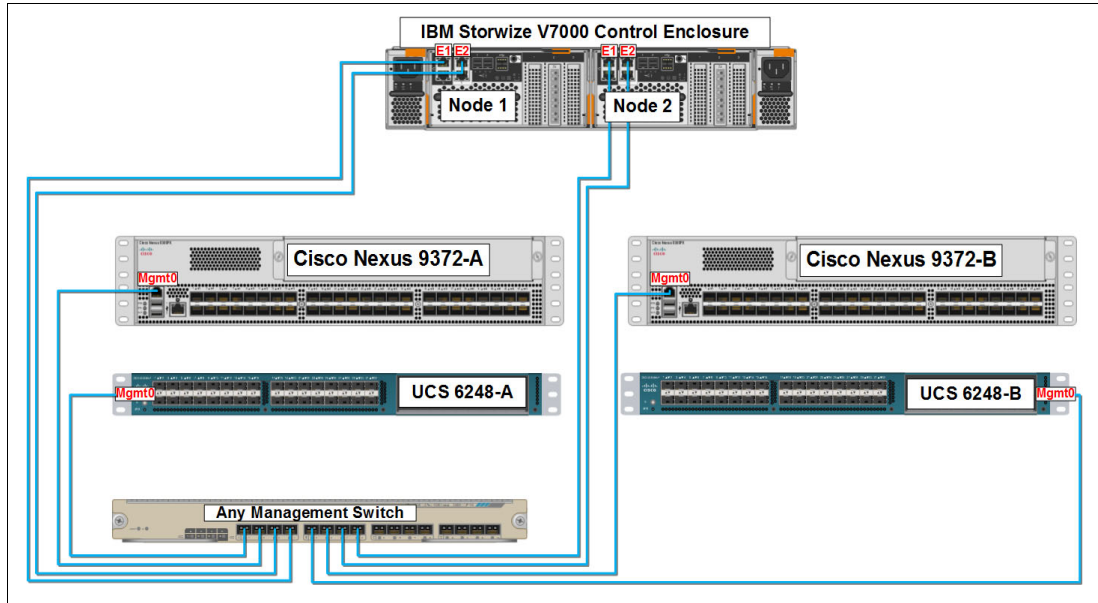


Figure 5-2 VersaStack management cable diagram

Table 5-1 shows the Cisco Nexus 9000-A cabling information.

Table 5-1 Cisco Nexus 9000-A cabling information

Local device	Local port	Connection	Remote device	Remote port
Cisco Nexus 9000-A	Eth1/25	10 GbE	Cisco UCS fabric interconnect-A	Eth1/25
	Eth1/26	10 GbE	Cisco UCS fabric interconnect-B	Eth1/26
	Eth1/47	10 GbE	Cisco Nexus 9000-B	Eth1/47
	Eth1/48	10 GbE	Cisco Nexus 9000-B	Eth1/48
	Eth1/1	GbE	GbE management switch	Any

Table 5-2 shows the Cisco Nexus 9000-B cabling information.

Table 5-2 Cisco Nexus 9000-B cabling information

Local device	Local port	Connection	Remote device	Remote port
Cisco Nexus 9000-B	Eth1/25	10 GbE	Cisco UCS fabric interconnect-A	Eth1/25
	Eth1/26	10 GbE	Cisco UCS fabric interconnect-B	Eth1/26
	Eth1/47	10 GbE	Cisco Nexus 9000-A	Eth1/47
	Eth1/48	10 GbE	Cisco Nexus 9000-A	Eth1/48
	Eth1/1	GbE	GbE management switch	Any

Table 5-3 shows the IBM Storwize V7000 Controller Node-1 cabling information.

Table 5-3 IBM Storwize V7000 Controller Node-1 cabling information

Local device	Local port	Connection	Remote device	Remote port
IBM Storwize V7000 Controller, Node-1	E1/E2	GbE	GbE management switch	Eth1/25
	FC1	8 Gbps	Cisco UCS fabric interconnect-A	FC1/29
	FC2	8 Gbps	Cisco UCS fabric interconnect-B	FC1/29
	FC3	8 Gbps	Cisco UCS fabric interconnect-B	FC1/31
	FC4	8 Gbps	Cisco UCS fabric interconnect-A	FC1/31

Table 5-4 shows the IBM Storwize V7000 Controller Node-2 cabling information.

Table 5-4 IBM Storwize V7000 Controller Node-2 cabling information

Local device	Local port	Connection	Remote device	Remote port
IBM Storwize V7000 Controller, Node-2	E1/E2	GbE	GbE management switch	Eth1/25
	FC1	8 Gbps	Cisco UCS fabric interconnect-A	FC1/30
	FC2	8 Gbps	Cisco UCS fabric interconnect-B	FC1/30
	FC3	8 Gbps	Cisco UCS fabric interconnect-B	FC1/32
	FC4	8 Gbps	Cisco UCS fabric interconnect-A	FC1/32

Table 5-5 shows the Cisco UCS Fabric Interconnect-A cabling information.

Table 5-5 Cisco UCS Fabric Interconnect-A cabling information

Local device	Local port	Connection	Remote device	Remote port
Cisco UCS fabric interconnect-A	Mgmt0	GbE	GbE management switch	Any
	Eth1/25	10 GbE	Cisco Nexus 9000-A	Eth1/25
	Eth1/26	10 GbE	Cisco Nexus 9000-B	Eth1/26
	Eth1/1	10 GbE	Cisco UCS Chassis FEX-A	IOM 1/1
	Eth1/2	10 GbE	Cisco UCS Chassis FEX-A	IOM 1/2
	FC1/29	8 Gbps	V7000 Controller, Node-1	FC1
	FC1/31	8 Gbps	V7000 Controller, Node-1	FC4
	FC1/29	8 Gbps	V7000 Controller, Node-2	FC2
	FC1/31	8 Gbps	V7000 Controller, Node-2	FC3
	L1	GbE	Cisco UCS fabric interconnect-B	L1
	L2	GbE	Cisco UCS fabric interconnect-B	L2

Table 5-6 shows Cisco UCS Fabric Interconnect-A cabling information.

Table 5-6 Cisco UCS Fabric Interconnect-A cabling information

Local device	Local port	Connection	Remote device	Remote port
Cisco UCS fabric interconnect-B	Mgmt0	GbE	GbE management switch	Any
	Eth1/25	10 GbE	Cisco Nexus 9000-B	Eth1/25
	Eth1/26	10 GbE	Cisco Nexus 9000-A	Eth1/26
	Eth1/1	10 GbE	Cisco UCS Chassis FEX-B	IOM 1/1
	Eth1/2	10 GbE	Cisco UCS Chassis FEX-B	IOM 1/2
	FC1/30	8 Gbps	V7000 Controller, Node-1	FC1
	FC1/32	8 Gbps	V7000 Controller, Node-1	FC4
	FC1/30	8 Gbps	V7000 Controller, Node-2	FC2
	FC1/32	8 Gbps	V7000 Controller, Node-2	FC3
	L1	GbE	Cisco UCS fabric interconnect-A	L1
	L2	GbE	Cisco UCS fabric interconnect-A	L2

5.2 Storage compatibility and interoperability

The IBM System Storage Interoperation Center (SSIC) provides information about supported external hardware and software for the specific Storwize V7000 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 Midrange Disk**, then click **Storwize V7000** or **Storwize V7000 Unified Host Attachment or Storage Controller Attachment**.

Software and hardware limitations for specific Storwize V7000 versions can be found at:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004923>

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

<http://www.ibm.com/support/docview.wss?uid=ssg1S1004941>

5.3 VersaStack system build process

Figure 5-3 illustrates the DB2 VersaStack system build workflow.

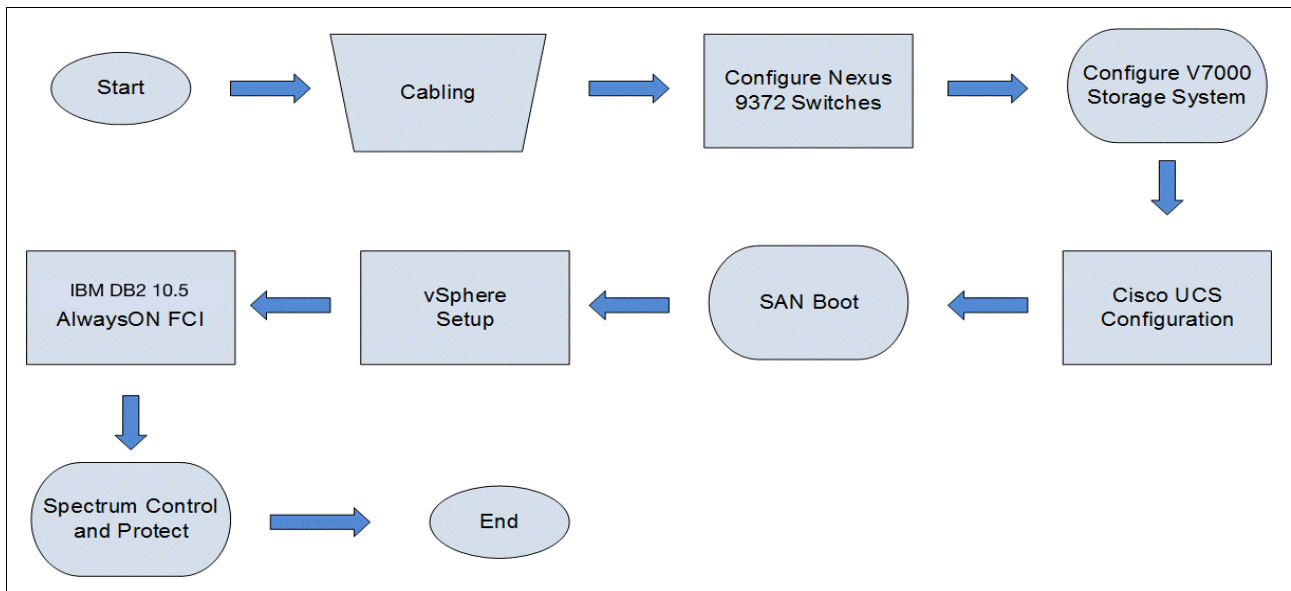


Figure 5-3 DB2 VersaStack installation workflow



VersaStack Cisco Nexus 9000 Series Switches configuration

This chapter provides detailed instructions for configuring Cisco Nexus 9000 Series Switches in a VersaStack environment. After the procedures are complete, the configuration provides higher throughput and redundant Layer 2 network connectivity for the Cisco Unified Computing System (UCS) environment to the upstream switches. Cisco Nexus 9000 Series Switches are Application Centric Infrastructure (ACI) ready, which provides a foundation for automating application deployments and delivering simplicity, agility, and flexibility. These deployment procedures are customized to include the environment variables.

6.1 Cisco Nexus 9000 Series Switches network initial configuration setup

This section provides details for the initial setup of two Cisco Nexus 9000 Series Switches.

6.1.1 Configuring Cisco Nexus A

To set up the initial configuration for the first Cisco Nexus 9000 Series Switch (named Cisco Nexus A in this example), complete the procedure that is shown in Example 6-1.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

Example 6-1 Configuration of Cisco Nexus A

```
Abort Auto Provisioning and continue with normal setup?(yes/no) [n]: y
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:
Enter the password for "admin":
Confirm the password for "admin":
---- Basic System Configuration Dialog VDC: 1 ----
This setup utility will guide you through the basic configuration of the system.
Setup configures only enough connectivity for management of the system.
Please register Cisco Nexus9000 Family devices promptly with your supplier.
Failure to register may affect response times for initial service calls. Nexus9000
devices must be registered to receive entitled support services.
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the
remaining dialogs.
Would you like to enter the basic configuration dialog (yes/no): y
Create another login account (yes/no) [n]: n
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : <<var_nexus_A_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Mgmt0 IPv4 address : <<var_nexus_A_mgmt0_ip>>
Mgmt0 IPv4 netmask : <<var_nexus_A_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway : <<var_nexus_A_mgmt0_gw>>
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
Type of ssh key you would like to generate (dsa/rsa) [rsa]:
Number of rsa key bits <1024-2048> [1024]: 2048
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address : <<var_global_ntp_server_ip>>
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut) [noshut]:
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
The following configuration will be applied:
password strength-check
switchname <<var_nexus_A_hostname>>
vrf context management
```

```

ip route 0.0.0.0/0 <<var_nexus_A_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 2048 force
feature ssh
ntp server <<var_global_ntp_server_ip>>
system default switchport
no system default switchport
copp profile strict
interface mgmt0
ip address <<var_nexus_A_mgmt0_ip>> <<var_nexus_A_mgmt0_netmask>>
no shutdown
Would you like to edit the configuration? (yes/no) [n]:
Use this configuration and save it? (yes/no) [y]:
[#####] 100%
Copy complete.

```

6.1.2 Configuring Cisco Nexus B

To set up the initial configuration for the second Cisco Nexus 9000 Series Switch (named Cisco Nexus B in this example), complete the procedure that is shown in Example 6-2.

Note: On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

Example 6-2 Configuration of Cisco Nexus B

```

Abort Auto Provisioning and continue with normal setup?(yes/no)[n]: y
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:
Enter the password for "admin":
Confirm the password for "admin":
---- Basic System Configuration Dialog VDC: 1 ----
This setup utility will guide you through the basic configuration of the system.
Setup configures only enough connectivity for management of the system.
Please register Cisco Nexus9000 Family devices promptly with your supplier.
Failure to register may affect response times for initial service calls. Nexus9000
devices must be registered to receive entitled support services.
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the
remaining dialogs.
Would you like to enter the basic configuration dialog (yes/no): y
Create another login account (yes/no) [n]: n
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : <<var_nexus_B_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Mgmt0 IPv4 address : <<var_nexus_B_mgmt0_ip>>
Mgmt0 IPv4 netmask : <<var_nexus_B_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway : <<var_nexus_B_mgmt0_gw>>
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
Type of ssh key you would like to generate (dsa/rsa) [rsa]:

```

```

Number of rsa key bits <1024-2048> [1024]: 2048
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address : <<var_global_ntp_server_ip>>
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut) [noshut]:
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
The following configuration will be applied:
password strength-check
switchname <<var_nexus_A_hostname>>
vrf context management
ip route 0.0.0.0/0 <<var_nexus_B_mgmt0_gw>>
exit
no feature telnet
ssh key rsa 2048 force
feature ssh
ntp server <<var_global_ntp_server_ip>>
system default switchport
no system default switchport
copp profile strict
interface mgmt0
ip address <<var_nexus_B_mgmt0_ip>> <<var_nexus_B_mgmt0_netmask>>
no shutdown
Would you like to edit the configuration? (yes/no) [n]:
Use this configuration and save it? (yes/no) [y]:
[#####] 100%
Copy complete.

```

6.1.3 Enabling the Cisco Nexus 9000 Series Switch features and settings

On *both* Cisco Nexus A and Cisco Nexus B, to enable the IP switching feature and set the default spanning tree behaviors, complete the following steps (this example shows only Cisco Nexus A; do the same exact procedure for Cisco Nexus B):

1. On *each* Cisco Nexus 9000 Series Switch, enter configuration mode by running the following command:

```
N9K-A# config terminal
```

2. To enable the necessary features, run the following commands:

```

N9K-A(config)# feature udld
N9K-A(config)# feature lacp
N9K-A(config)# feature vpc

```

3. Configure the spanning tree and save the running configuration to start by running the following commands:

```

N9K-A(config)# spanning-tree port type network default
N9K-A(config)# spanning-tree port type edge bpduguard default
N9K-A(config)# spanning-tree port type edge bpdufilter default
N9K-A(config)# copy run start

```

6.1.4 Creating VLANs for VersaStack traffic

This section describes how to create the VLANs for VersaStack traffic.

Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), run the following commands on *both* switches when in configuration mode:

```
N9K-A(config)# vlan 30
N9K-A(config)# name vMotion
N9K-A(config)# vlan 40
N9K-A(config)# name WinClus
N9K-A(config)# vlan 50
N9K-A(config)# name WinCSV
N9K-A(config)# vlan 60
N9K-A(config)# name Backup
```

6.1.5 Configuring the Virtual Port Channel Domain

This section describes how to create the Virtual Port Channel Domain.

Cisco Nexus A

To configure virtual port channels (vPCs) for Cisco Nexus A, complete the following steps:

1. From the global configuration mode, create a vPC domain by running the following command:

```
N9K-A(config-vpc-domain)# vpc domain 101
```

2. Make Nexus 9000 A the primary vPC peer by defining a low priority value by running the following command:

```
N9K-A(config-vpc-domain)# role priority 10
```

3. Use the management interfaces on the supervisors of the Nexus 9000 A to establish a keepalive link by running the following command:

```
N9K-A(config-vpc-domain)# peer-keepalive destination 192.168.10.32 source
192.168.10.31
```

4. Enable the features for this vPC domain by running the following commands:

```
N9K-A(config-vpc-domain)# peer-switch
N9K-A(config-vpc-domain)# delay restore 150
N9K-A(config-vpc-domain)# peer-gateway
N9K-A(config-vpc-domain)# ip arp synchronize
N9K-A(config-vpc-domain)# auto-recovery
```

Cisco Nexus B

To configure vPCs for Cisco Nexus B, complete the following steps:

1. From the global configuration mode, create a vPC domain by running the following command:

```
N9K-B(config-vpc-domain)# vpc domain 101
```

2. Make Nexus 9000 B the primary vPC peer by defining a low priority value by running the following command:

```
N9K-B(config-vpc-domain)# role priority 20
```

3. Use the management interfaces on the supervisors of Nexus 9000 B to establish a keepalive link by running the following command:

```
N9K-B(config-vpc-domain)# peer-keepalive destination 192.168.10.31 source
192.168.10.32
```

4. Enable the features for this vPC domain by running the following commands:

```
N9K-B(config-vpc-domain)# peer-switch
N9K-B(config-vpc-domain)# delay restore 150
N9K-B(config-vpc-domain)# peer-gateway
N9K-B(config-vpc-domain)# ip arp synchronize
N9K-B(config-vpc-domain)# auto-recovery
```

6.1.6 Configuring network interfaces for the vPC peer links

This section describes how to configure the network interfaces for the vPC peer links.

Cisco Nexus A

To configure the network interfaces for the vPC peer links for Cisco Nexus A, complete the following steps:

1. Define a port description for the interfaces connecting to vPC Peer N9K-B by running the following commands:

```
N9K-A(config)# interface eth1/47
N9K-A(config-if)# description vPC Peer N9K-B:1/47
N9K-A(config-if)# interface eth1/48
N9K-A(config-if)# description vPC Peer N9K-B:1/48
```

2. Apply a port channel to both vPC peer links and start the interfaces by running the following commands:

```
N9K-A(config-if)# interface eth1/47,eth1/48
N9K-A(config-if-range)# channel-group 10 mode active
N9K-A(config-if-range)# no shutdown
```

3. Define a description for the port channel connecting to N9K-B by running the following commands:

```
N9K-A(config-if-range)# interface Po10
N9K-A(config-if)# description vPC peer-link
```

4. Make the port channel a switchport and configure a trunk to allow all VLANs by running the following commands:

```
N9K-A(config-if)# switchport
N9K-A(config-if)# switchport mode trunk
N9K-A(config-if)# switchport trunk allowed vlan all
```

5. Make this port channel the vPC peer link and start it by running the following commands:

```
N9K-A(config-if)# vpc peer-link
N9K-A(config-if)# no shutdown
```

Cisco Nexus B

To configure the network interfaces for the vPC peer links for Cisco Nexus B, complete the following steps:

1. Define a port description for the interfaces connecting to vPC Peer N9K-A by running the following commands:

```
N9K-B(config-vpc-domain)# interface eth1/47
N9K-B(config-if)# description vPC Peer N9K-A:1/47
N9K-B(config-if)# interface eth1/48
N9K-B(config-if)# description vPC Peer N9K-A:1/48
```

2. Apply a port channel to both vPC peer links and start the interfaces by running the following commands:


```
N9K-B(config-if)# interface eth1/47,eth1/48
N9K-B(config-if-range)# channel-group 10 mode active
N9K-B(config-if-range)# no shutdown
```
3. Define a description for the port channel connecting to N9K-A by running the following commands:


```
N9K-B(config-if-range)# interface Po10
N9K-B(config-if)# description vPC peer-link
```
4. Make the port channel a switchport and configure a trunk to allow all VLANs by running the following commands:


```
N9K-B(config-if)# switchport
N9K-B(config-if)# switchport mode trunk
N9K-B(config-if)# switchport trunk allowed vlan all
```
5. Make this port channel the vPC peer link and start it by running the following commands:


```
N9K-B(config-if)# vpc peer-link
N9K-B(config-if)# no shutdown
```
6. Verify the status of vPC by running **sh vpc brief**:


```
N9K-B(config-if)# sh vpc brief
Legend:
(*) - local vpc is down, forwarding via vPC peer-link
```

```
vPC domain id           : 10
Peer status             : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : Consistency Check Not Performed
vPC role                : secondary
Number of vPCs configured : 0
Peer Gateway            : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status    : Enabled (timeout = 240 seconds)
```

```
vPC Peer-link status
-----
id  Port Status Active vlans
--  -
1   Po10 up      30,40,50,60
```

6.1.7 Configuring network interfaces to the Cisco UCS Fabric Interconnect

This section describes how to configure the network interfaces to the Cisco UCS fabric interconnect.

Cisco Nexus A

To configure the network interfaces to the Cisco UCS fabric interconnect for Cisco Nexus A, complete the following steps:

1. Define a description for the port channel connecting to FI-A by running the following commands:

```
N9K-A(config-if)# interface Po13
N9K-A(config-if)# description to FI-A
```

2. Make the port channel a switchport and configure a trunk to allow all VLAN traffic by running the following commands:

```
N9K-A(config-if)# switchport
N9K-A(config-if)# switchport mode trunk
N9K-A(config-if)# switchport trunk allowed vlan all
```

3. Make the port channel and associated interfaces into spanning tree edge ports by running the following command:

```
N9K-A(config-if)# spanning-tree port type edge trunk
```

4. Set the MTU to be 9216 to support jumbo frames by running the following command:

```
N9K-A(config-if)# mtu 9216
```

5. Make a vPC port channel and start it by running the following commands:

```
N9K-A(config-if)# vpc 13
N9K-A(config-if)# no shutdown
```

6. Define a port description for the interface connecting to FI-A by running the following commands:

```
N9K-A(config-if)# interface eth1/25
N9K-A(config-if)# description FI-A:1/25
```

7. Start the interface by running the following commands:

```
N9K-A(config-if)# channel-group 13 mode active
N9K-A(config-if)# no shutdown
```

8. Define a description for the port channel connecting to FI-B by running the following commands:

```
N9K-A(config-if)# interface Po14
N9K-A(config-if)# description to FI-B
```

9. Make the port channel a switchport and configure a trunk to allow all VLAN traffic by running the following commands:

```
N9K-A(config-if)# switchport
N9K-A(config-if)# switchport mode trunk
N9K-A(config-if)# switchport trunk allowed vlan all
```

10. Make the port channel and associated interfaces into spanning tree edge ports by running the following command:

```
N9K-A(config-if)# spanning-tree port type edge trunk
```

11. Set the MTU to be 9216 to support jumbo frames by running the following command:

```
N9K-A(config-if)# mtu 9216
```

12. Make a vPC port channel and start it by running the following commands:

```
N9K-A(config-if)# vpc 14
N9K-A(config-if)# no shutdown
```


13. Define a port description for the interface connecting to FI-B by running the following commands:

```
N9K-A(config-if)# interface eth1/26
N9K-A(config-if)# description FI-B:1/26
```

14. Start the interface by running the following commands:

```
N9K-A(config-if)# channel-group 14 mode active
N9K-A(config-if)# no shutdown
N9K-A(config-if)# copy run start
[#####] 100%
Copy complete.
```

Cisco Nexus B

To configure the network interfaces to the Cisco UCS fabric interconnect for Cisco Nexus B, complete the following steps:

1. Define a description for the port channel connecting to FI-B by running the following commands:

```
N9K-B(config-if)# interface Po14
N9K-B(config-if)# description to FI-B
```

2. Make the port channel a switchport and configure a trunk to allow all VLAN traffic by running the following commands:

```
N9K-B(config-if)# switchport
N9K-B(config-if)# switchport mode trunk
N9K-B(config-if)# switchport trunk allowed vlan all
```

3. Make the port channel and associated interfaces into spanning tree edge ports by running the following command:

```
N9K-B(config-if)# spanning-tree port type edge trunk
```

4. Set the MTU to be 9216 to support jumbo frames by running the following command:

```
N9K-B(config-if)# mtu 9216
```

5. Make a vPC port channel and start it by running the following commands:

```
N9K-B(config-if)# vpc 14
N9K-B(config-if)# no shutdown
```

6. Define a port description for the interface connecting to FI-B by running the following commands:

```
N9K-B(config-if)# interface eth1/25
N9K-B(config-if)# description FI-B:1/25
```

7. Start the interface by running the following commands:

```
N9K-B(config-if)# channel-group 14 mode active
N9K-B(config-if)# no shutdown
```

8. Define a description for the port channel connecting to FI-A by running the following commands:

```
N9K-B(config-if)# interface Po13
N9K-B(config-if)# description to FI-A
```

9. Make the port channel a switchport and configure a trunk to allow all VLAN traffic by running the following commands:


```
N9K-B(config-if)# switchport
N9K-B(config-if)# switchport mode trunk
N9K-B(config-if)# switchport trunk allowed vlan all
```
10. Make the port channel and associated interfaces into spanning tree edge ports by running the following command:


```
N9K-B(config-if)# spanning-tree port type edge trunk
```
11. Set the MTU to be 9216 to support jumbo frames by running the following command:


```
N9K-B(config-if)# mtu 9216
```
12. Make a vPC port channel and start it by running the following commands:


```
N9K-B(config-if)# vpc 13
N9K-B(config-if)# no shutdown
```
13. Define a port description for the interface connecting to FI-A by running the following commands:


```
N9K-B(config-if)# interface eth1/26
N9K-B(config-if)# description FI-A:1/26
```
14. Start the interface by running the following commands:


```
N9K-B(config-if)# channel-group 13 mode active
N9K-B(config-if)# no shutdown
N9K-B(config-if)# copy run start
[#####] 100%
Copy complete.
```

6.1.8 Linking in to an existing network infrastructure

Depending on the available network infrastructure, you can use several methods and features to uplink to the VersaStack environment. If an existing Cisco Nexus environment is present, Cisco recommends using vPCs to uplink the Cisco Nexus 9000 Series Switches that are included in the VersaStack environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment.



IBM Storwize V7000 storage configuration

This chapter describes the steps that are necessary to configure the Storwize V7000 storage system in the VersaStack environment.

7.1 Secure web access to the IBM Storwize V7000 service and management GUI

Browser access to all system and service IPs is automatically configured to connect securely by using HTTPS and SSL. Attempts to connect through HTTP are redirected to HTTPS.

The system generates its own self-signed SSL certificate. On first connection to the system, your browser might present a security exception because it does not trust the signer; you should allow the connection to proceed.

Figure 7-1 shows the rear of the Storwize V7000 Gen2 storage system.




Figure 7-1 Rear of the Storwize V7000 Gen2 storage system

7.2 IBM Storwize V7000 initial configuration setup

To accomplish the initial configuration setup of the IBM Storwize V7000 storage system, complete the following steps:

1. Configure an Ethernet port of a PC or notebook to allow DHCP to configure its IP address and DNS settings.
2. Connect an Ethernet cable from the PC or notebook Ethernet port to the Ethernet port labeled "T" on the rear of either node canister in the Storwize V7000 control enclosure.
3. A few moments after the connection is made, the node uses DHCP to configure the IP address and DNS settings of the PC or notebook.

Note: This step will likely disconnect you from any other network connections that you have on the PC or notebook. If you do not have DHCP on your PC or notebook, you can manually configure it with the following network settings:

- ▶ IPv4 address: 192.168.0.2
- ▶ Mask: 255.255.255.0
- ▶ Gateway: 192.168.0.1
- ▶ DNS: 192.168.0.1

4. Open a browser and go to `https://install`, which opens the initialization wizard. Figure 7-2 shows the Welcome window for the wizard.




Figure 7-2 System initialization wizard

5. When asked how the node will be used, select **As the first node in a new system**.
6. Follow the instructions that are presented by the initialization tool to configure the system with a management IP address of `<<var_cluster_mgmt_ip>>`, `<<var_cluster_mgmt_mask>>`, and `<<var_cluster_mgmt_gateway>>`.
7. After you complete the initialization process, disconnect the cable between the PC and notebook and the technician port as directed, and reconnect to your network with your previous settings.
8. Click **OK** to redirect your browser to the management GUI at the IP address you configured.

Note: You might have to wait up to 5 minutes for the management GUI to start and become accessible.

9. Read and agree to the license agreement by selecting the check box next to it, and then click **Next** to proceed.
10. Log in as superuser with a password of `passw0rd`.
11. Change the password for superuser, and then click **Log In**.

12. Figure 7-3 shows the Storwize V7000 welcome window, which is the first window of System Setup. Click **Next**.

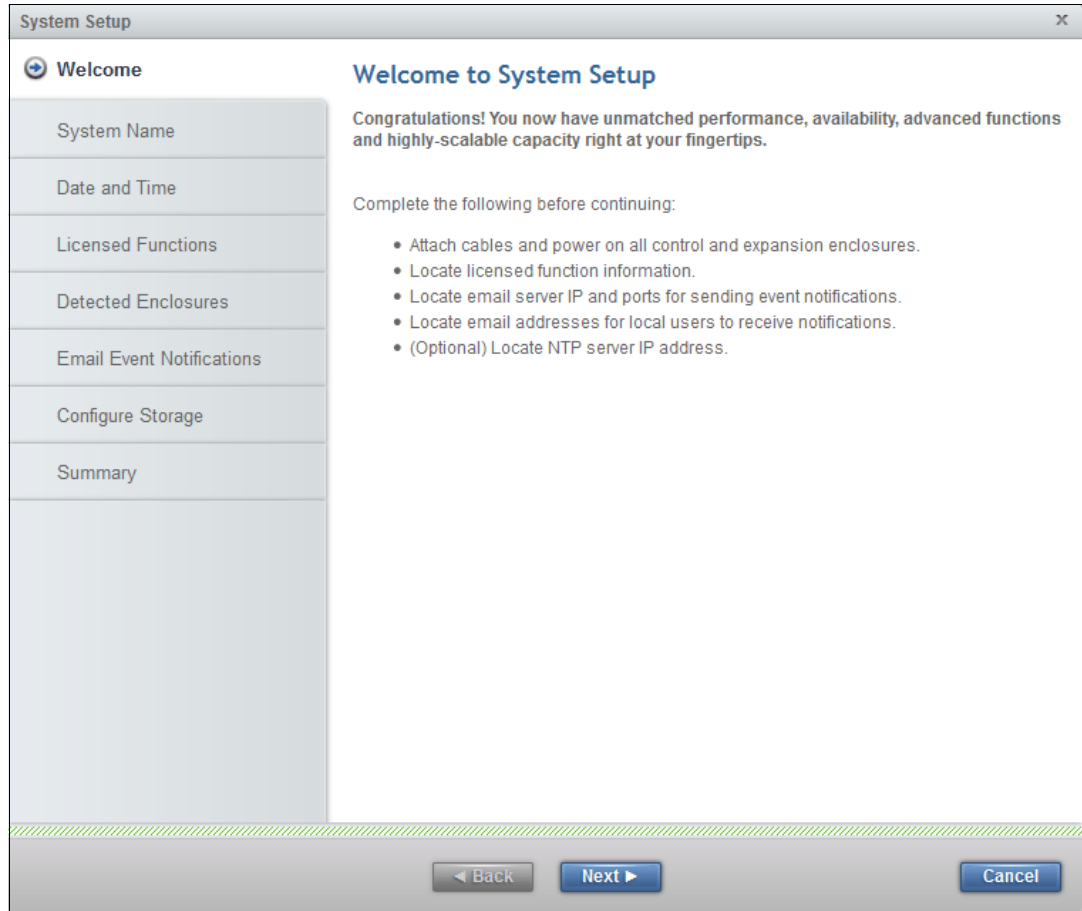


Figure 7-3 StorwizeV7000 welcome window

13. The window that is shown in Figure 7-4 opens, where you can change the system name. Change the system name if required, and click **Apply and Next**.

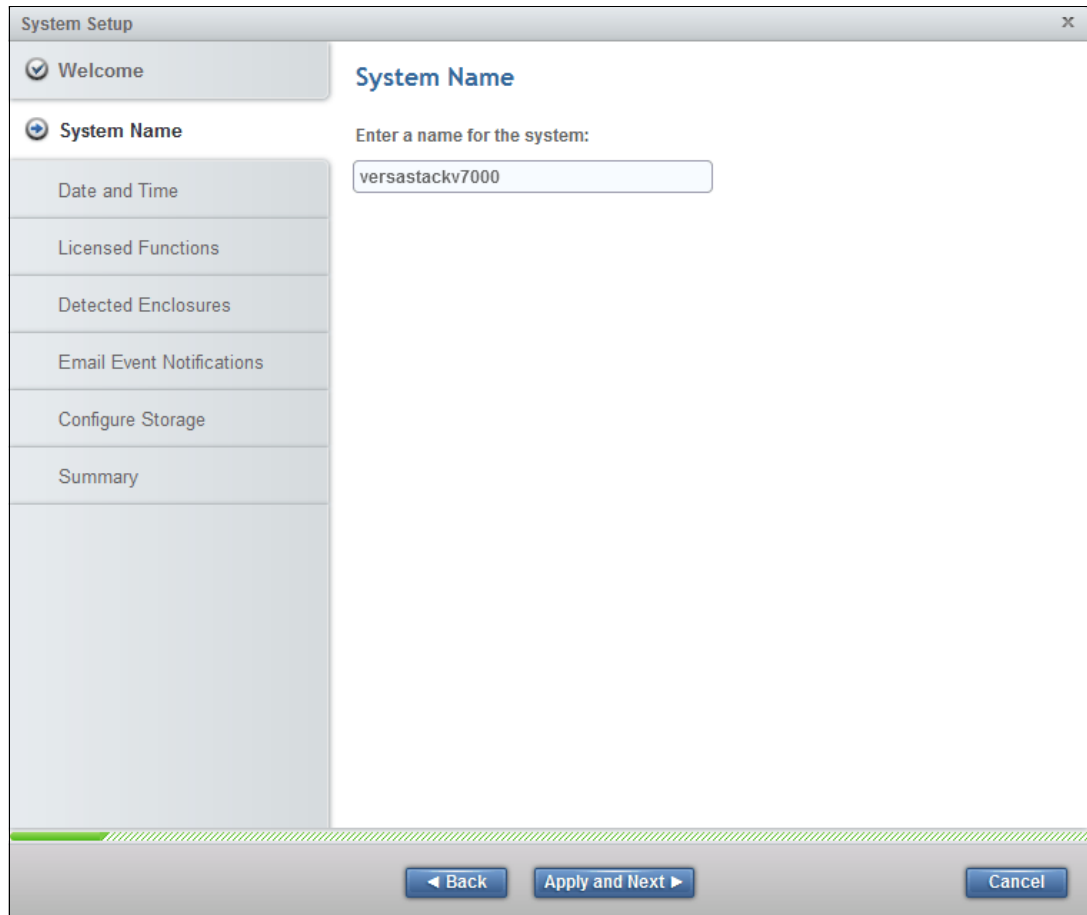
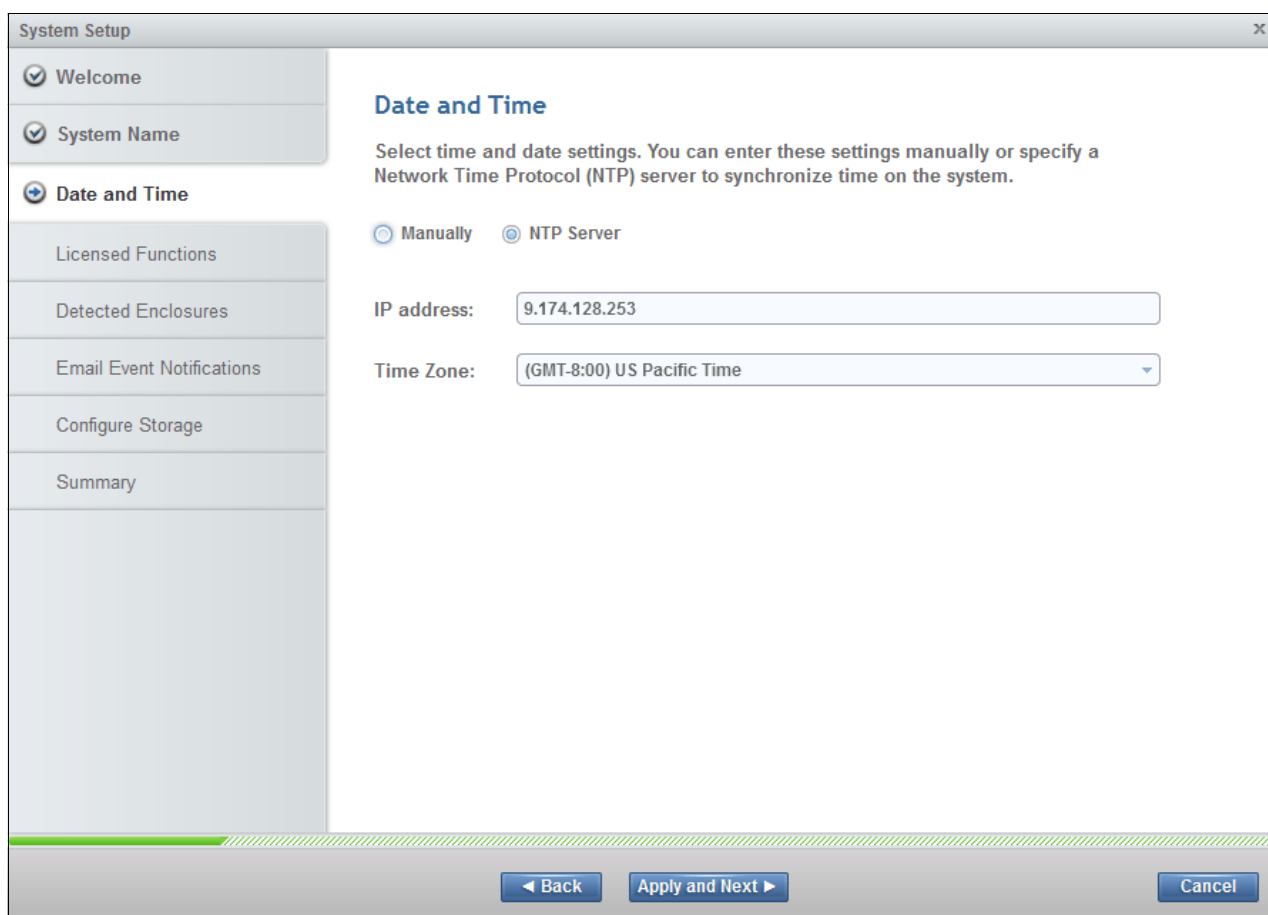


Figure 7-4 System Name window

14. The window that is shown in Figure 7-5 opens, where you can set the data and time manually or configure an NTP server. Select **NTP Server** and enter the NTP server address. Click **Apply and Next** to view and close the Tasks Completed window.



The screenshot shows a web-based configuration window titled "System Setup". On the left is a vertical navigation menu with the following items: "Welcome" (checked), "System Name" (checked), "Date and Time" (selected with a blue arrow), "Licensed Functions", "Detected Enclosures", "Email Event Notifications", "Configure Storage", and "Summary". The main content area is titled "Date and Time" and contains the following text: "Select time and date settings. You can enter these settings manually or specify a Network Time Protocol (NTP) server to synchronize time on the system." Below this text are two radio buttons: "Manually" (unselected) and "NTP Server" (selected). There are two input fields: "IP address:" with the value "9.174.128.253" and "Time Zone:" with a dropdown menu showing "(GMT-8:00) US Pacific Time". At the bottom of the window are three buttons: "Back", "Apply and Next", and "Cancel".

Figure 7-5 Date and Time setting window

15. Validate that the enclosures that you connected are detected. If there are any discrepancies, review 5.1, "VersaStack cabling" on page 24. Click **Apply and Next**. View and close the Tasks Completed window.
16. The window that is shown in Figure 7-6 on page 47 opens, where you can specify the number of each license that you possess. Licenses are granted based on the number of enclosures. Enter the number for the licensed functions, and then click **Apply and Next**. View and close the Task Completed window.

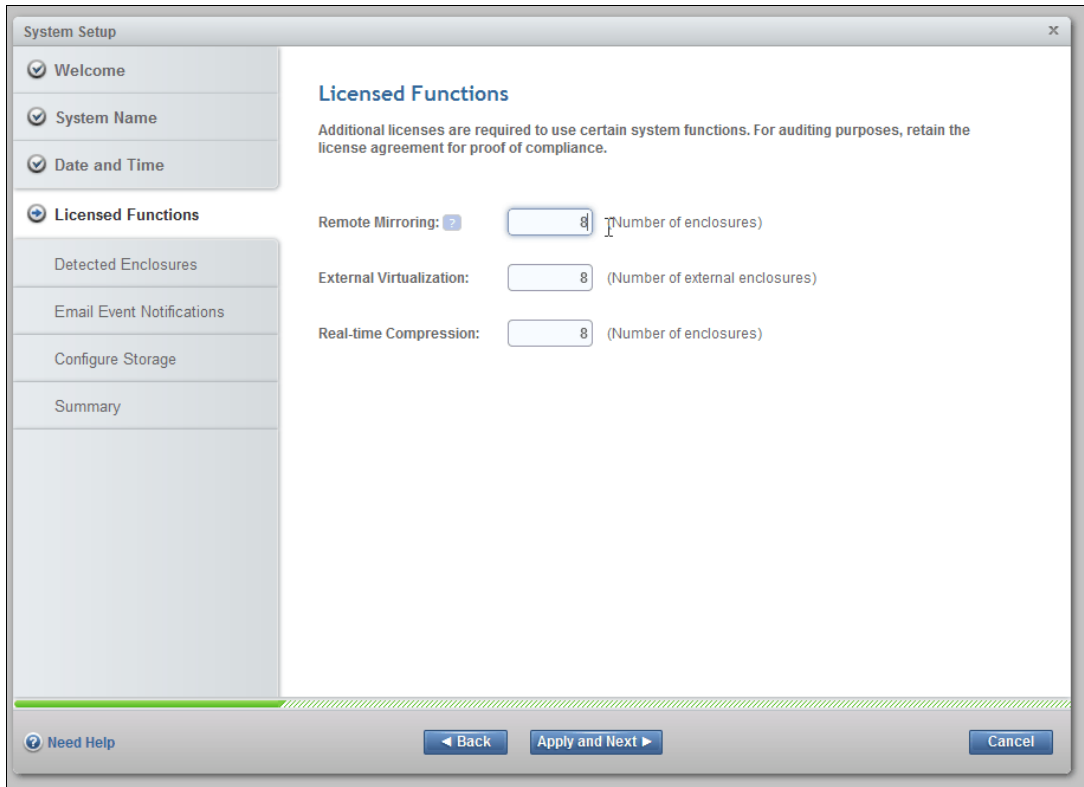


Figure 7-6 Licensed Functions

17. The window that is shown in Figure 7-7 opens, where you can set up event notifications through email. Select **Yes** to enter the email information for event notification.

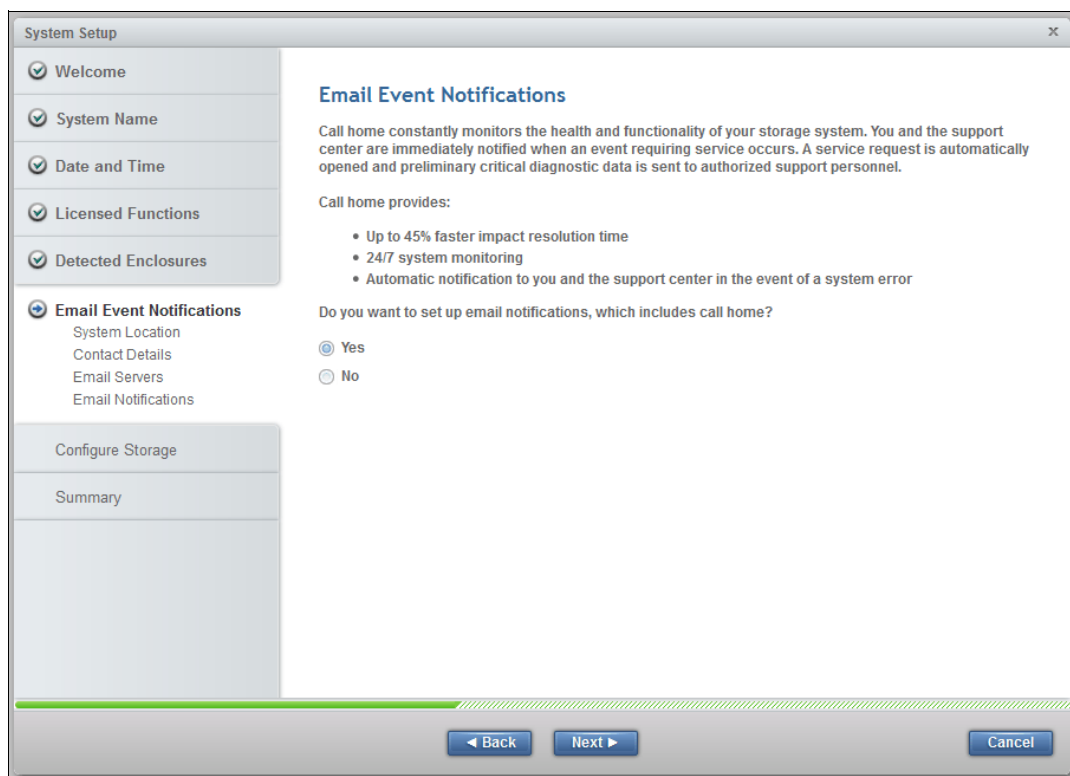


Figure 7-7 Email Event Notifications window

18. Enter the system location and contact details of <<var_org>>, <<var_street_address>>, <<var_city>>, <<var_state>>, <<var_zip>>, and <<var_country_code>>, and then click **Apply and Next**. View and close the Tasks Completed window.
19. Insert the contact details of <<var_contact_name>>, <<var_email_contact>>, <<var_admin_phone>>, and <<var_city>>, click **Apply and Next**, and click **Close**.
20. Enter the email server IP address <<var_mailhost_ip>> and change the port if necessary, and then click **Apply and Next**. View and close the Tasks Completed window. Click **Apply and Next**.
21. In the Call Home validation window, click **Apply and Next**, and then click **Close**.

22. Enter the email addresses for all administrators that should be notified when issues occur and any other parties that need information or inventory by using <<var_email_contact>>. Click **Apply and Next**. Review and close the Tasks Completed window. Figure 23 on page 50 shows where you can specify whom receives each email and for what they receive notifications.

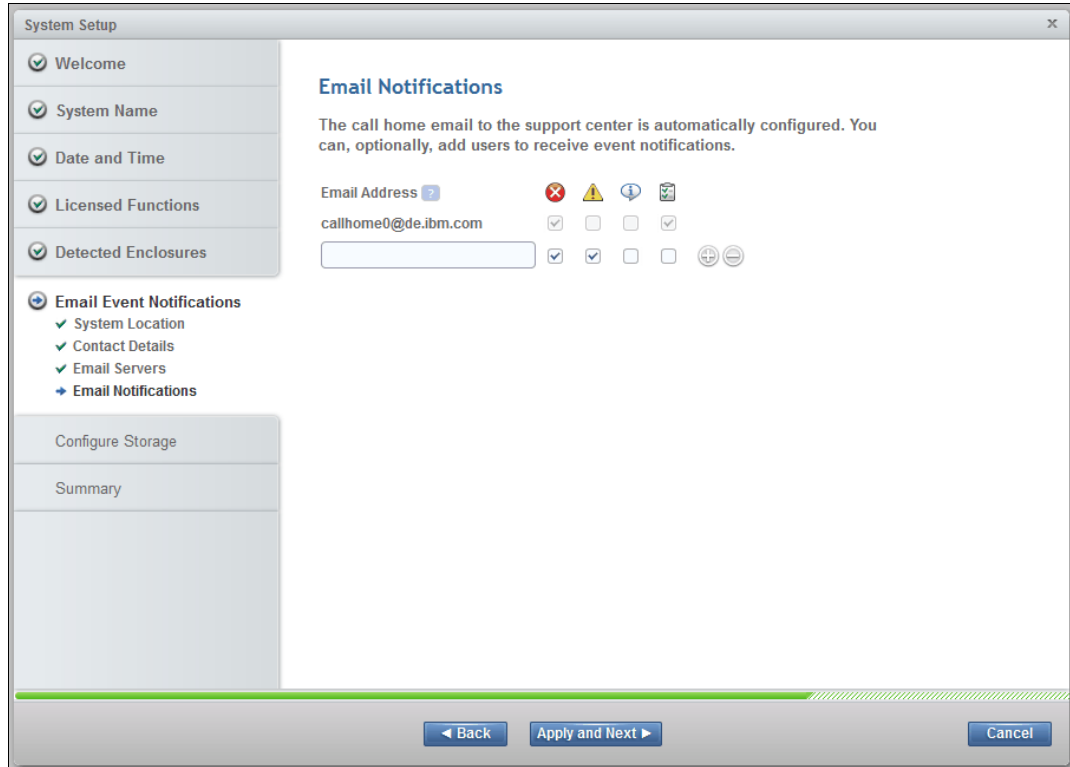


Figure 7-8 Email Notifications window

23. The window that is shown in Figure 7-9 opens, where you choose to configure your external storage automatically now or to wait until later. Select **Configure storage now** and then click **Next**.

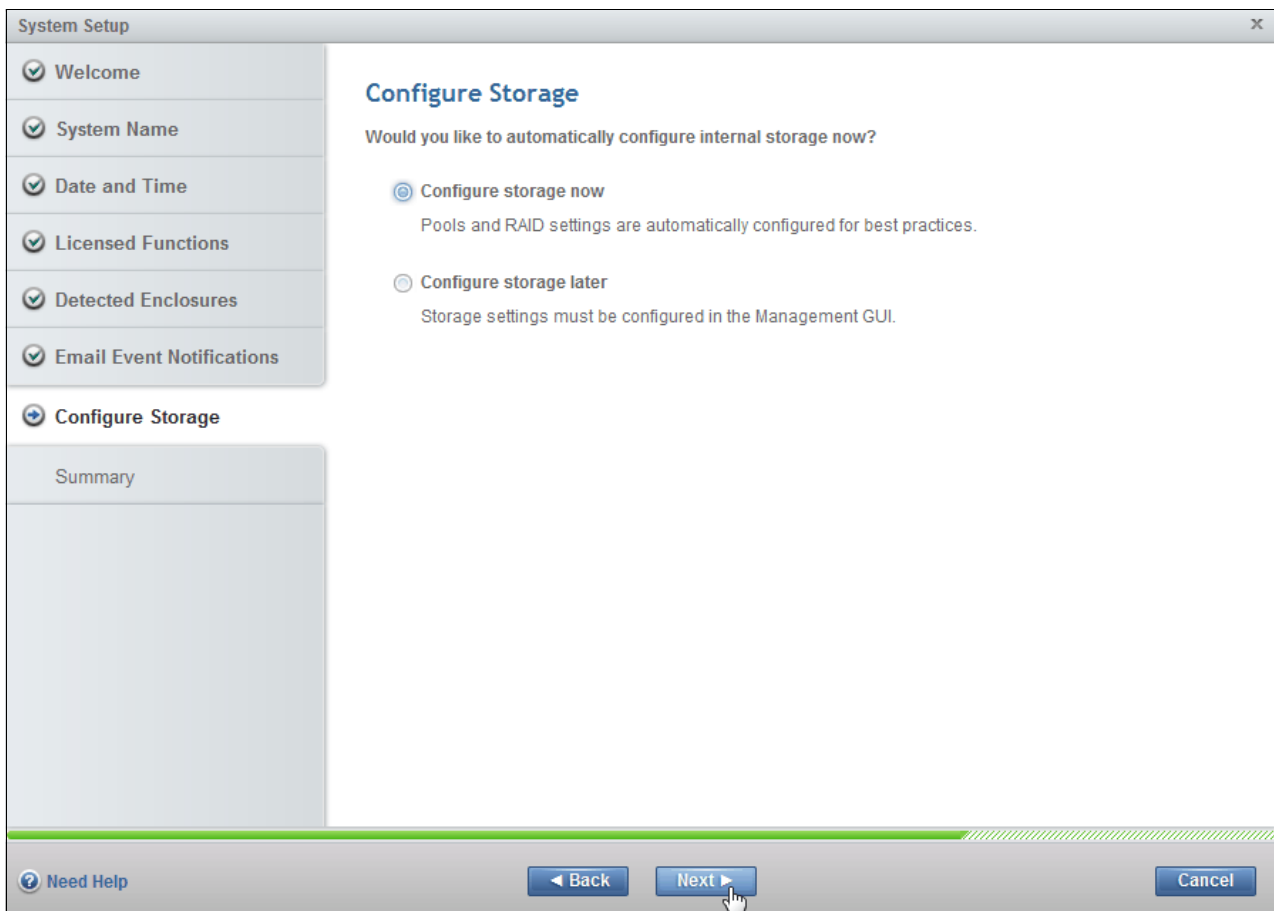


Figure 7-9 Configure Storage window

24. Review the summary and click **Finish**. View and close the Tasks Completed window.

25. Click **Cancel** on the Create Hosts window, as these hosts are created after the Cisco Fabric Interconnects are configured. Optionally, you can view an introductory tour of the management GUI by using the link.

26. In the window that is shown in Figure 7-10 on page 51, click the Settings icon in the lower left of the window, and then select **Network**.

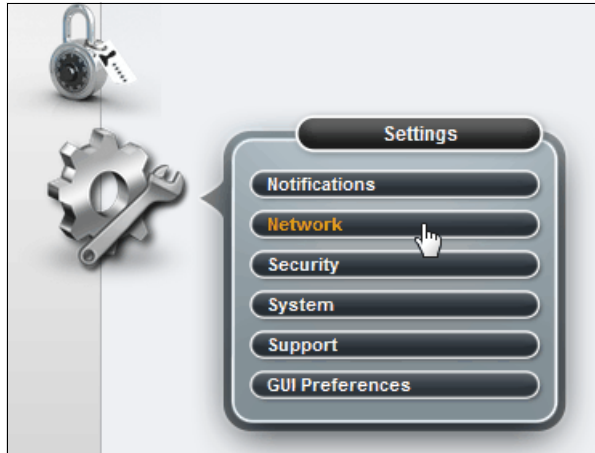


Figure 7-10 Network menu

27. In the window that is shown in Figure 7-11, click the **Service IP Addresses** menu and click port 1 to enter the node management port IP address (`<<var_node01_mgmt_ip>>`), netmask (`<<var_node01_mgmt_mask>>`), and gateway (`<<var_node01_mgmt_gateway>>`). Click **OK** and then click **Close**.



Figure 7-11 Set the service IP for port 1

28. Click the **Node Canister** drop-down menu item, change the selection to **Left**, and click the port 1 picture to enter the node management port IP address (`<<var_node02_mgmt_ip>>`), netmask (`<<var_node02_mgmt_mask>>`), and gateway (`<<var_node02_mgmt_gateway>>`). Click **OK** and then **Close**.

29. In the left menu, hover over each of the icons to become familiar with the GUI options.

30. To create a separate administrator user, click the lock icon in the left pane, which opens the Users pane. Click **Create User** and enter the user name (`<<var_admin>>`) and a password (`<<var_password>>`). Click **Create** and then **Close**.

31. Log off by selecting the superuser account in the upper right pane and clicking **Log Out**. Log back in by using the admin account that you created.

32. Click the fourth icon from the top in the left pane to open the Volumes pane. Click **Create Volume** at the upper left to open the Create Volume wizard.

33. Click **Thin-Provision** in the Select a Preset section. Click **mdiskgroup0** in the Select a Pool section.

34. To create SAN boot volumes for ESX, in the Volume Details section, enter the following values:

- Quantity: 2
- Capacity: 32 GiB
- Name: vm_host_boot
- Change the starting number to 1,

Click **Create**, and then click **Close**.

Figure 7-12 shows the Create Volumes window, which shows the creation of two vm-host-boot volumes.

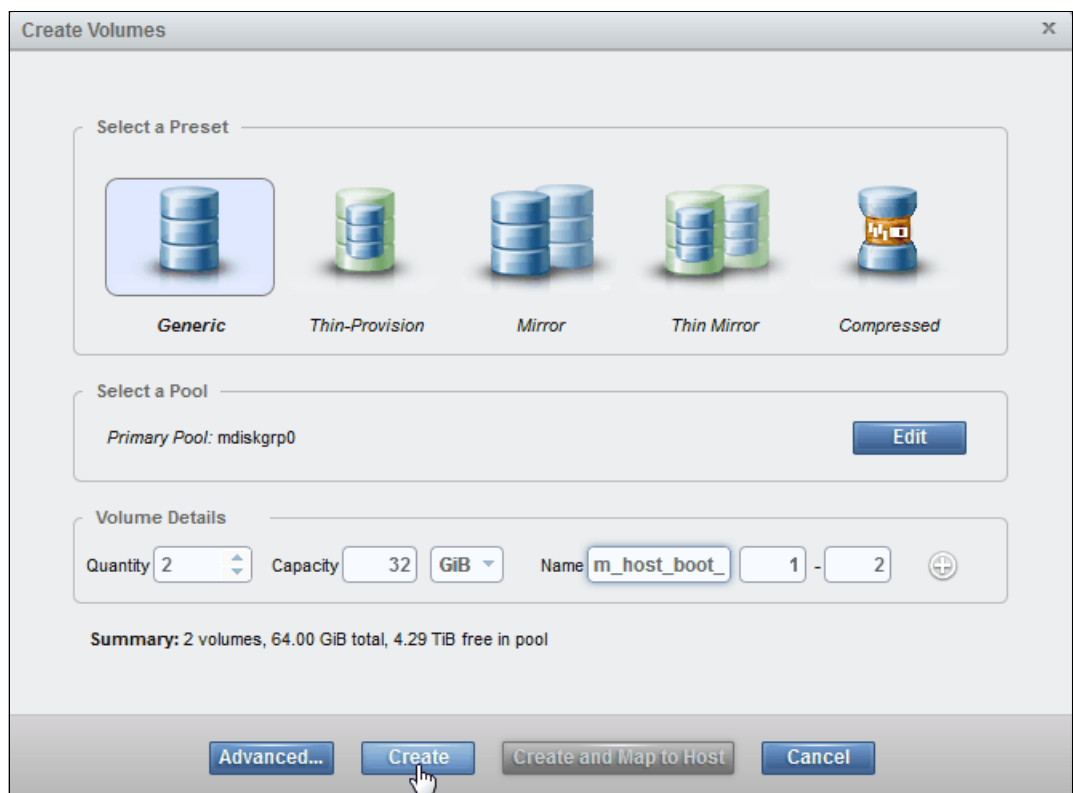


Figure 7-12 Create two vm-host-boot volumes

35. To create a generic VMFS datastore for virtual machines, click **Create Volume**, select **Generic** or another preset that you want, and select **mdiskgroup0** for the pool. Enter the following values:

- Quantity: 1
- Capacity: 2048 GiB
- Name: infra_datastore_1

Click **Create** and then click **Close**.

36. To create a thin-provisioned VMFS datastore, click **Volume**, select **Thin-Provision**, and select **mdiskgroup0** for the pool.

Figure 7-13 shows the Create Volumes window, which shows the creation of a thin-provisioned datastore.

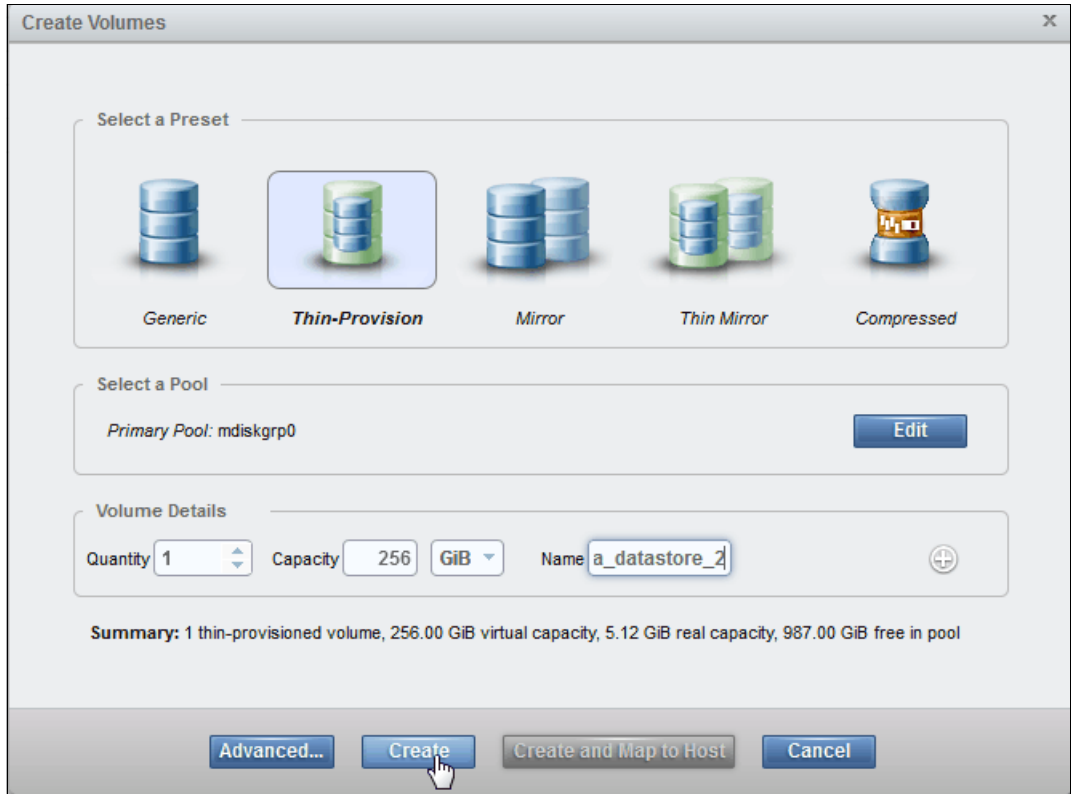


Figure 7-13 Create a thin-provisioned datastore

37. Create the rest of volumes by using the values that are shown in Table 7-1.

Table 7-1 Table of volume names and sizes

Volume name	Size (GiB)	FlashCopy
sql_rdm_data	256	Yes
sql_rdm_log	64	Yes
sql_rdm_quorum	1	No
vm_datastore_1	1024	Yes
vm_datastore_2	256 (thin-provisioned)	Yes
vm_host_boot_1	32	No
vm_host_boot_2	32	No

Figure 7-14 shows the final result of volume creation.

Name	State	Capacity	Pool	Host Mappings	UID
infra_datastore_1	✓ Online (formatting)	1.00 TiB	mdiskgrp0	No	6005076802C480B2C400000
infra_datastore_2	✓ Online	256.00 GiB	mdiskgrp0	No	6005076802C480B2C400000
sp_datastore_1	✓ Online (formatting)	2.00 TiB	mdiskgrp0	No	6005076802C480B2C400000
sql_rdm_data	✓ Online (formatting)	256.00 GiB	mdiskgrp0	No	6005076802C480B2C400000
sql_rdm_log	✓ Online (formatting)	64.00 GiB	mdiskgrp0	No	6005076802C480B2C400000
sql_rdm_quorum	✓ Online (formatting)	1.00 GiB	mdiskgrp0	No	6005076802C480B2C400000
vm_host_boot_1	✓ Online (formatting)	32.00 GiB	mdiskgrp0	No	6005076802C480B2C400000
vm_host_boot_2	✓ Online (formatting)	32.00 GiB	mdiskgrp0	No	6005076802C480B2C400000

Figure 7-14 Final result of volume creation

Note: You might want to create a separate VMFS volume for the swap file. To do so, click **Create Volume**, select **Thin-Provision**, and enter the following details:

- ▶ mdiskgrp0
- ▶ Quantity: 1
- ▶ Capacity: 100 GiB
- ▶ Name: infra_swap

Click **Create** and then **Close**.

38. Collect the information for the Fibre Channel WWPNs that are used later for SAN boot by selecting the cog icon in the left pane, which opens the Settings menu. Click **Network**, as shown in Figure 7-15.

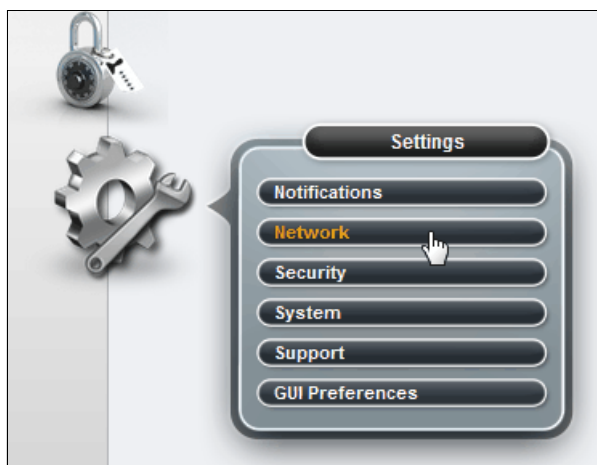


Figure 7-15 Access the FC ports menu

39. Select Fibre Channel Ports in the lower left menu, and then expand the ports 1 - 4 to show the WWPNs, as shown in Figure 7-16.

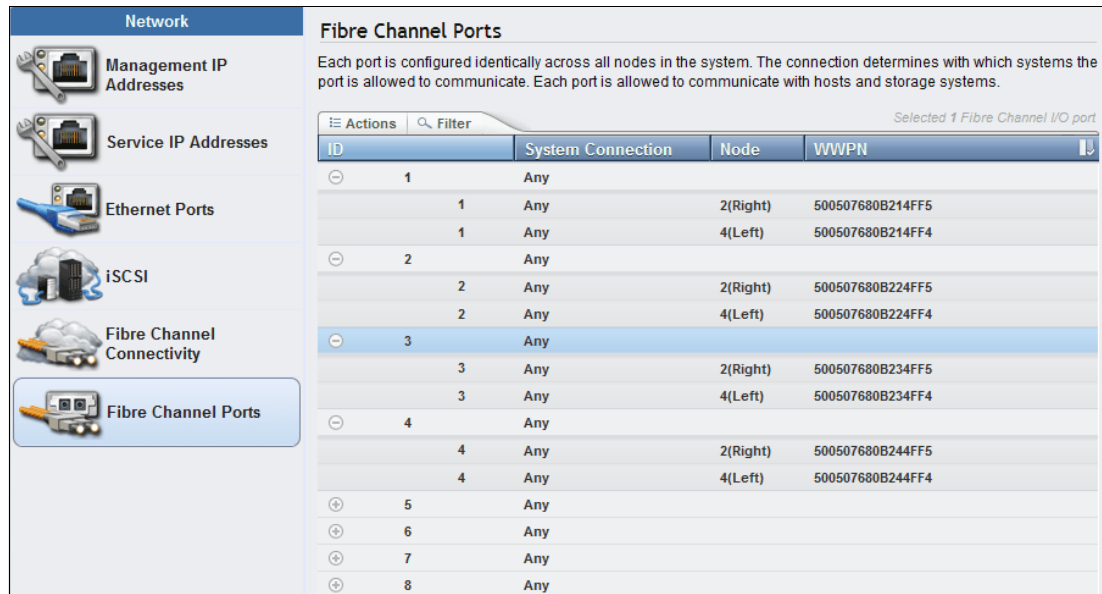


Figure 7-16 Find the FC WWPNs

Note: In this example configuration, we deployed a bloc-only system by using all four Fibre Channel ports. The right node is the configuration node with an ID 2 and the left node is the partner node with an ID of 4.

40. Enter the WWPN numbers that are shown in Table 7-2. They will be required later when you configure the FC zones. The data for the hosts will be collected later in this document. For example, ID 3 on node 1 in Figure 7-16 corresponds to FC_Node1-3 in the spreadsheet.

Table 7-2 WWPNs for IBM Storwize V7000 Gen2 storage system

Source	Switch target	Variable	Customer WWPN
FC_Node1-1	FI-A	<i>var_wwpn_Node1-1-F I-A</i>	
FC_Node1-2	FI-A	<i>var_wwpn_Node1-2-F I-A</i>	
FC_Node1-3	FI-B	<i>var_wwpn_Node1-1-F I-B</i>	
FC_Node1-4	FI-B	<i>var_wwpn_Node1-2-F I-B</i>	
FC_Node2-1	FI-A	<i>var_wwpn_Node2-1-F I-A</i>	
FC_Node2-2	FI-A	<i>var_wwpn_Node2-2-F I-A</i>	
FC_Node2-3	FI-B	<i>var_wwpn_Node2-1-F I-B</i>	

Source	Switch target	Variable	Customer WWPN
FC_Node2-4	FI-B	<i>var_wwpn_Node2-2-FI-B</i>	
vm-host-infra-01-a	FI-A	<i>var_wwpn_VM-Host-Infra-01-A</i>	
vm-host-infra-01-b	FI-B	<i>var_wwpn_VM-Host-Infra-01-B</i>	
vm-host-infra-02-a	FI-A	<i>var_wwpn_VM-Host-Infra-02-A</i>	
vm-host-infra-02-b	FI-B	<i>var_wwpn_VM-Host-Infra-02-B</i>	

The Storwize V7000 storage system is now configured.

For the examples in this book, we will not be using the Real-time Compression (RtC) feature or the extensive range of replication services. For more information about RtC, see *IBM Real-time Compression in IBM SAN Volume Controller and IBM Storwize V7000*, REDP-4859

For more information about replication services, see *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574.

For more information about the Storwize V7000 Gen2 storage system, see *Implementing the IBM Storwize V7000 Gen2*, SG24-8244.

For more information about Storwize software Version 7.4, see *Implementing the IBM Storwize V7000 V7.4*, SG24-7938.



Cisco Unified Computing System configuration

This chapter describes how to configure the Cisco Unified Computing System (Cisco UCS) for use in a VersaStack environment.

8.1 Performing the initial setup of Cisco UCS 6248 Fabric Interconnect for VersaStack environments

This section provides detailed procedures for configuring the Cisco UCS for use in a VersaStack environment. The steps are necessary to provision the Cisco UCS C-Series and B-Series servers and should be followed precisely to avoid improper configuration.

8.1.1 Cisco UCS 6248 A

To configure the Cisco UCS 6248 A server for use in a VersaStack environment, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect and complete the following prompts with the provided information:

```
Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? Setup
You have chosen to set up a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y
Enter the password for "admin": <<var_password>>
Enter the same password for "admin": <<var_password>>
Is this fabric interconnect part of a cluster (select 'no' for standalone)?
(yes/no) [n]: y
Which switch fabric (A|B): A
Enter the system name: <<var_ucs_clustername>>
Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>
Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>
IPv4 address of the default gateway: <<var_ucsa_mgmt_gateway>>
Cluster IPv4 address: <<var_ucs_cluster_ip>>
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: <<var_nameserver_ip>>
Configure the default domain name? y
Default domain name: <<var_dns_domain_name>>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
```

2. Review the settings that are output to the console. If they are correct, answer yes to apply and save the configuration.
3. Wait for the login prompt and make sure that the configuration process completes before proceeding. It can take approximately 3 - 5 minutes.

8.1.2 Cisco UCS 6248 B

To configure the Cisco UCS 6248 B server for use in a VersaStack environment, complete the following steps:

1. Power on the second module and connect to the console port on the second Cisco UCS 6248 fabric interconnect and complete the following prompts with the provided information:

```
Enter the configuration method: console
Installer has detected the presence of a peer Fabric interconnect. This Fabric
interconnect will be added to the cluster. Do you want to continue {y|n}? y
Enter the admin password for the peer fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>
```

Apply and save the configuration (select 'no' if you want to reenter)?
(yes/no): y

8.2 Cisco UCS for IBM Storwize V7000 storage system

This section describes the steps to install the Cisco UCS for Storwize V7000 storage system.

8.2.1 Logging in to Cisco UCS Manager

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and go to the Cisco UCS 6248 fabric interconnect cluster address.
2. Click the **Launch UCS Manager** link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password of <<var_password>>.
5. Click **Login** to log in to Cisco UCS Manager.
6. Enter the information for Anonymous Reporting if you want and click **OK**.

8.2.2 Upgrading Cisco UCS Manager software to Version 2.2(3d)

This book assumes the use of Cisco UCS Manager Software Version 2.2(3d). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to Version 2.2(3d), see the Cisco UCS Manager Install and Upgrade Guides, found at the following website:

<http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html>

8.2.3 Adding a block of IP addresses for KVM access

To create a block of IP addresses for server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps.

Note: This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Click **Pools** → **root** → **IP Pools** → **IP Pool ext-mgmt**.
3. In the Actions pane, select **Create Block of IP Addresses**.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information (<<var_In-band_mgmtblock_net>>).
5. Click **OK** to create the IP block.
6. Click **OK** in the confirmation message.

8.2.4 Adding a block of IPv4 addresses for KVM access

To create a block of IP addresses for server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, complete the following steps:

Note: This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select **Pools** → **root** → **IP Pools** → **IP Pool ext-mgmt**, as shown in Figure 8-1.




Figure 8-1 IP pool ext-mgmt

3. In the Actions pane, select **Create Block of IPv4 Addresses**.
4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information (<<var_In-band_mgmtblock_net>>).

Figure 8-2 on page 61 shows the creation of the block of IPv4 addresses.




Figure 8-2 Create a block of IPv4 addresses

5. Click **OK** to create the IP block.
6. Click **OK** in the confirmation message.

8.2.5 Synchronizing the Cisco UCS environment to NTP

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

1. In Cisco UCS Manager, click the **Admin** tab in the navigation pane.
2. Click **All** → **Timezone Management**, as shown in Figure 8-3.

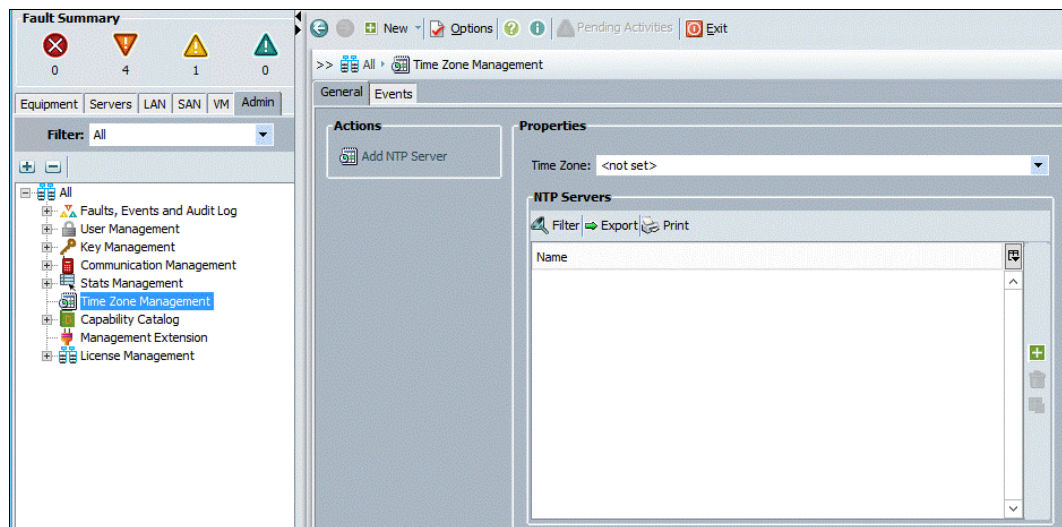


Figure 8-3 Timezone Management

3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click **Save Changes**, and then click **OK**.
5. Click **Add NTP Server**.

6. Enter `<<var_global_ntp_server_ip/FQDN>>` and click **OK**, as shown in Figure 8-4.




Figure 8-4 Add NTP server

7. Click **OK**.

Editing the chassis discovery policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and of additional fabric extenders for further Cisco UCS C-Series connectivity.

To modify the chassis discovery policy, complete the following steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane and select **Equipment** in the list on the left.
2. In the right pane, click the **Policies** tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to **Port Channel**, as shown in Figure 8-5.

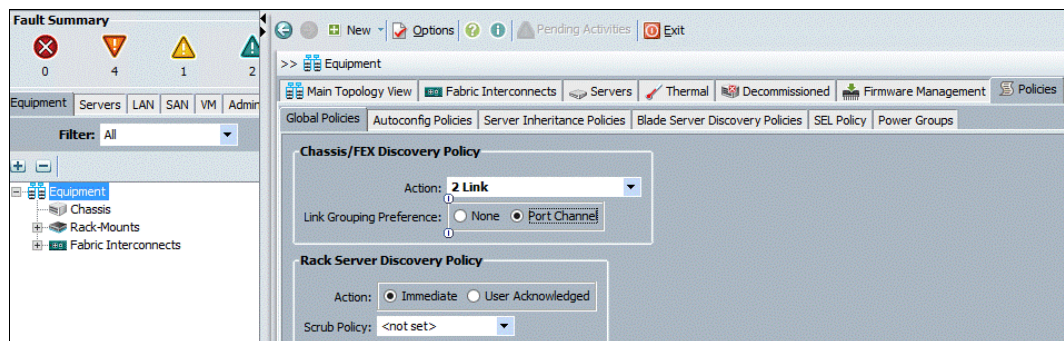


Figure 8-5 Link grouping preference

5. Click **Save Changes**.
6. Click **OK**.

8.2.6 Enabling the server and uplink ports

To enable the server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
2. Click **Equipment** → **Fabric Interconnects** → **Fabric Interconnect A (primary)** → **Fixed Module**.
3. Expand **Ethernet Ports**.
4. Select the ports that are connected to the chassis, right-click them, and select **Configure as Server Port**, as shown in Figure 8-6.

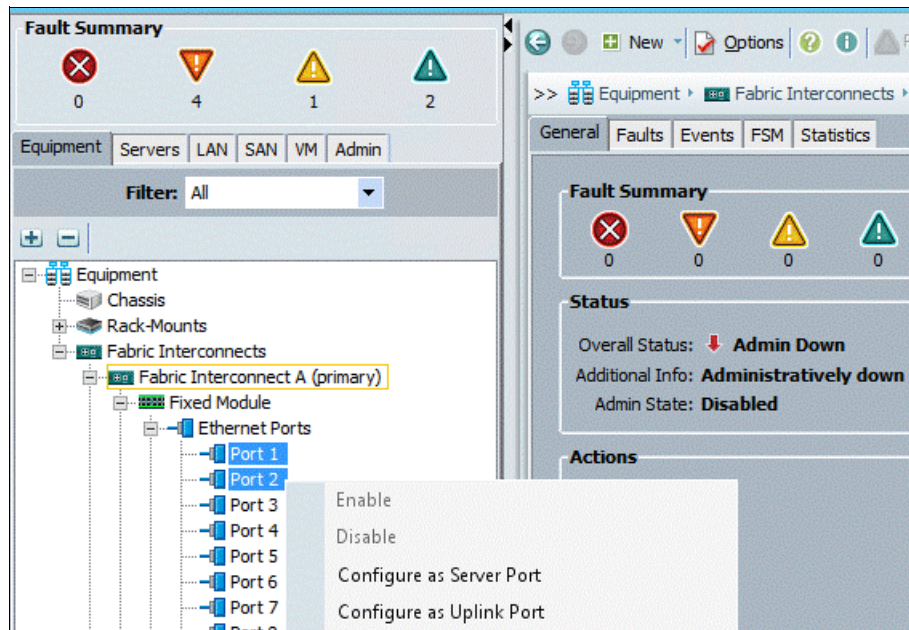


Figure 8-6 Configure as a server port

5. Click **Yes** to confirm server ports and click **OK**.
6. Verify that the ports that are connected to the chassis are now configured as server ports, as shown in Figure 8-7.

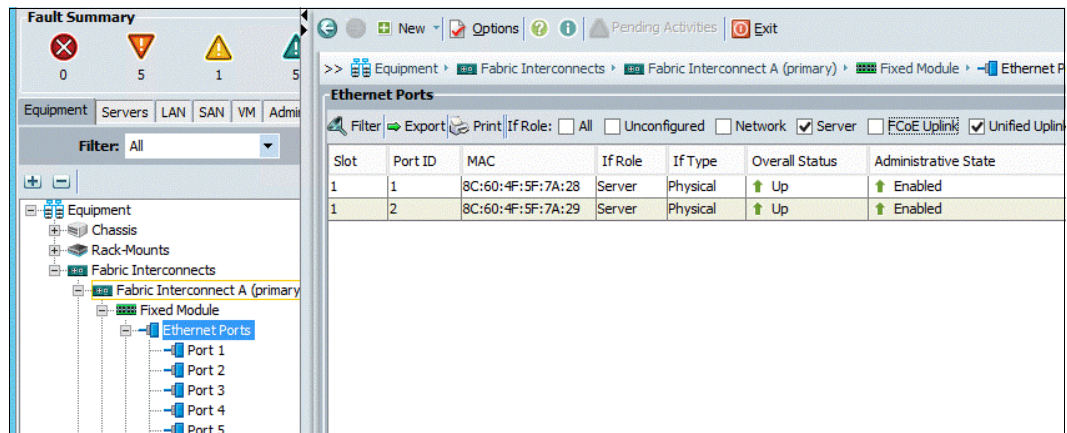


Figure 8-7 Verify the server ports

7. Select ports 25 and 26 that are connected to the Cisco Nexus switches, right-click them, and select **Configure as Uplink Port**.
8. Click **Yes** to confirm the uplink ports and click **OK**.
9. Click **Equipment** → **Fabric Interconnects** → **Fabric Interconnect B (subordinate)** → **Fixed Module**.
10. Expand **Ethernet Ports**.
11. Select the ports that are connected to the chassis, right-click them, and select **Configure as Server Port**.
12. Click **Yes** to confirm server ports and click **OK**.
13. Select ports 25 and 26 that are connected to the Cisco Nexus switches, right-click them, and select **Configure as Uplink Port**.
14. Click **Yes** to confirm the uplink ports and click **OK**.

Changing FI to FC Switching Mode on FI-A and FI-B

Switching FC modes requires the Fabric Interconnects to restart. The restart takes place automatically. When the Fabric Interconnects complete the restart process, a new management session must be established to continue with management and configuration.

Complete the following steps:

1. In the window that is shown in Figure 8-8 on page 65, go to the Equipment tab in the left pane and expand the **Fabric Interconnects** object.

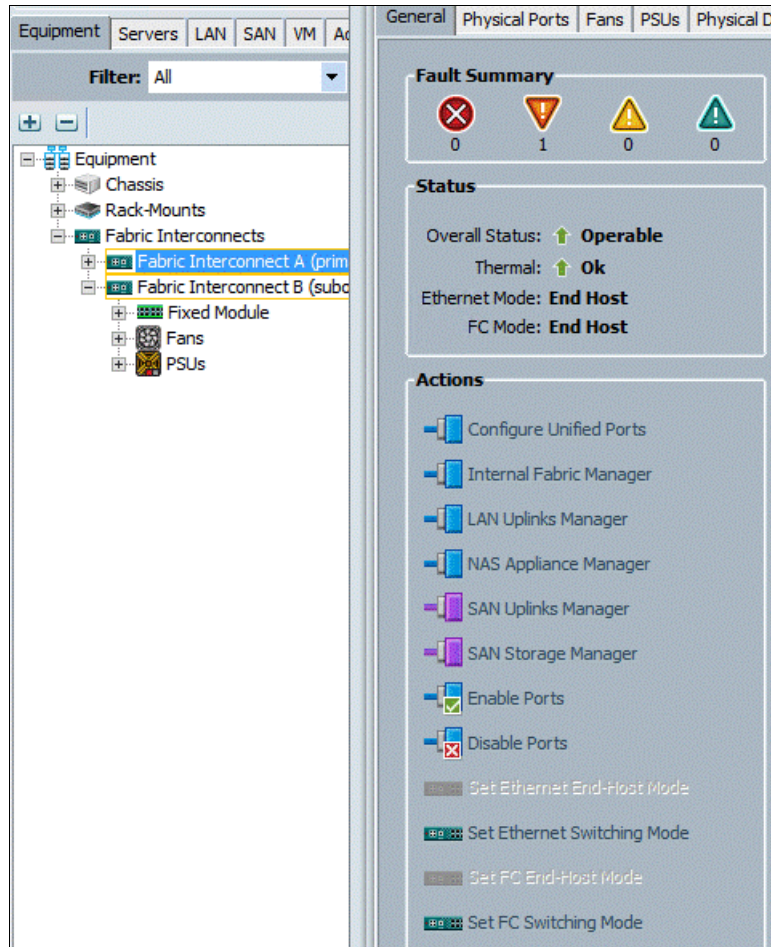


Figure 8-8 Check fabric interconnects

2. Select **Fabric Interconnect A**, in the left pane, click the **General** tab, and click **Set FC Switch Mode** in the left pane.
3. Click **Yes** and then **OK**.
4. Wait for the Fabric Interconnects to restart before proceeding. This process takes approximately 5 minutes for the restart of both nodes.

8.2.7 Enabling Fibre Channel ports

To enable the server and FC uplink ports, complete the following steps:

1. On the Equipment tab, select **Fabric Interconnect B**, which should be the subordinate FI, select **Configure Unified Ports**, and click **Yes**, as shown in Figure 8-9.




Figure 8-9 Configure Unified Ports

2. Slide the bar to the left to select ports 31 - 32 for FC (purple), click **Finish**, and click **Yes** in response to the restart message, as shown in Figure 8-10 on page 67. You must log in to the client again after the restart of the FI completes.

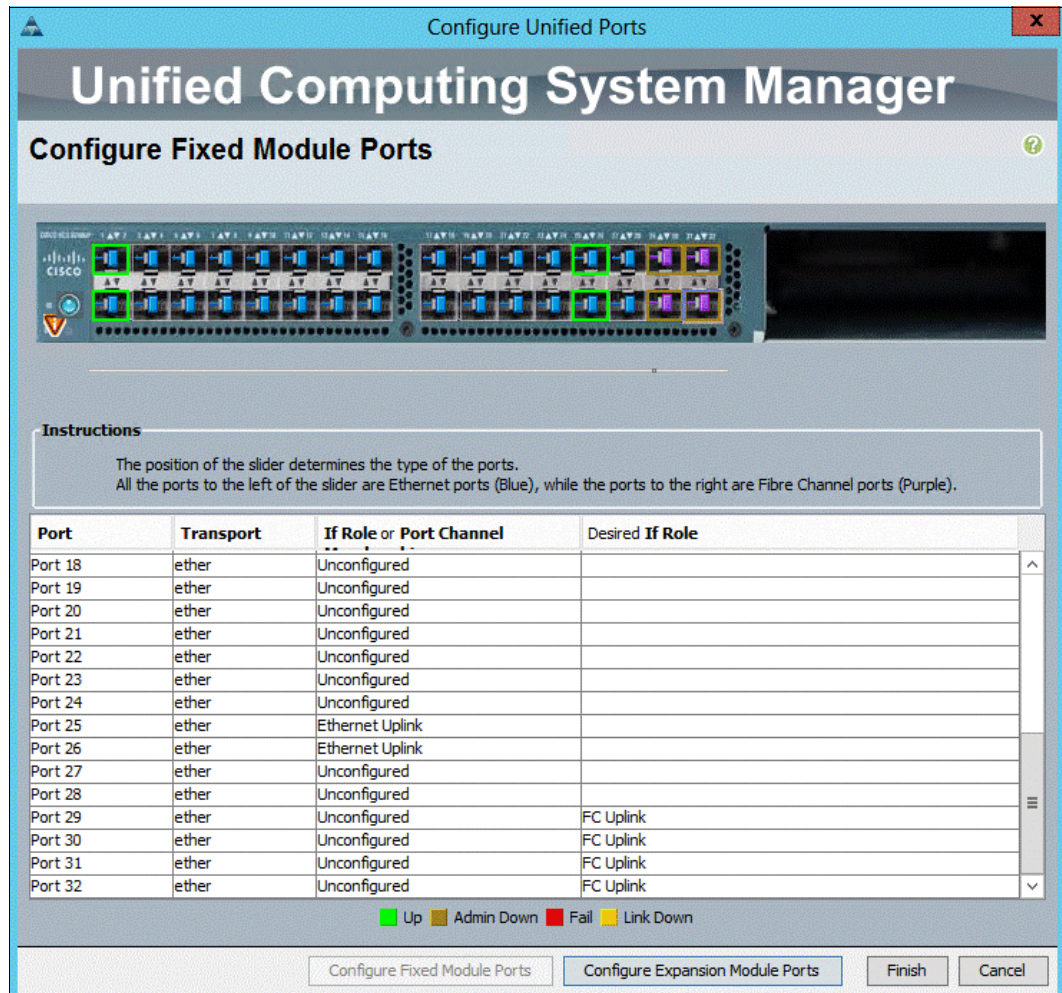


Figure 8-10 Check ports

3. Select **Fabric Interconnect A (primary)**, select **Configure Unified Ports**, and click **Yes**.
4. Slide the bar to the left to select ports 29 - 32 for FC (purple), click **Finish**, and click **Yes** in response to the restart message. You must log in to the client again after the restart of the FI completes.

8.2.8 Creating storage VSANs

To configure the necessary VSANs and FC Port Channels for the Cisco UCS environment, complete the following steps:

1. Select the **SAN** tab at the upper left of the window.
2. Expand the **Storage Cloud** tree.

3. Right-click VSANs, as shown in Figure 8-11.




Figure 8-11 Click VSANs

4. Select **Create Storage VSAN**. The window that is shown in Figure 8-12 opens.

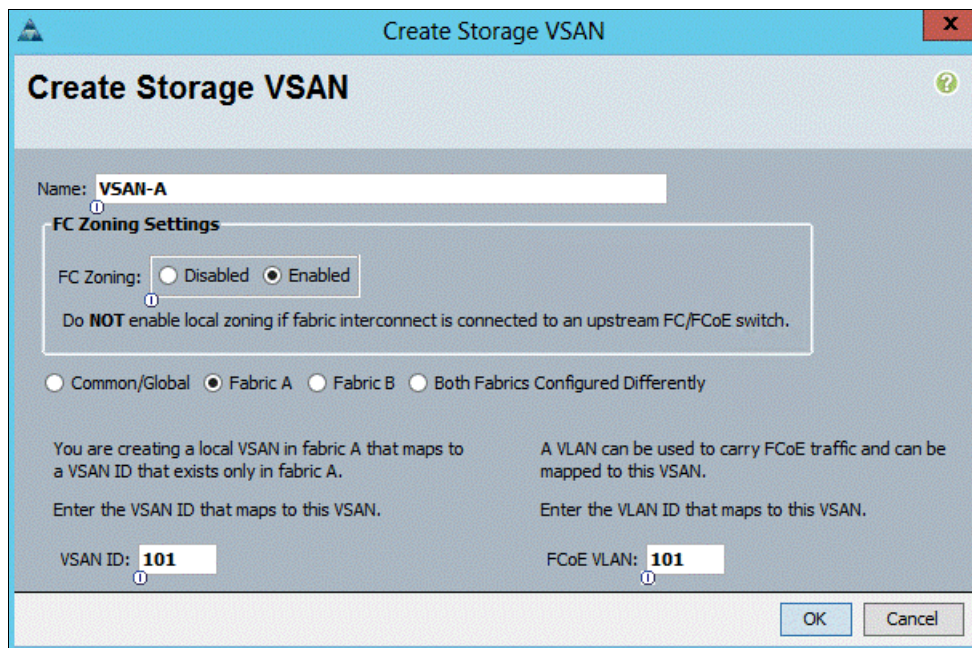


Figure 8-12 Create storage VSAN-A

5. Enter VSAN_A as the VSAN name for Fabric A.
6. Select **Enabled** under the FC Zoning Settings.
7. Select **Fabric A**.
8. Enter the VSAN ID (<<var_vsan_a_id>>) for Fabric A.
9. Enter the FCoE VLAN ID (<<var_vsan_a_id>>) for Fabric A.
10. Click **OK**, and then **OK** again to create the VSAN.
11. In the window that is shown in Figure 8-11, right-click **VSANs** and select **Create Storage VSAN**. The window that is shown in Figure 8-13 on page 69 opens.

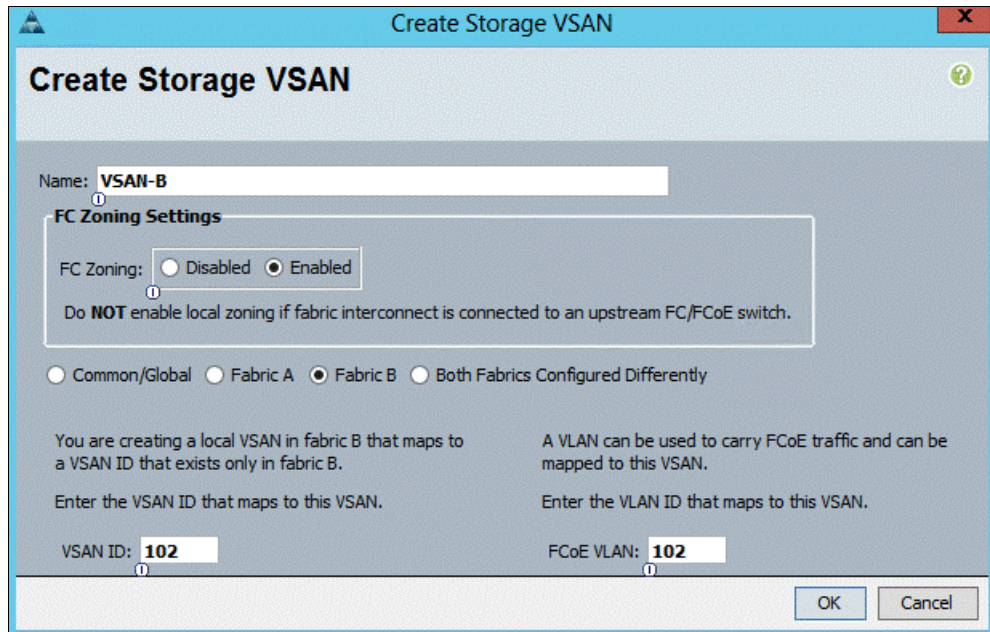


Figure 8-13 Create storage VSAN-B

12. Enter VSAN_B as the VSAN name for Fabric B.
13. Select **Enabled** under the FC Zoning Settings.
14. Select **Fabric B**.
15. Enter the VSAN ID (<<var_vsan_b_id>>) for Fabric B.
16. Enter the FCoE VLAN ID (<<var_vsan_b_id>>) for Fabric B.
17. Click **OK**, and then **OK** to create the VSAN.

8.2.9 Configuring the FC storage ports

To configure the FC storage ports, complete the following steps:

1. Click the **Equipment** tab at the upper left of the window.
2. Click **Equipment** → **Fabric Interconnects** → **Fabric Interconnect A (primary)** → **Fixed Module**. The window that is shown in Figure 8-14 opens.

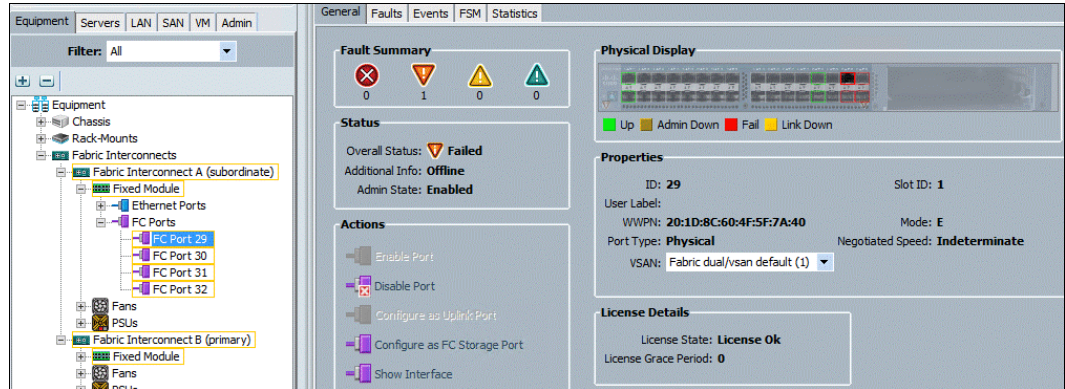


Figure 8-14 Configure the FC port

3. Expand the **FC Ports** object.
4. Select **FC Port 29**, which is connected to the IBM storage array.
5. Under the General tab, click **Configure as FC Storage Port**.
6. Click **Yes**, and then click **OK**.
7. Repeat steps 2 - 6 for FC ports 30 – 32.
8. Click **Equipment** → **Fabric Interconnects** → **Fabric Interconnect B (primary)** → **Fixed Module**. The window that is shown in Figure 8-15 opens.

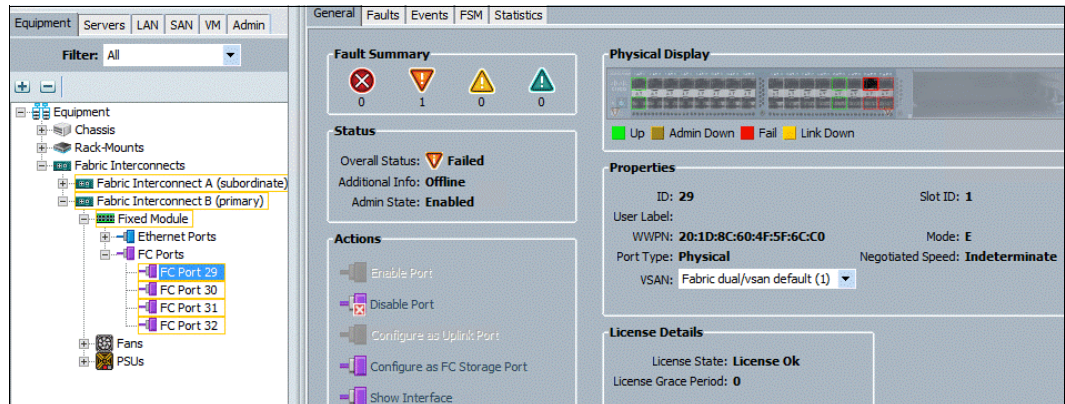


Figure 8-15 General tab

9. Expand the **FC Ports** object.
10. Select **FC Port 29**, which is connected to the IBM storage array.
11. Under the General tab, select **Configure as FC Storage Port**.
12. Click **Yes**, and then click **OK**.
13. Repeat the steps 8 - 12 for FC ports 30 – 32.

8.2.10 Configuring the VSAN for the FC storage ports

To configure VSAN-A and VSAN-B, complete the following steps:

1. Click the **Equipment** tab at the upper left of the window.
2. Click **Equipment** → **Fabric Interconnects** → **Fabric Interconnect A (primary)** → **Fixed Module**. The window that is shown in Figure 8-16 opens.

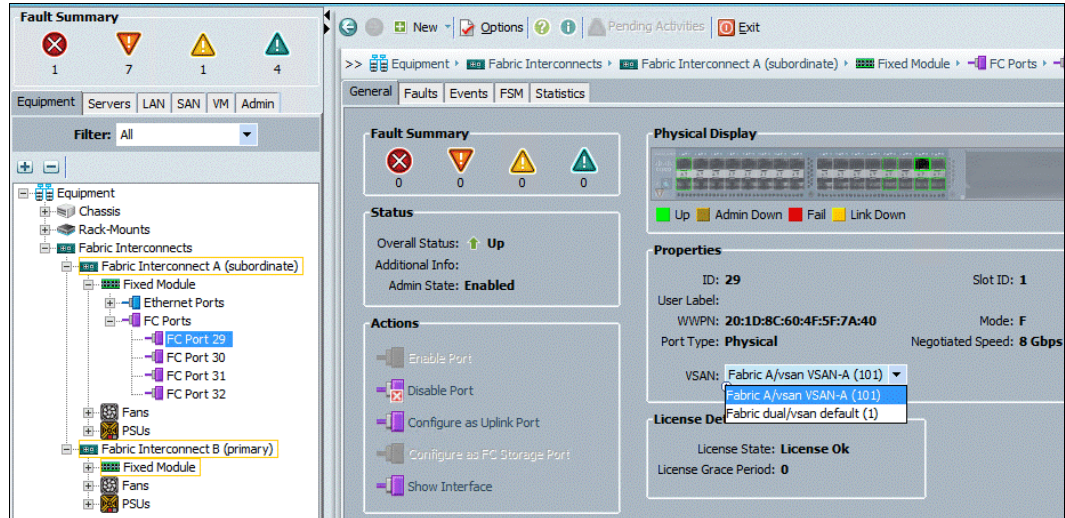


Figure 8-16 Select VSAN

3. Expand the **FC Ports** object.
4. Select **FC Port 29**, which is connected to the IBM storage array.
5. In the right pane, click the **VSAN** drop-down menu and select **Fabric A / vsan VSAN-A (101)**.
6. Click **Save Changes** and then click **OK**.
7. Repeat the steps 2 to 6 for FC ports 30 – 32.
8. Click **Equipment** → **Fabric Interconnects** → **Fabric Interconnect B (subordinate)** → **Fixed Module**. The window that is shown in Figure 8-17 opens.

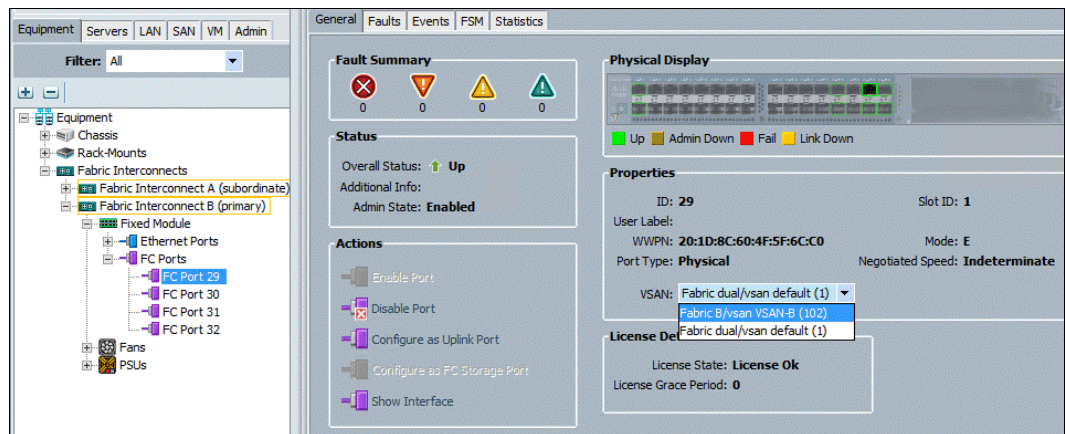


Figure 8-17 Select VSAN

9. Expand the **FC Ports** object.

10. Select **FC Port 29**, which is connected to the storage array.
11. In the right pane, click the **VSAN** drop-down menu and select **Fabric B / vsan VSAN-B (102)**.
12. Click **Save Changes**, and then click **OK**.
13. Repeat steps 8 on page 71 - 12 for FC ports 30 – 32.
14. Verify the roles and statuses of the ports, as shown in Figure 8-18.

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Administrative State
1	29	20:1D:8C:60:4F:5F:6C:C0	Storage	Physical	Up	Enabled
1	30	20:1E:8C:60:4F:5F:6C:C0	Storage	Physical	Up	Enabled
1	31	20:1F:8C:60:4F:5F:6C:C0	Storage	Physical	Up	Enabled
1	32	20:20:8C:60:4F:5F:6C:C0	Storage	Physical	Up	Enabled

Figure 8-18 Verify roles and statuses

15. You should see the storage arrays WWPN flogi in the Cisco UCS Fabric Interconnect. You can view flogi by connecting to the Cisco UCS Manager cluster IP through SSH:

```
Versastack-FI-B# connect nxos b
Versastack-FI-B(nxos)# sh flogi database
INTERFACE VSAN FCID PORT NAME NODE NAME
fc1/29 102 0x0c0000 50:05:07:68:0b:21:4f:f5 50:05:07:68:0b:00:4f:f5
fc1/30 102 0x0c0020 50:05:07:68:0b:22:4f:f4 50:05:07:68:0b:00:4f:f4
fc1/31 102 0x0c0040 50:05:07:68:0b:22:4f:f5 50:05:07:68:0b:00:4f:f5
fc1/32 102 0x0c0060 50:05:07:68:0b:23:4f:f4 50:05:07:68:0b:00:4f:f4
Total number of f logi = 4.
Versastack-FI-B (nxos) f exit
Versastack-FI-BI connect nxos a
Versastack-FI-A(nxos)# sh flogi database
INTERFACE VSAN FCID PORT NAME NODE NAME
fc1/29 101 0xa90000 50:05:07:68:0b:23:4f:f550:05:07:68:0b:00:4f:f5
fc1/30 101 0xa90020 50:05:07:68:0b:24:4f:f4 50:05:07:68:0b:00:4f:f4
fc1/31101 0xa90040 50:05:07:68:0b:24:4f:f550:05:07:68:0b:00:4f:f5
fc1/32 101 0xa90160 50:05:07:68:0b:23:4f:f4 50:05:07:68:0b:00:4f:f4
Total number of f logi = 4.
```

This is also shown in Figure 8-19 on page 73.

```

Versastack-FI-B# connect nxos b
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
Versastack-FI-B(nxos)# sh flogi database
-----
INTERFACE          VSAN    FCID          PORT NAME          NODE NAME
-----
fc1/29              102     0x0c0000      50:05:07:68:0b:21:4f:f5  50:05:07:68:0b:00:4f:f5
fc1/30              102     0x0c0020      50:05:07:68:0b:22:4f:f4  50:05:07:68:0b:00:4f:f4
fc1/31              102     0x0c0040      50:05:07:68:0b:22:4f:f5  50:05:07:68:0b:00:4f:f5
fc1/32              102     0x0c0060      50:05:07:68:0b:23:4f:f4  50:05:07:68:0b:00:4f:f4

Total number of flogi = 4.

Versastack-FI-B(nxos)# exit
Versastack-FI-B# connect nxos a
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
Versastack-FI-A(nxos)# sh flogi database
-----
INTERFACE          VSAN    FCID          PORT NAME          NODE NAME
-----
fc1/29              101     0xa90000      50:05:07:68:0b:23:4f:f5  50:05:07:68:0b:00:4f:f5
fc1/30              101     0xa90020      50:05:07:68:0b:24:4f:f4  50:05:07:68:0b:00:4f:f4
fc1/31              101     0xa90040      50:05:07:68:0b:24:4f:f5  50:05:07:68:0b:00:4f:f5
fc1/32              101     0xa90160      50:05:07:68:0b:23:4f:f4  50:05:07:68:0b:00:4f:f4

Total number of flogi = 4.

```

Figure 8-19 WWPNs flogi

8.2.11 Creating WWNN pools

To configure the necessary World Wide Node Name (WWNN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Click **Pools** → **root**.

3. Right-click **WWNN Pools**, as shown in Figure 8-20.

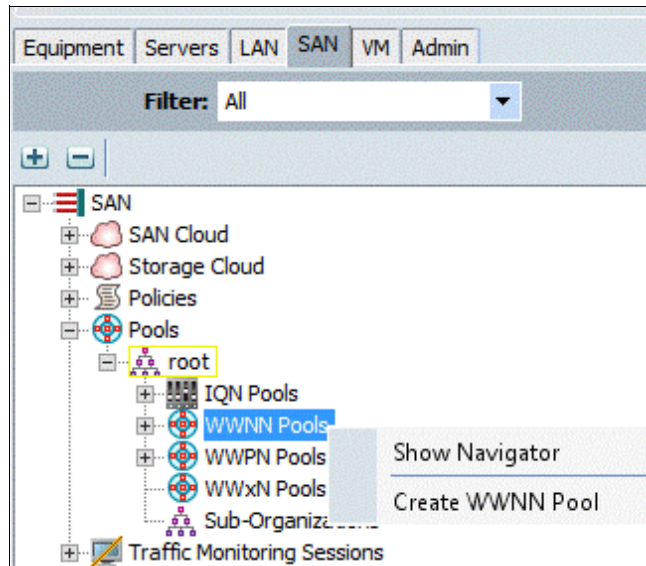


Figure 8-20 Create WWNN Pool

4. Select **Create WWNN Pool**.
5. Enter WWNN_Pool as the name of the WWNN pool.
6. (Optional) Add a description for the WWNN pool.
7. Click **Next**.
8. Click **Add** to add a block of WWNNs.
9. Keep the default block of WWNNs, or specify a base WWNN.
10. Specify a size for the WWNN block that is sufficient to support the available blade or server resources, as shown in Figure 8-21.

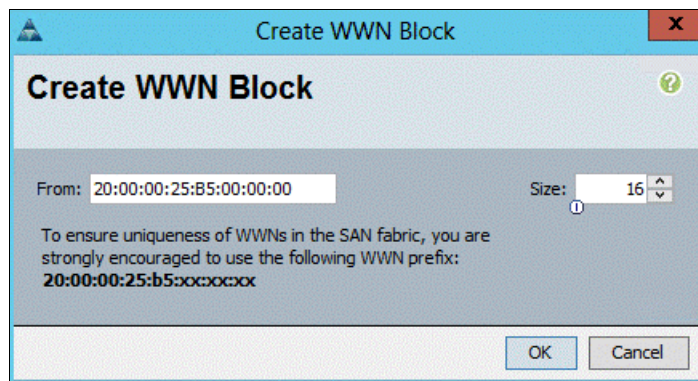


Figure 8-21 Create WWNN block

11. Click **OK**.
12. Click **Finish**, as shown in Figure 8-22 on page 75.

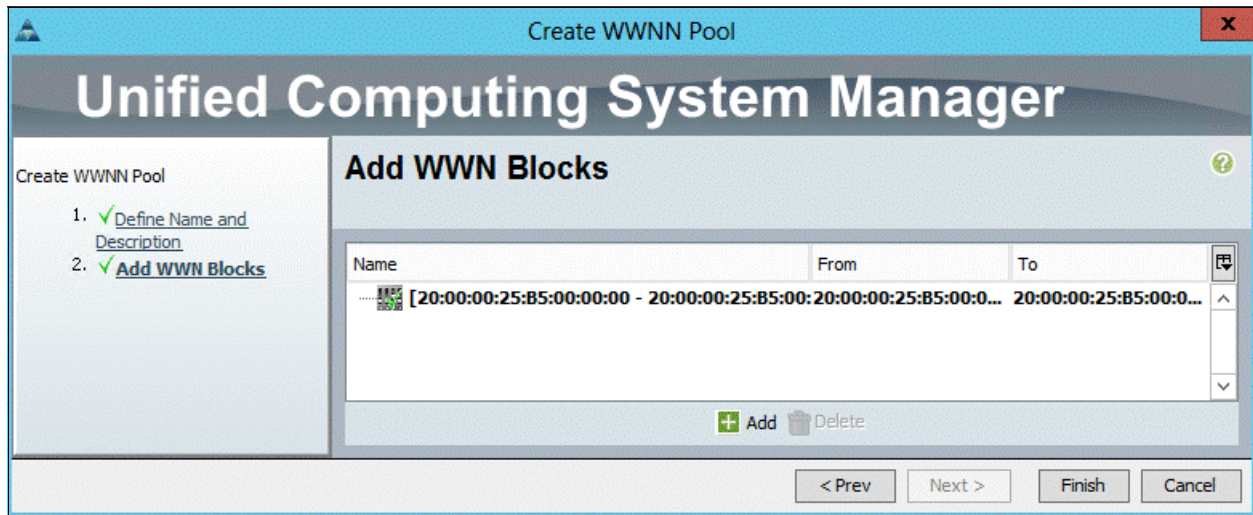


Figure 8-22 Create WWNN pool

13. Click **OK**.

8.2.12 Creating WWPN pools

To configure the necessary worldwide port name (WWPN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Click **Pools** → **root**.

Note: In this procedure, two WWPN pools are created: one for fabric A and one for fabric B.

3. Right-click **WWPN Pools**, as shown in Figure 8-23.

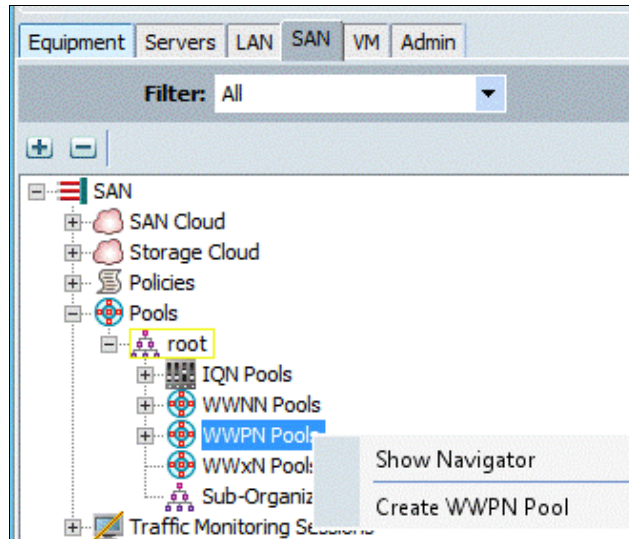


Figure 8-23 Create WWPN Pool

4. Select **Create WWPN Pool**.
5. Enter WWPN_Pool_A as the name of the WWPN pool for Fabric A.
6. (Optional) Enter a description for this WWPN pool.
7. Click **Next**.
8. Click **Add** to add a block of WWPNs.
9. Specify the starting WWPN in the block for Fabric A, as shown in Figure 8-24.

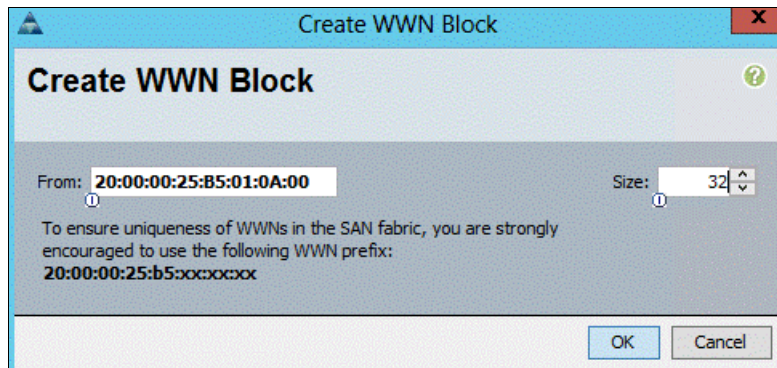


Figure 8-24 Create WWN block

Note: For the VersaStack solution, the recommendation is to place 0A in the next-to-last octet of the starting WWPN to identify all the WWPNs in this pool as Fabric A addresses.

10. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.
11. Click **OK**.
12. Click **Finish** to create the WWPN pool, as shown in Figure 8-25 on page 77.

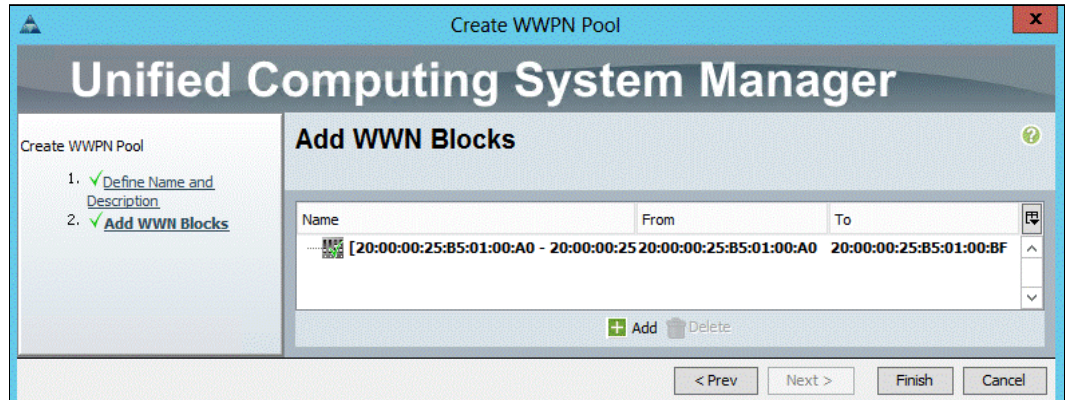


Figure 8-25 Create the WWPN pool

13. Click **OK**.
14. Right-click **WWPN Pools**.
15. Click **Create WWPN Pool**.
16. Enter WWPN_Pool_B as the name for the WWPN pool for Fabric B.
17. (Optional) Enter a description for this WWPN pool.
18. Click **Next**.
19. Click **Add** to add a block of WWPNs.
20. Enter the starting WWPN address in the block for Fabric B.

Note: For the VersaStack solution, the recommendation is to place 0B in the next to last octet of the starting WWPN to identify all the WWPNs in this pool as Fabric B addresses.

21. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.
22. Click **OK**.
23. Click **Finish**.
24. Click **OK**.

Figure 8-26 shows successful pool creation.

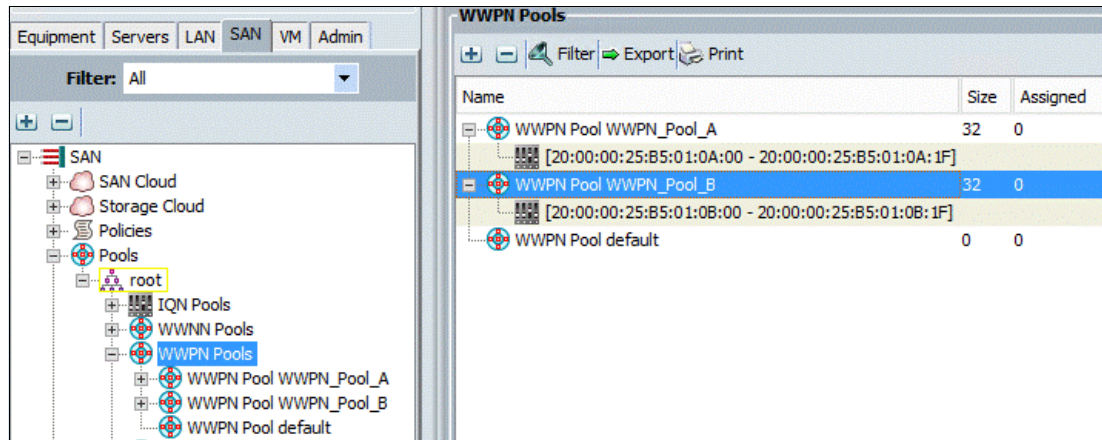


Figure 8-26 Check pool creation

8.2.13 Creating vHBA templates for Fabric A and Fabric B

To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **vHBA Templates**, as shown in Figure 8-27.

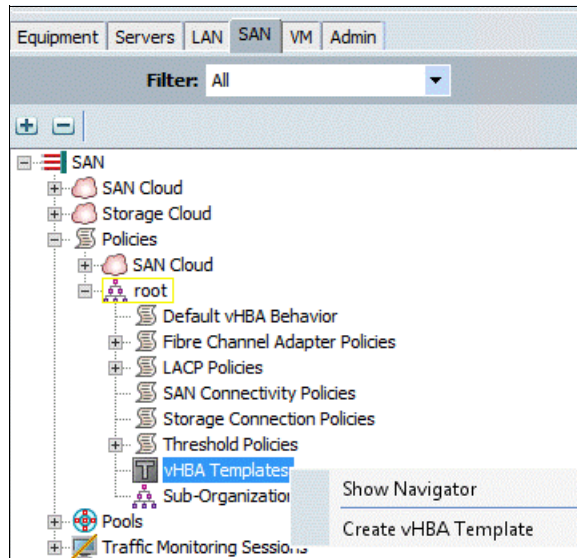


Figure 8-27 Create vHBA Template

4. Select **Create vHBA Template**. The window that is shown in Figure 8-28 on page 79 opens.




Figure 8-28 Create vHBA Template

5. Enter vHBA_Template_A as the vHBA template name.
6. Select **A** for Fabric ID.
7. In the Select VSAN list, select **VSAN_A**.
8. In the WWPN Pool list, select **WWPN_Pool_A**.
9. Click **OK** to create the vHBA template, and click **OK** again.
10. In the navigation pane, click the **SAN** tab.
11. Click **Policies** → **root**.
12. Right-click **vHBA Templates**.

13. Select **Create vHBA Template**. The window that is shown in Figure 8-29 opens.

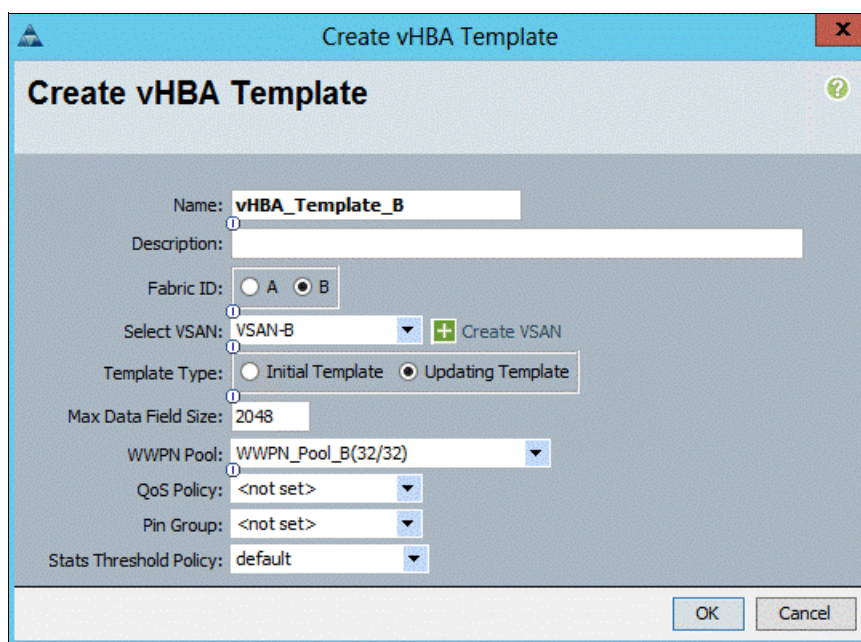


Figure 8-29 Create vHBA Template

14. Enter vHBA_Template_B as the vHBA template name.
15. Select **B** for Fabric ID.
16. In the Select VSAN list, select **VSAN_B**.
17. In the WWPN Pool, select **WWPN_Pool_B**.
18. Click **OK** to create the vHBA template, and click **OK** again.

8.2.14 Creating the storage connection policy for Fabric-A

To create a storage policy for Fabric-A that helps create the FC fabric zoning, complete the following steps:

1. Select the **SAN** tab at the upper left of the window.
2. Click **Policies** → **root**.
3. Right-click **Storage Connection Policies**. The window that is shown in Figure 8-30 on page 81 opens.

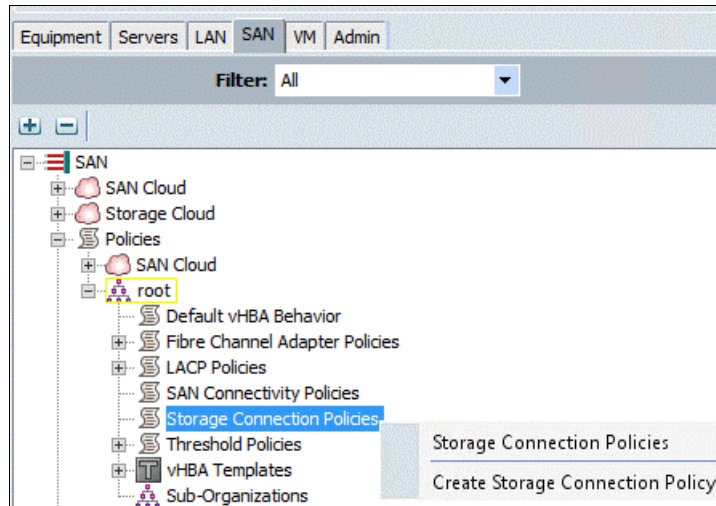


Figure 8-30 Create Storage Connection Policy

4. Select **Create Storage Connection Policy**. The window that is shown in Figure 8-31 opens.

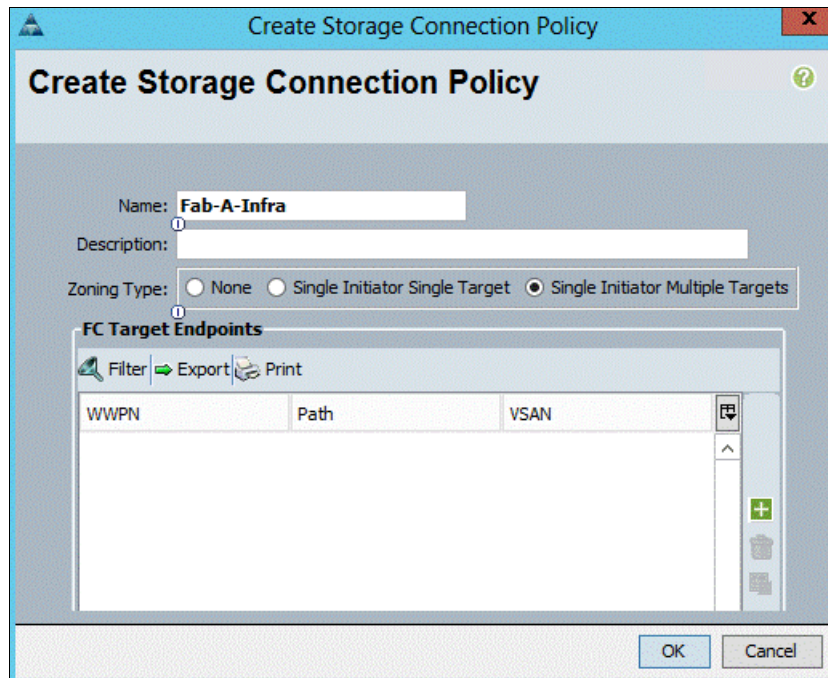


Figure 8-31 Create Storage Connection Policy

5. In the Name field, enter the Storage Connection Policy named Fab-A-Infra.
6. Select **Single Initiator Multiple Targets** for Zoning Type.

- Click the plus icon to add the FC Target Endpoint. The window that is shown in Figure 8-32 opens.

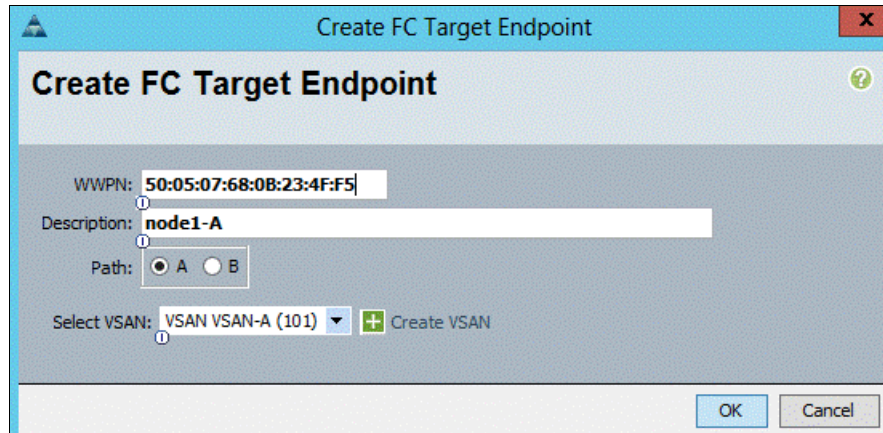


Figure 8-32 Create FC Target Endpoint

- In the WWPN field, enter <<var_wwpn_Node1-switch-A>> for Node 1 Fabric A.
- Select **A** for Path.
- Select **VSAN_A** for the Select VSAN field.
- Click **OK** to create the FC Target Endpoint.
- Repeat steps 7 - 11 to create the remaining FC target endpoints on fabric path A, as shown in Figure 8-33.

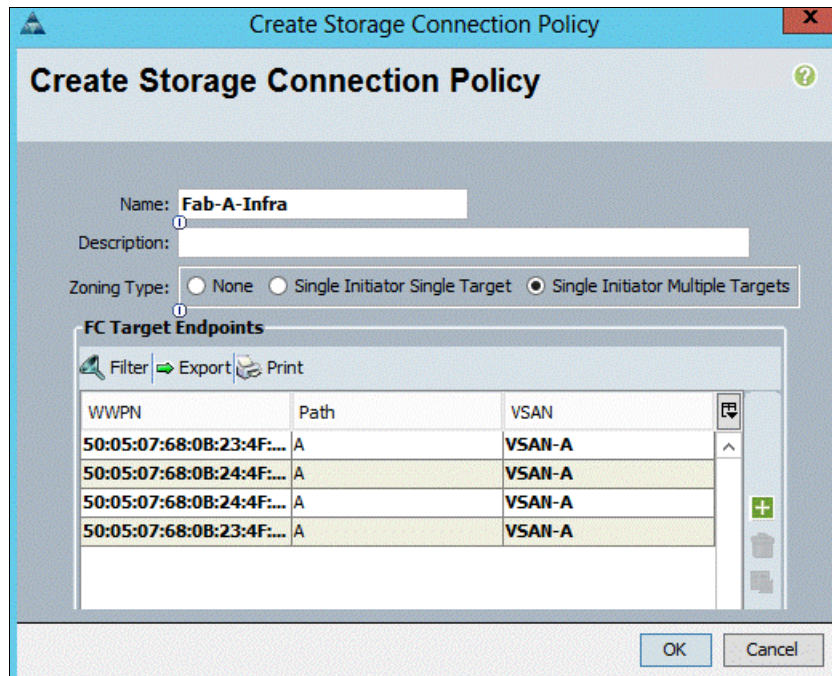


Figure 8-33 Ensure that all the policies are created

8.2.15 Creating the Storage Connection Policy for Fabric-B

To create a storage policy for Fabric-B that helps create the FC fabric zoning, complete the following steps:

1. Select the **SAN** tab at the upper left of the window.
2. Click **Policies** → **root**.
3. Right-click **Storage Connection Policies**.
4. Select **Create Storage Connection Policy**. The window that is shown in Figure 8-34 opens.

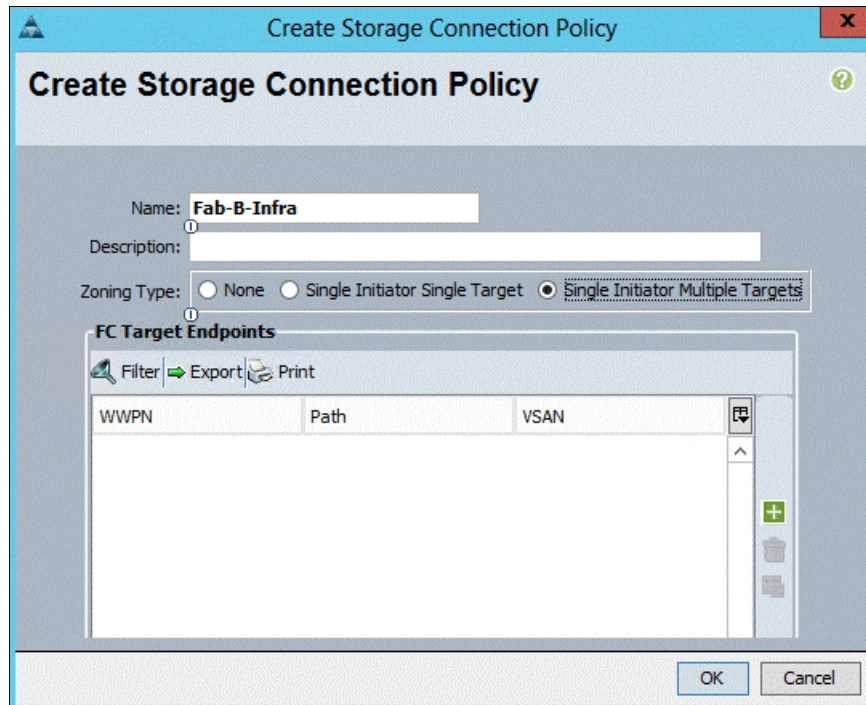


Figure 8-34 Create Storage Connection Policy

5. In the Name field, enter the Storage Connection Policy named Fab-B-Infra.
6. Select **Single Initiator Multiple Targets** for the Zoning Type.

- Click the plus icon to add the FC Target Endpoint. The window that is shown in Figure 8-35 opens.

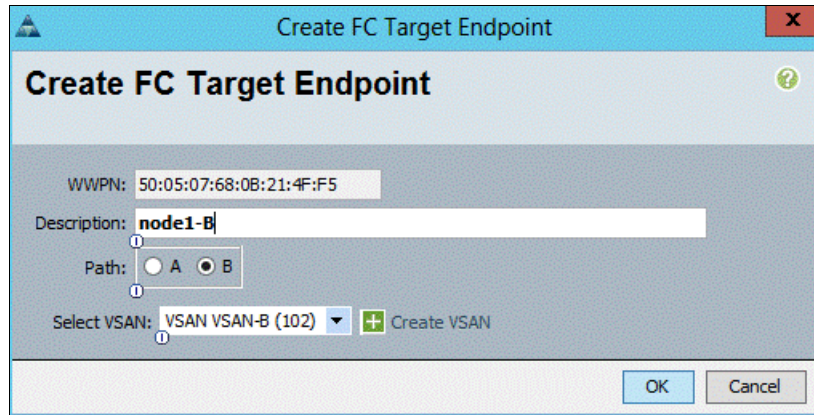


Figure 8-35 Create FC Target Endpoint

- In the WWPN field, enter <<var_wwpn_Node1-switch-A>> for Node 1 Fabric B.
- Select **B** for Path.
- For the Select VSAN field, select **VSAN_B**.
- Click **OK** to create the FC Target Endpoint.
- Repeat steps 7 - 11 to create the remaining FC target endpoints on fabric path B, as shown in Figure 8-36.

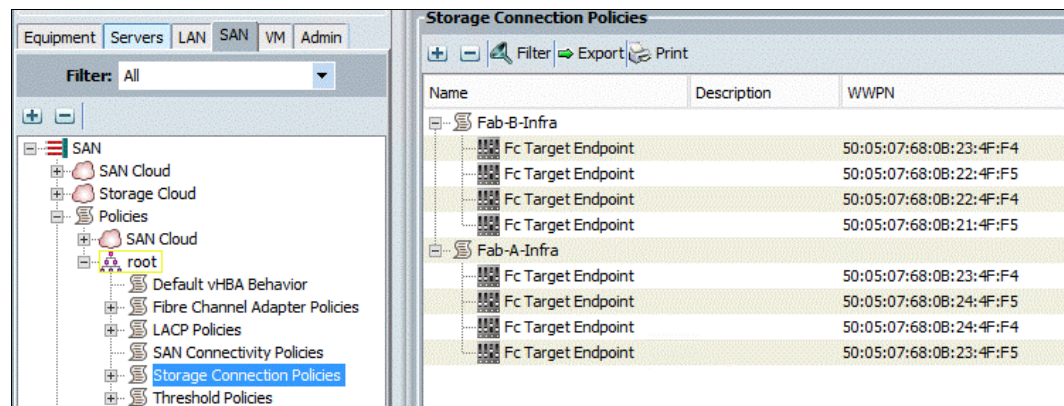


Figure 8-36 Create the remaining FC target endpoints on fabric path B

Creating a SAN connectivity policy

To create a SAN connectivity policy that is used for automated Fibre Channel zone creation on the Fabric interconnects, complete the following steps:

- Select the **SAN** tab at the upper left of the window.
- Click **Policies** → **root**.
- Right-click **SAN Connectivity Policies** and select **Create SAN Connectivity Policy**, as shown in Figure 8-37 on page 85. The window that is shown in Figure 8-38 on page 85 opens.

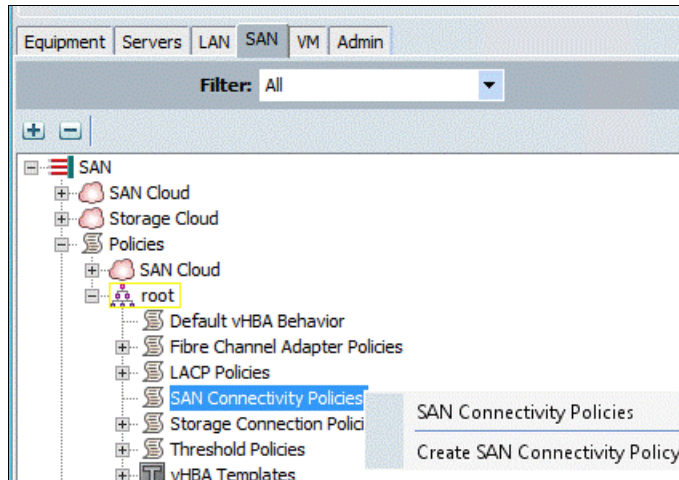


Figure 8-37 Click Create SAN Connectivity Policy

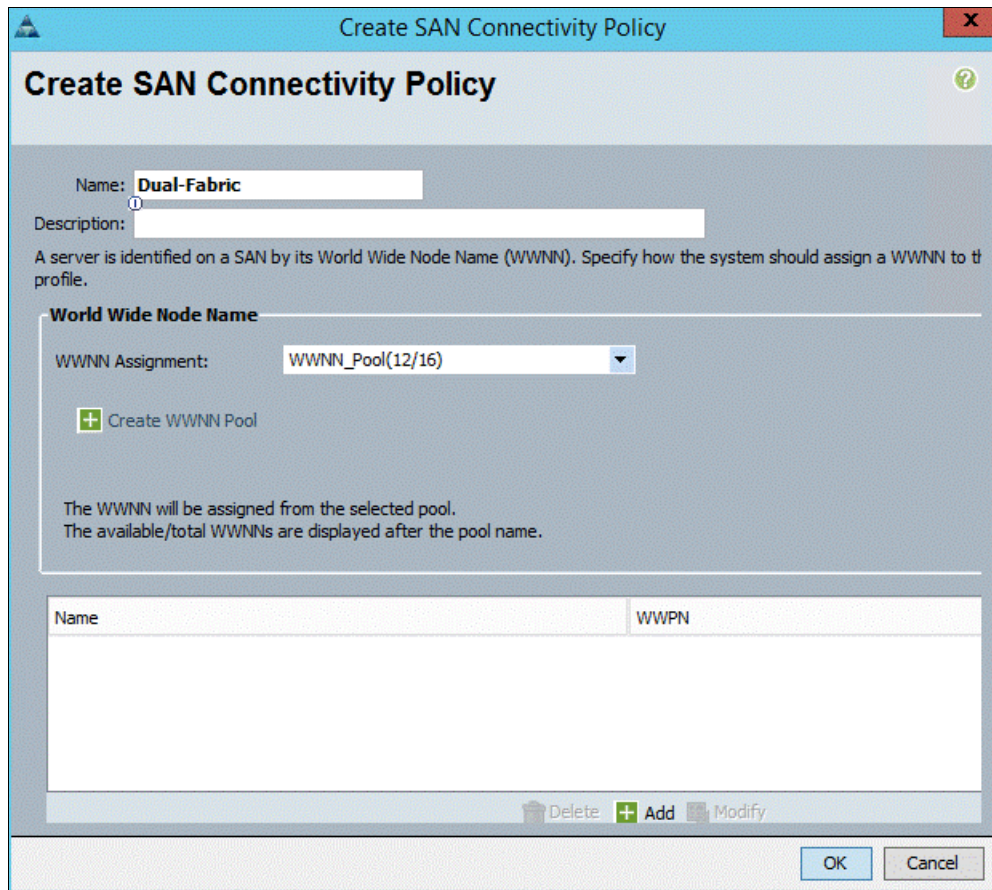


Figure 8-38 Select WWNN_Pool for WWNN assignment

4. In the Name field, enter Dual-Fabric.
5. For WWNN Assignment, select **WWNN_Pool**.

6. Click **Add** at the bottom of the window. The window that is shown in Figure 8-39 opens.

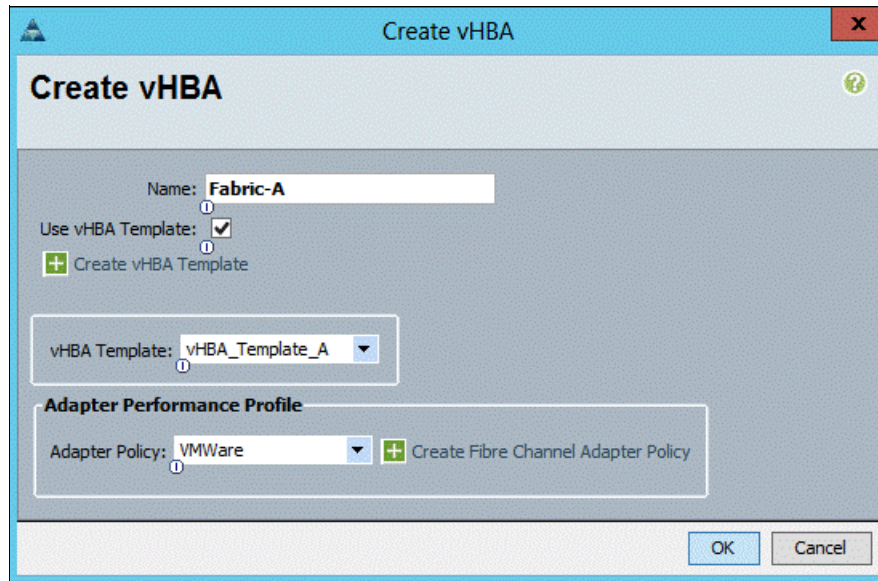


Figure 8-39 Create vHBA

7. For the Name field, enter Fabric-A.
8. Select the **Use vHBA Template** check box.
9. In the vHBA Template menu, select **vHBA_Template_A**.
10. In the Adapter Policy menu, select **VMWare**.
11. Click **OK**. The window that is shown in Figure 8-38 on page 85 opens again. Click **Add**. The window that is shown in Figure 8-40 opens.

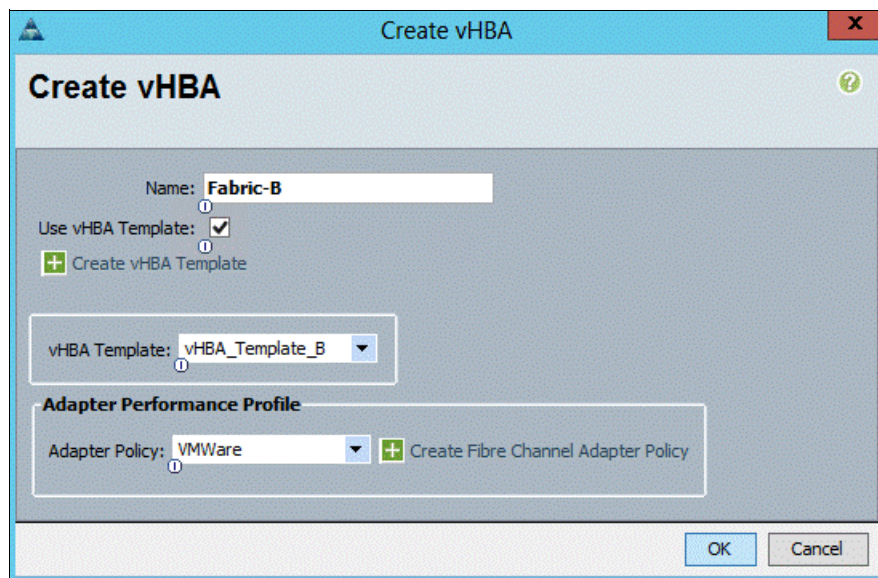


Figure 8-40 Select Adapter Policy VMWare

12. For the Name field, enter Fabric B.
13. Select the **Use vHBA Template** check box.

14. In the vHBA Template menu, select **vHBA_Template_B**.
15. In the Adapter Policy menu, select **VMware**.
16. Click **OK** to complete the policy creation, and click **OK** again. The window that is shown in Figure 8-41 opens.

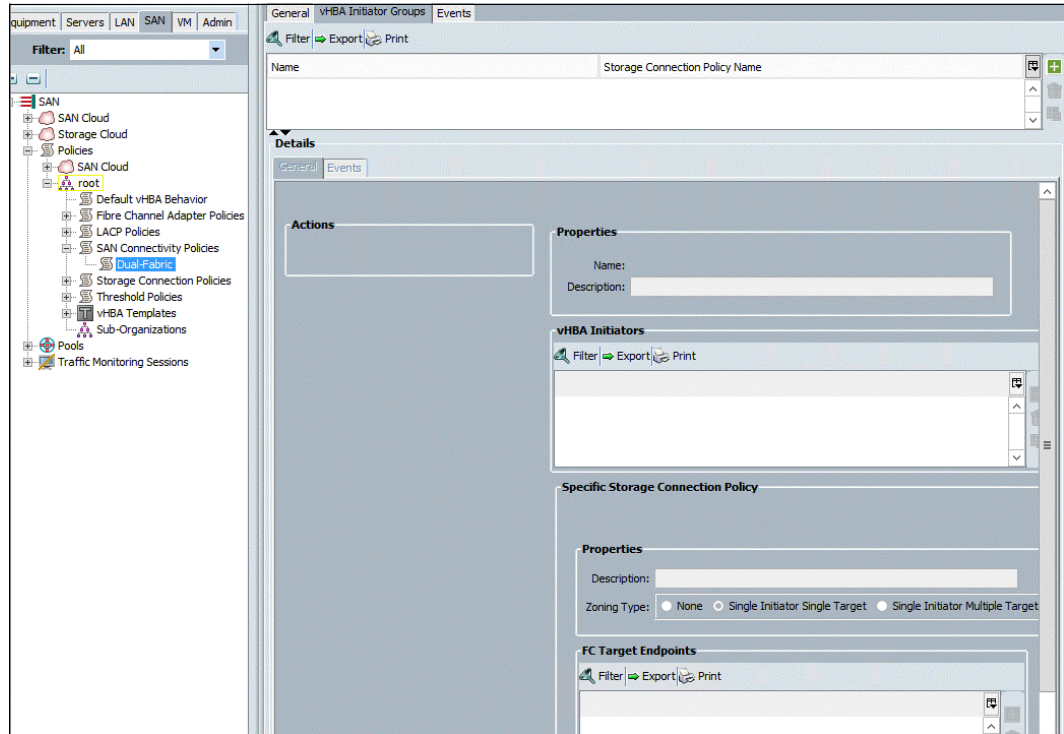


Figure 8-41 Dual-Fabric policy

17. Expand the **SAN Connectivity Policies** and click the **Dual-Fabric** policy.
18. In the right pane, click the **vHBA Initiator Groups** tab.

19. Click the green add button on the right side. The window that is shown in Figure 8-42 opens.

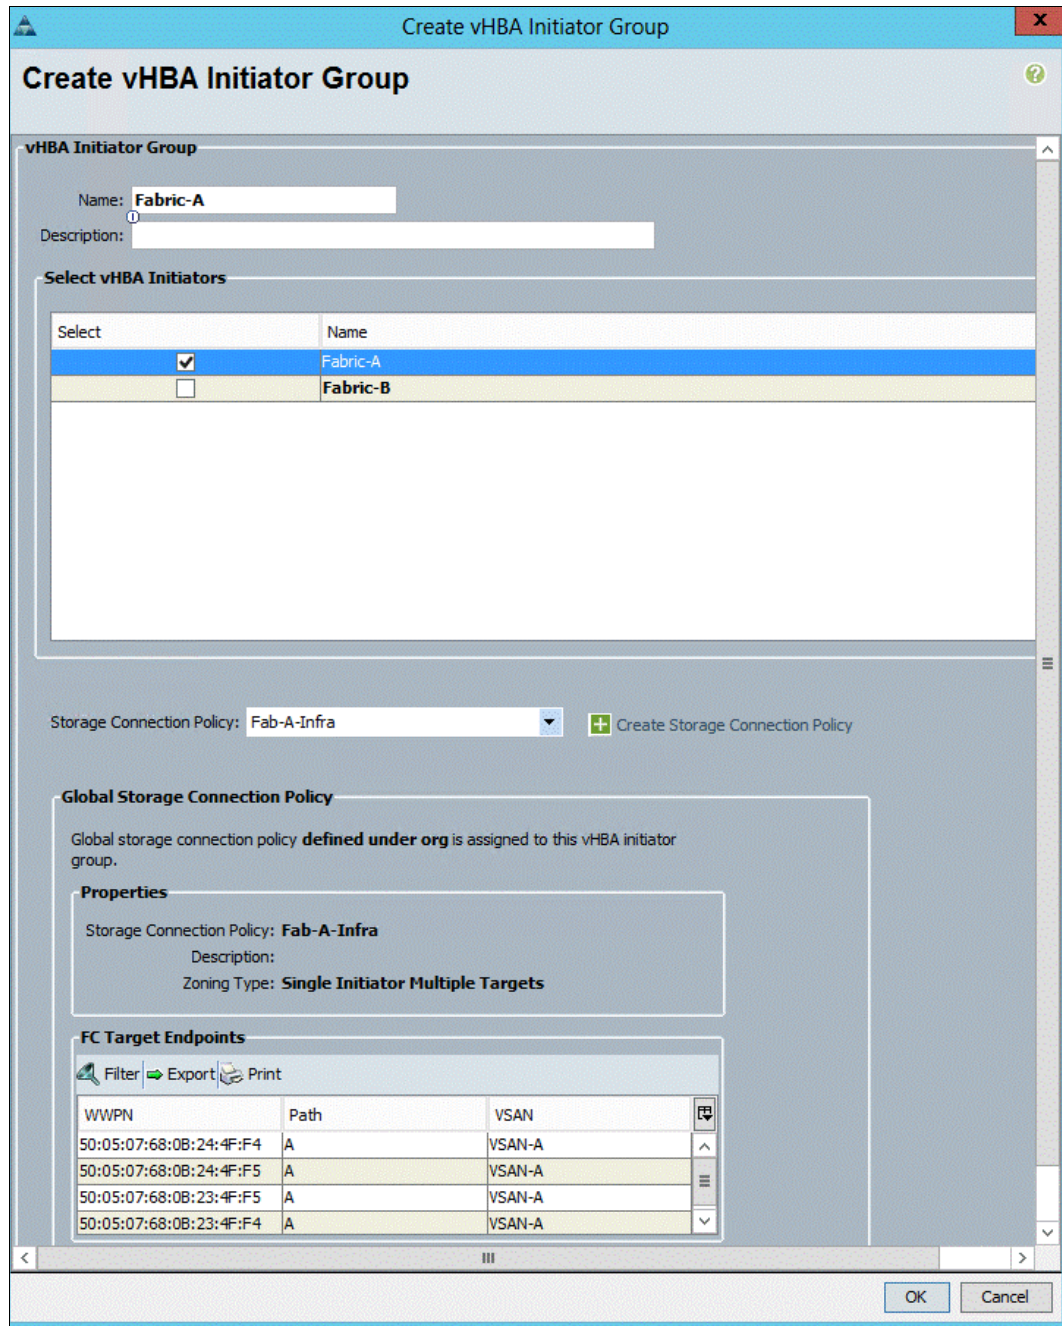


Figure 8-42 Create vHBA Initiator Group

20. In the Name field, enter Fabric-A.

21. Select the **Fabric-A** check box.

22. In the Storage Connection Policy menu, select **Fab-A-Infra**.

23. Click **OK**, and then click **OK** again. The window that is shown in Figure 8-41 on page 87 opens again. Click the green add button on the right side. The window that is shown in Figure 8-43 on page 89 opens.

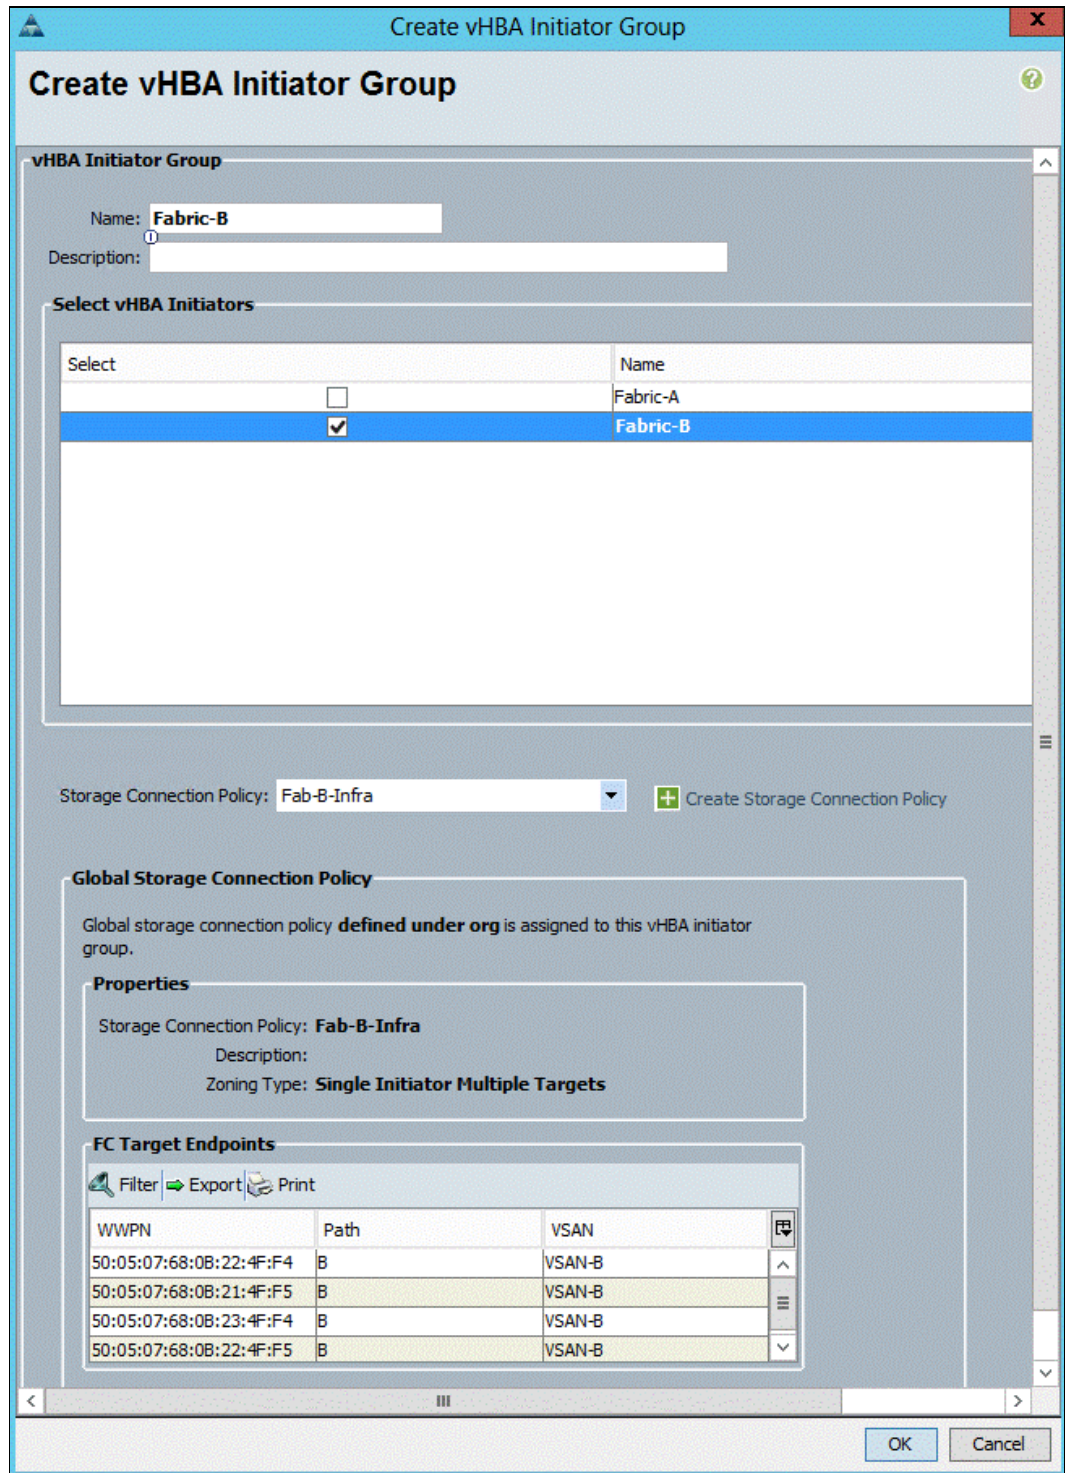


Figure 8-43 Create vHBA Initiator Group

24. In the Name field, enter Fabric-B.
25. In the Select vHBA Initiators pane, select the **Fabric-B** check box.
26. In the Storage Connection Policy menu, Select **Fab-B-Infra**.
27. Click **OK**, and then click **OK** again.

8.2.16 Acknowledging Cisco UCS chassis and FEX modules

To acknowledge all the Cisco UCS chassis and external 2232 FEX modules, complete the following steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane, as shown in Figure 8-44.

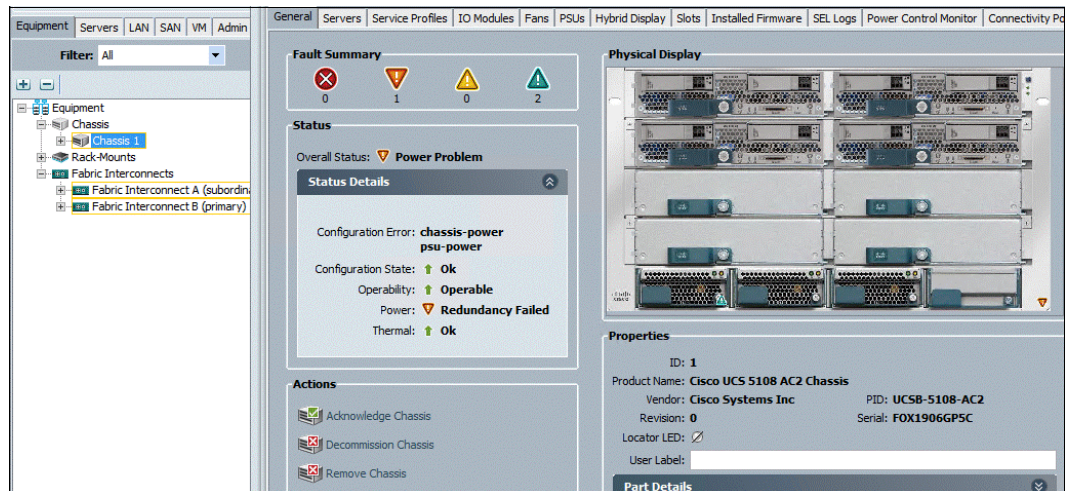


Figure 8-44 Equipment tab in the navigation pane

2. Expand **Chassis** and select each chassis that is listed.
3. Right-click each chassis and select **Acknowledge Chassis**, click **Yes**, and then click OK.
4. If C-Series servers are part of the configuration, expand **Rack Mounts** and **FEX**.
5. Right-click each FEX that is listed and select **Acknowledge FEX**.
6. Click **Yes**, and then click **OK**.

8.2.17 Creating uplink port channels to Cisco Nexus switches

To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.

Note: In this procedure, two port channels are created: one from fabric A to both Cisco Nexus switches and one from fabric B to both Cisco Nexus switches.

2. Click **LAN** → **LAN Cloud** and expand **Fabric A**.
3. Right-click **Port Channels** and select **Create Port Channel**, as shown in Figure 8-45 on page 91. The window that is shown in Figure 8-46 on page 91 opens.

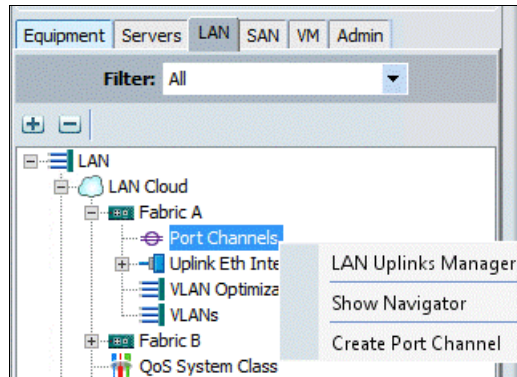


Figure 8-45 Create Port Channel

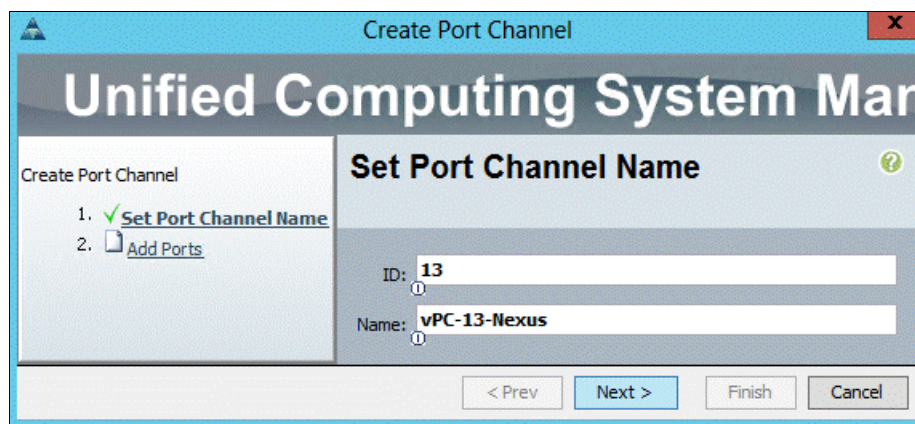


Figure 8-46 Set Port Channel Name

4. Enter 13 as the unique ID of the port channel.
5. Enter vPC-13-Nexus as the name of the port channel.
6. Click **Next**. The window that is shown in Figure 8-47 opens.

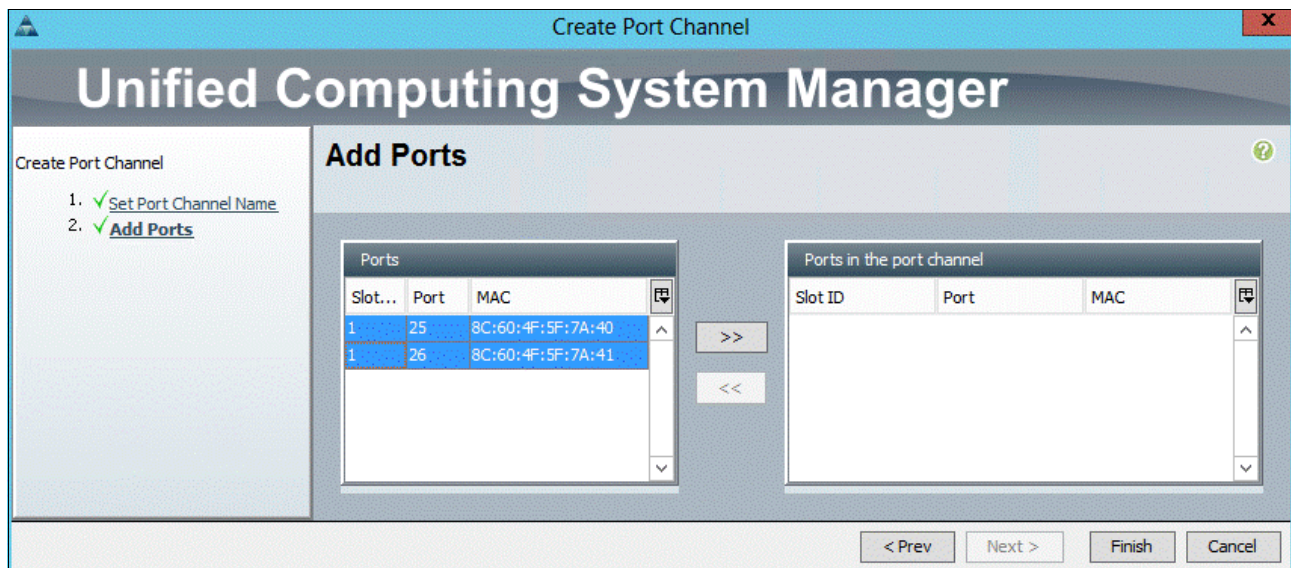


Figure 8-47 Add Ports

7. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 25
 - Slot ID 1 and port 26
8. Click >> to add the ports to the port channel.
9. Click **Finish** to create the port channel, as shown in Figure 8-48.

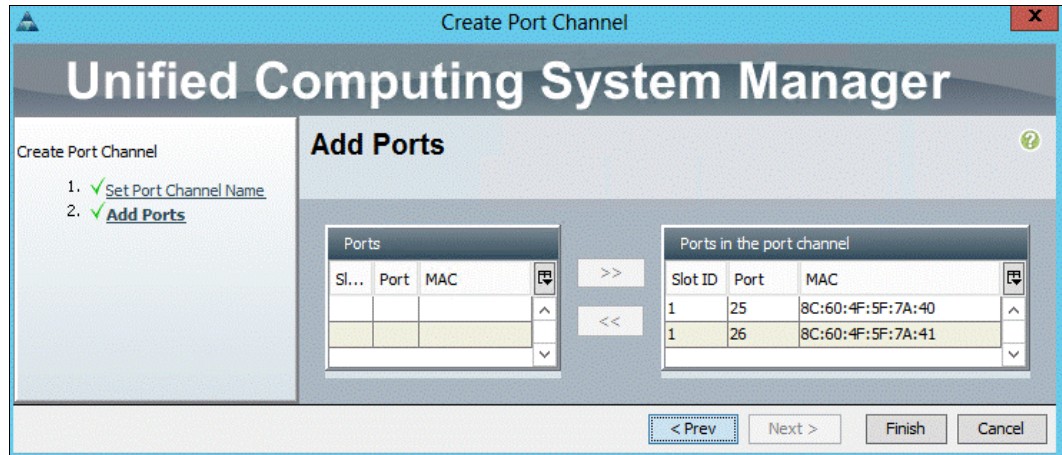


Figure 8-48 Click Finish

10. Click **OK**.
11. In the navigation pane, click **LAN** → **LAN Cloud** and expand **Fabric B**.
12. Right-click **Port Channels** and select **Create Port Channel**. The window that is shown in Figure 8-49 opens.

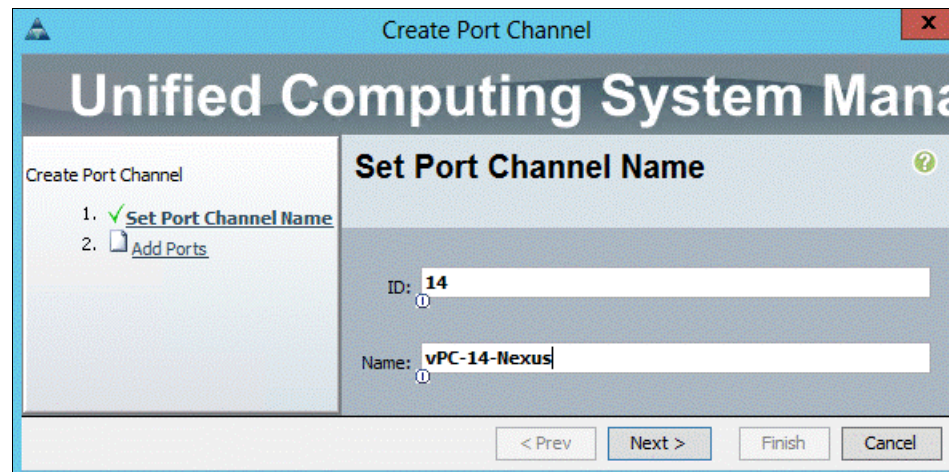


Figure 8-49 Set Port Channel Name

13. Enter 14 as the unique ID of the port channel.
14. Enter vPC-14-NEXUS as the name of the port channel.
15. Click **Next**.
16. Select the following ports to be added to the port channel:
 - Slot ID 1 and port 25
 - Slot ID 1 and port 26

17. Click **>>** to add the ports to the port channel.
18. Click **Finish** to create the port channel.
19. Click **OK**.

8.2.18 Creating MAC address pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Click **Pools** → **root**.

Note: In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click **MAC Pools** under the root organization and select **Create MAC Pool** to create the MAC address pool, as shown in Figure 8-50.

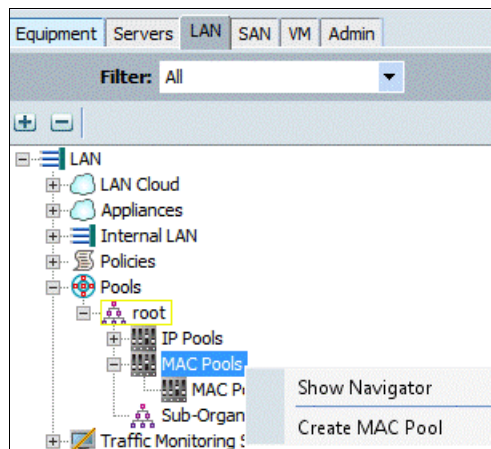


Figure 8-50 Create MAC Pool

4. Enter **MAC_Pool_A** as the name of the MAC pool.
5. (Optional) Enter a description for the MAC pool.
6. Click **Next**.

7. Click **Add**. The window that is shown in Figure 8-51 opens.

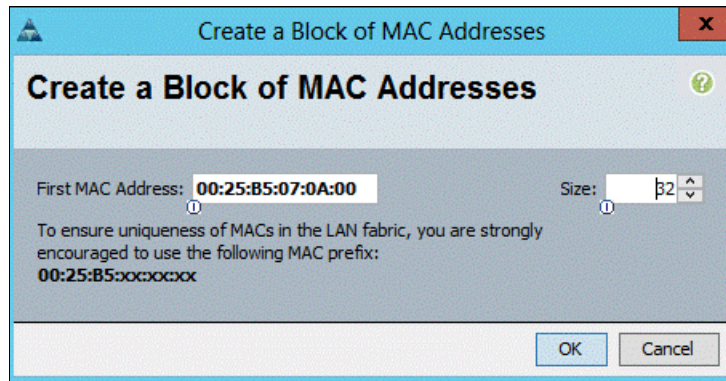


Figure 8-51 MAC address size pool

8. Specify a starting MAC address.

Note: For the VersaStack solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as Fabric A addresses.

9. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.

10. Click **OK**.

11. Click **Finish**.

12. In response to the confirmation message, click **OK**.

13. Right-click **MAC Pools** under the root organization and select **Create MAC Pool** to create the MAC address pool.

14. Enter MAC_Pool_B as the name of the MAC pool.

15. (Optional) Enter a description for the MAC pool.

16. Click **Next**.

17. Click **Add**. The window that is shown in Figure 8-52 opens.

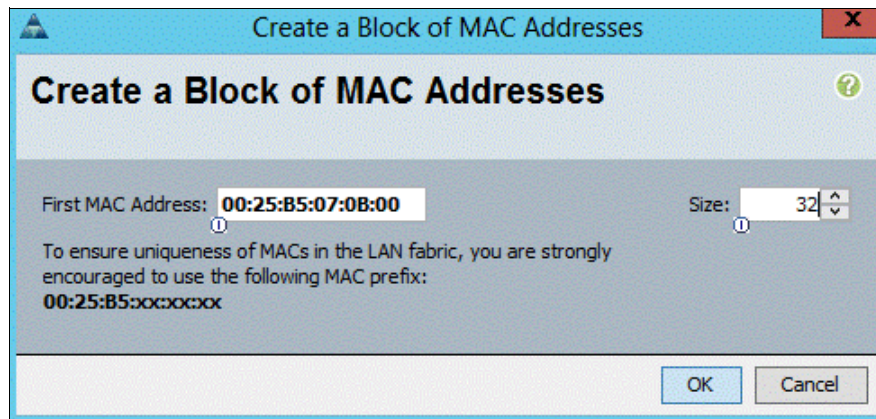


Figure 8-52 MAC address size pool (2)

18. Specify a starting MAC address.

Note: For the VersaStack solution, the recommendation is to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as Fabric B addresses.

19. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
20. Click **OK**.
21. Click **Finish**.
22. In response to the confirmation message, click **OK**.

Figure 8-53 shows the results of MAC pool creation.

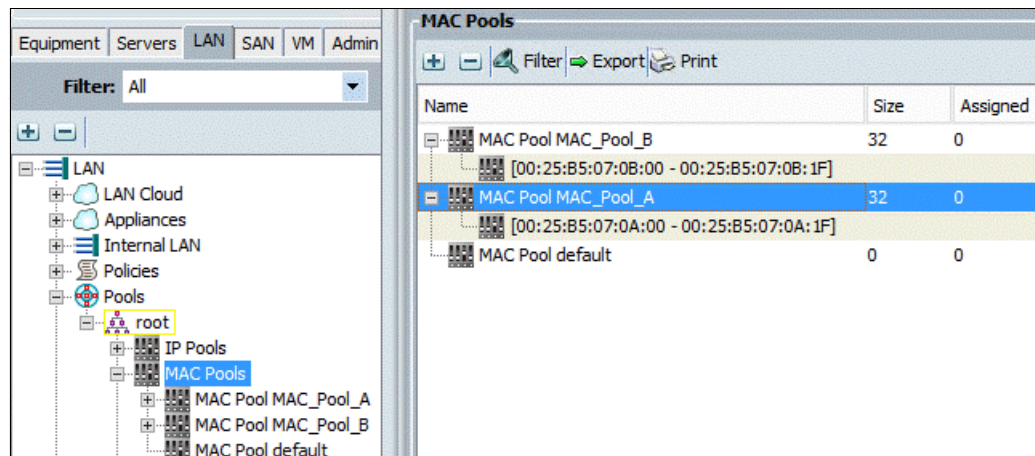


Figure 8-53 MAC pools created

8.2.19 Creating an UUID suffix pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Pools** → **root**.

3. Right-click **UUID Suffix Pools** and select **Create UUID Suffix Pool**, as shown in Figure 8-54.

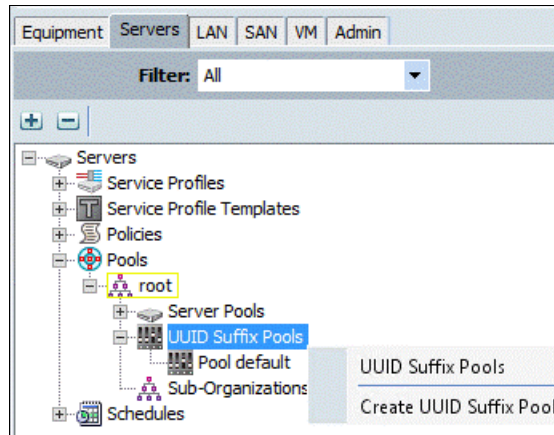


Figure 8-54 Create UUID Suffix Pool

4. Enter **UUID_Pool** as the name of the UUID suffix pool.
5. (Optional) Enter a description for the UUID suffix pool.
6. Keep the prefix at the derived option.
7. Click **Next**.
8. Click **Add** to add a block of UUIDs.
9. Keep the From field at the default setting.
10. Specify a size for the UUID block that is sufficient to support the available blade or server resources, as shown in Figure 8-55.

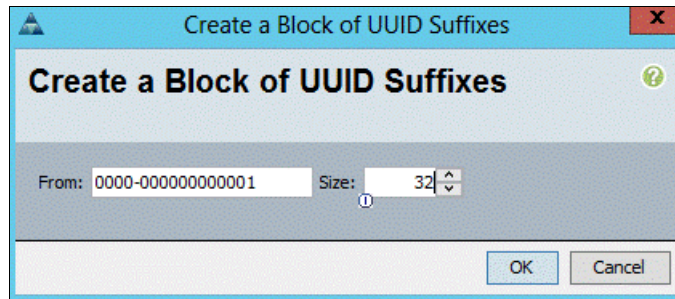


Figure 8-55 Add a block of UUIDs

11. Click **OK**.
12. Click **Finish**, as shown in Figure 8-56 on page 97, and click **OK**.

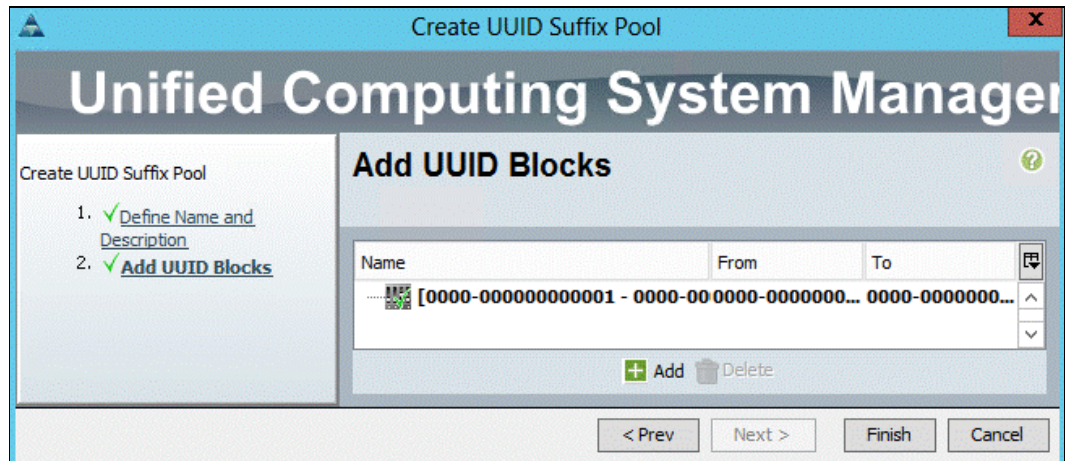


Figure 8-56 Add UUID blocks

8.2.20 Creating a server pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps.

Note: Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Pools** → **root**.
3. Right-click **Server Pools** and select **Create Server Pool**, as shown in Figure 8-57.

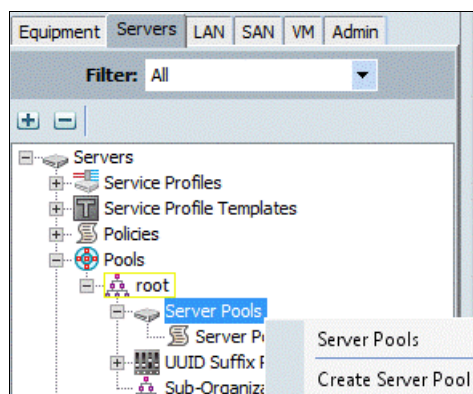


Figure 8-57 Create Server Pool

4. Enter **Infra_Pool** as the name of the server pool.
5. (Optional) Enter a description for the server pool.
6. Click **Next**.
7. Select two (or more) servers to be used for the VMware management cluster and click **>>** to add them to the **Infra_Pool** server pool.
8. Click **Finish**.
9. Click **OK**.

Figure 8-58 shows the results of this procedure.

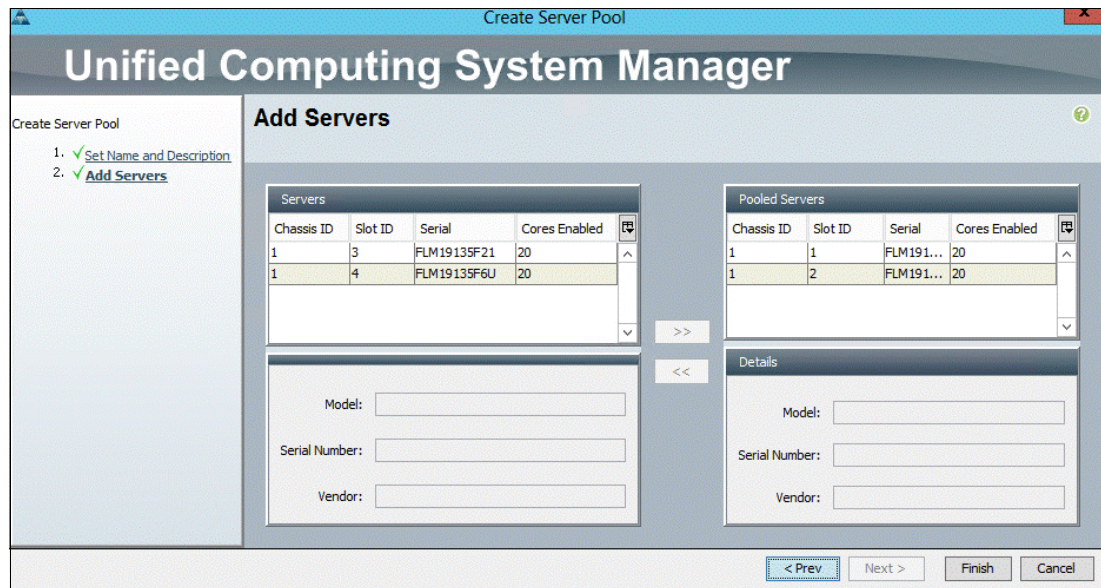


Figure 8-58 Create a server pool

8.2.21 Creating virtual local area networks

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.

Note: In this procedure, four VLANs are created. The default VLAN ID 1 is used for Management, VLAN ID 30 is used for vMotion traffic, VLAN ID 40 is for Windows Cluster traffic, VLAN ID 50 is used for CSV traffic, and VLAN ID 60 is used for Backup traffic.

2. Click **LAN** → **LAN Cloud**.
3. Right-click **VLANs** and select **Create VLANs**, as shown in Figure 8-59. The window that is shown in Figure 8-60 on page 99 opens.




Figure 8-59 Create VLANs

4. Enter vMotion as the name of the VLAN to be used for vMotion traffic.
5. Keep the Common/Global option selected for the scope of the VLAN.
6. Enter <<var_vMotion_vlan_id>> as the ID of the management VLAN.
7. Keep the Sharing Type as None.
8. Click **OK**, and then click **OK** again.

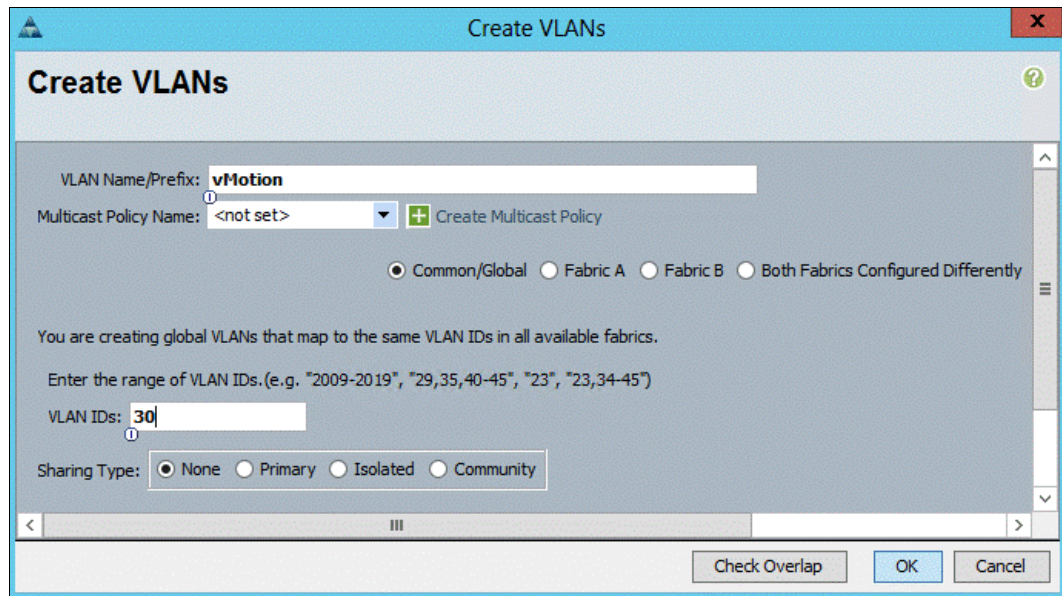


Figure 8-60 Create VLANs

9. Right-click **VLANs** and select **Create VLANs**.
10. Enter WinClus as the name of the VLAN to be used for Windows Cluster heartbeat traffic.
11. Keep the Common/Global option selected for the scope of the VLAN.
12. Enter the <<var_WinClus_vlan_id>> for the Windows Cluster VLAN.
13. Keep the Sharing Type as None.
14. Click **OK**, and then click **OK** again.
15. Right-click **VLANs** and select **Create VLANs**.
16. Enter WinCSV as the name of the VLAN to be used for Cluster Shared Volume traffic.
17. Keep the Common/Global option selected for the scope of the VLAN.
18. Enter the <<var_WinCSV_vlan_id>> as the ID of the CSV VLAN.
19. Keep the Sharing Type as None.
20. Click **OK**, and then click **OK** again.
21. Right-click **VLANs** and select **Create VLANs**.
22. Enter Backup as the name of the VLAN to be used for the Backup traffic.
23. Keep the Common/Global option selected for the scope of the VLAN.
24. Enter the <<var_Backup_vlan_id>> for the Backup VLAN.
25. Keep the Sharing Type as None.
26. Click **OK**, and then click **OK** again.

Figure 8-61 shows the final result of this procedure.

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN Backup (60)	60	Lan	Ether	No	None
VLAN WinCSV (50)	50	Lan	Ether	No	None
VLAN WinClus (40)	40	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN vMotion (30)	30	Lan	Ether	No	None

Figure 8-61 VLANs created

8.2.22 Creating a host firmware package

The administrator can use firmware management policies to select the corresponding packages for a server configuration. These policies often include packages for adapter, BIOS, system board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **Host Firmware Packages** and select **Create Host Firmware Package**, as shown in Figure 8-62.

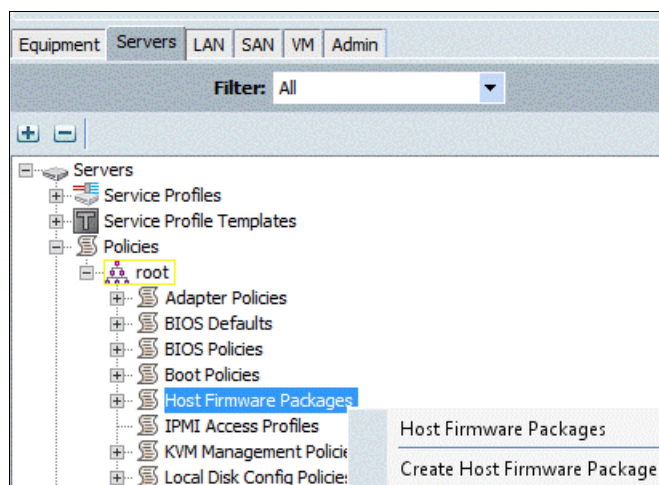


Figure 8-62 Create Host Firmware Package

4. Enter VM-Host-Infra as the name of the host firmware package.
5. Leave **Simple** selected.

6. Select the Version 2.2(3c) for Blade Servers. Also, select Version 2.2(3c) for the Rack Package if you use rack servers.
7. Click **OK** to create the host firmware package, and click OK again.

Figure 8-63 shows the final result of this procedure.

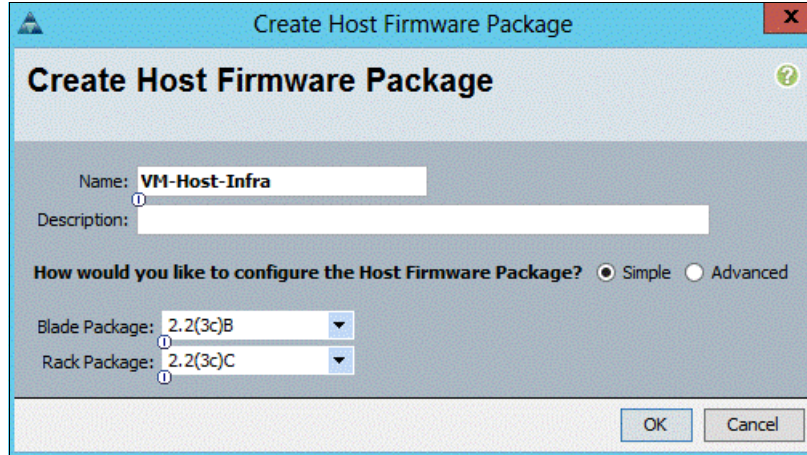


Figure 8-63 Firmware packages that are created

8.2.23 Setting jumbo frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service (QoS) in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Click **LAN** → **LAN Cloud** → **QoS System Class**. The window that is shown in Figure 8-64 opens.

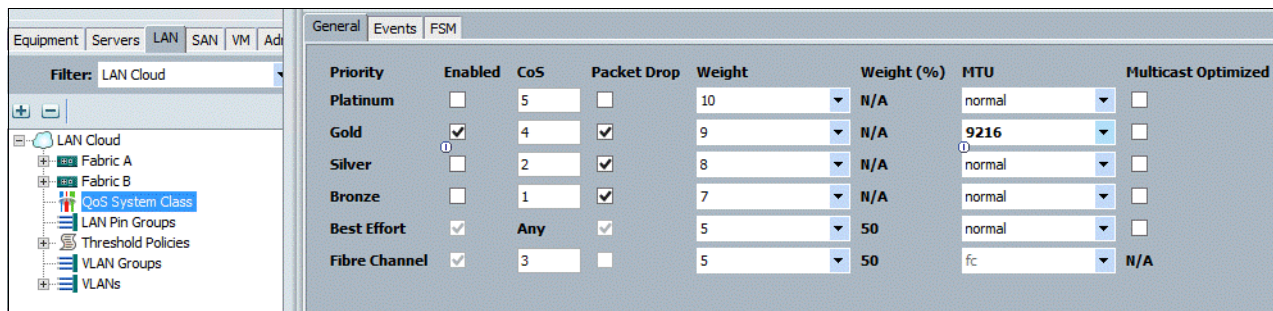


Figure 8-64 QoS System Class

3. In the right pane, click the **General** tab.
4. On the Best Effort row, enter 9216 in the field under the MTU column.
5. Click **Save Changes** at the bottom of the window.

6. Click **OK**. The window that is shown in Figure 8-65 opens.

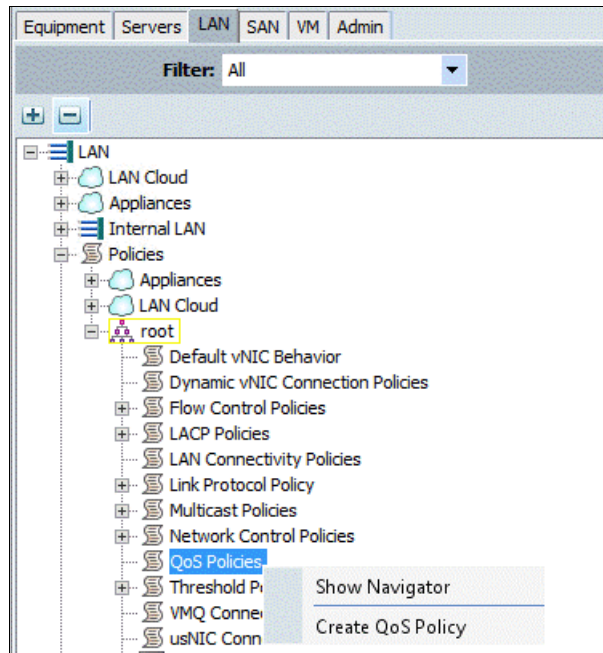


Figure 8-65 Create a QoS Policy

7. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.

8. Click **LAN** → **Policies** → **root** → **QoS Policies**.

9. Right-click **QoS Policies** and select **Create QoS Policy**.

10. Enter a name and select **Gold** from the drop-down list as the Priority and leave the rest of the settings at their defaults, as shown in Figure 8-66.

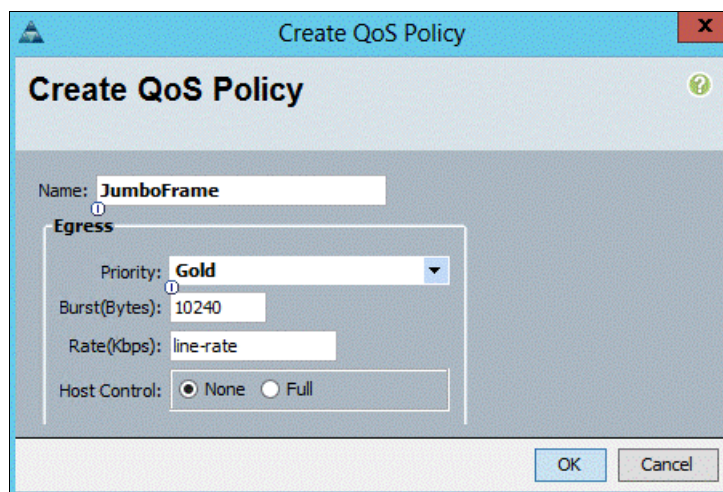


Figure 8-66 Select Gold

8.2.24 Creating a local disk configuration policy

The procedure in this section creates a SAN boot disk policy. A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.

Note: This policy should not be used on servers that contain local disks.

To create a local disk configuration policy for SAN-Boot, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **Local Disk Config Policies** and select **Create Local Disk Configuration Policy**, as shown in Figure 8-67.

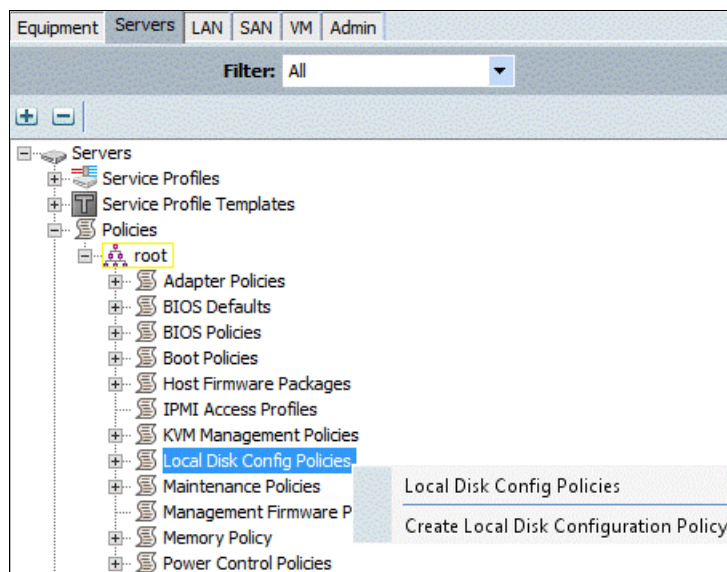


Figure 8-67 Create Local Disk Configuration Policy

The window that is shown in Figure 8-68 opens.

The screenshot shows a dialog box titled "Create Local Disk Configuration Policy". The main heading is "Create Local Disk Configuration Policy". The form fields are: Name: SAN-Boot; Description: (empty); Mode: No Local Storage (dropdown menu); FlexFlash section with FlexFlash State: Disable (selected) and Enable; and FlexFlash RAID Reporting State: Disable (selected). There are OK and Cancel buttons at the bottom right.

Figure 8-68 Create the policy

4. Enter SAN-Boot as the local disk configuration policy name.
5. Change the mode to **No Local Storage**.
6. Click **OK** to create the local disk configuration policy, and click **OK** again.

8.2.25 Creating a network control policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **Network Control Policies** and select Create Network Control Policy, as shown in Figure 8-69 on page 105. The window that is shown in Figure 8-70 on page 105 opens.

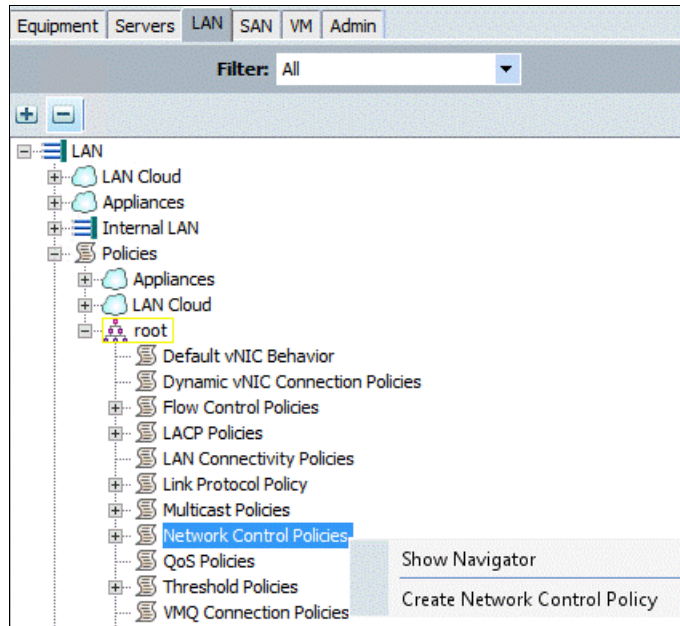


Figure 8-69 Create Network Control Policy

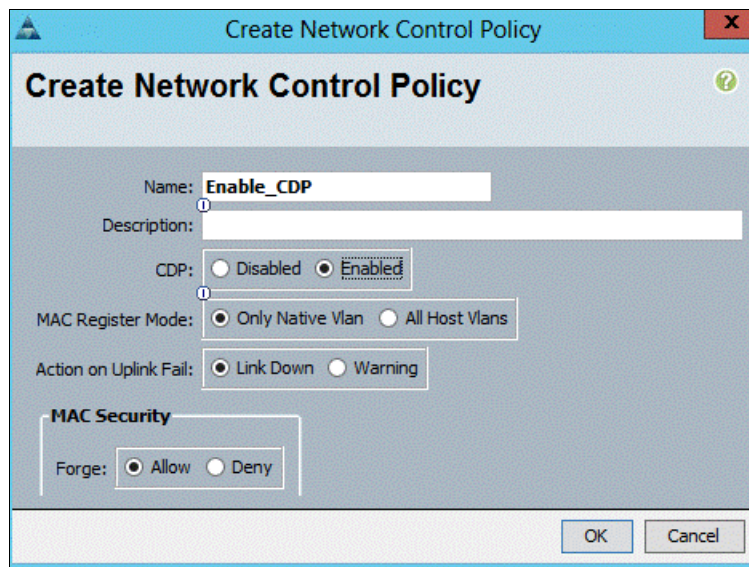


Figure 8-70 Create Network Control Policy

4. Enter Enable_CDP as the policy name.
5. For CDP, select the **Enabled** option.
6. Click **OK** to create the network control policy, and click OK again.

8.2.26 Creating a power control policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.

3. Right-click **Power Control Policies** and select Create Power Control Policy. The window that is shown in Figure 8-71 opens.

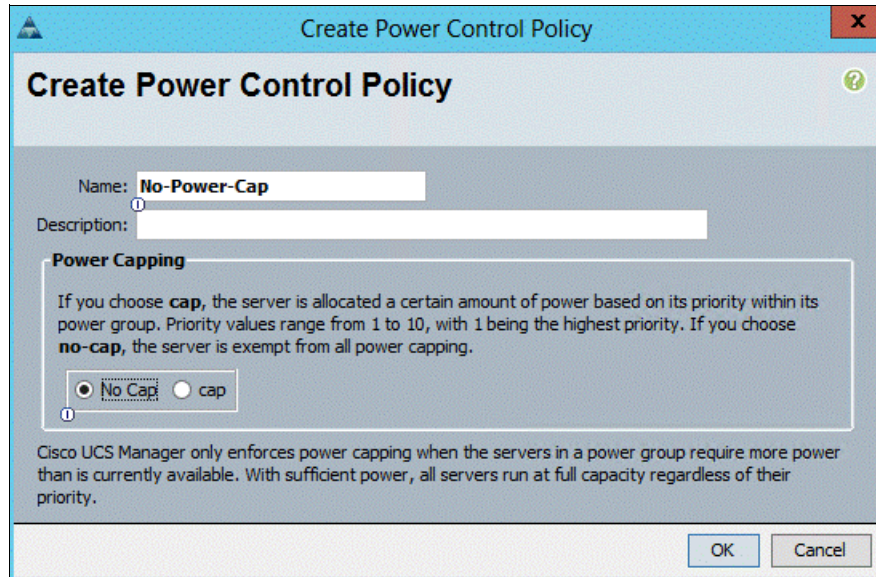


Figure 8-71 Create Power Control Policy

4. Enter No-Power-Cap as the power control policy name.
5. Change the Power Capping setting to **No Cap**.
6. Click **OK** to create the power control policy, and click **OK**.

8.2.27 Creating a server pool qualification policy (optional)

To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **Server Pool Policy Qualifications** and select **Create Server Pool Policy Qualification**. The window that is shown in Figure 8-72 on page 107 opens.

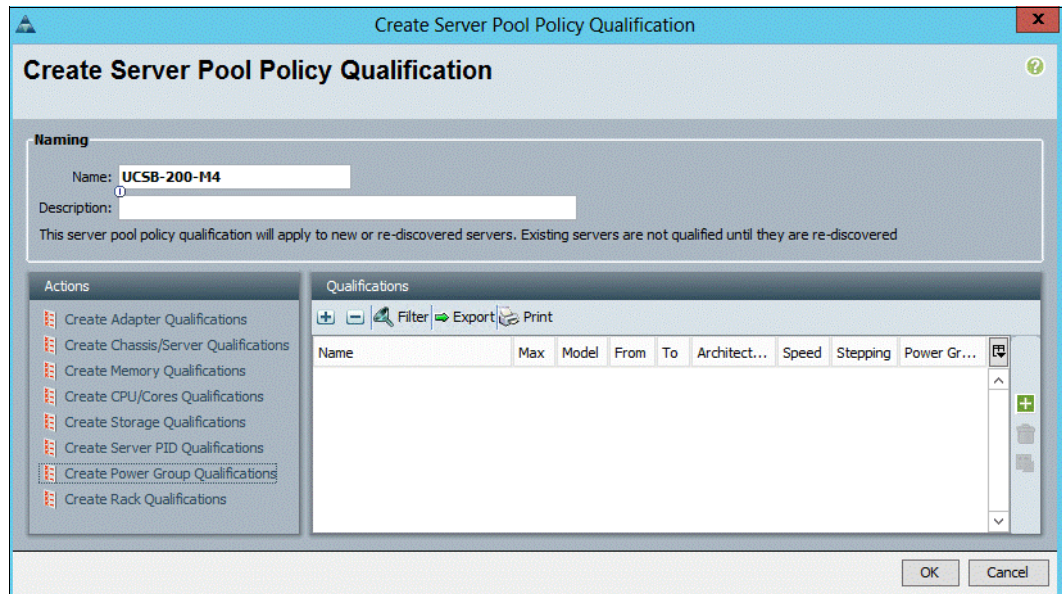


Figure 8-72 Create Server Pool Policy Qualification

4. Enter UCSB-B200-M4 as the name for the policy.
5. Select **Create Server PID Qualifications**. The window that is shown in Figure 8-73 opens.

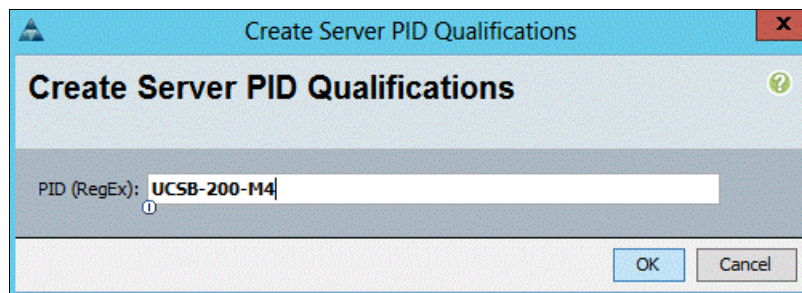


Figure 8-73 Create Server PID Qualifications

6. Enter UCSB-B200-M4 as the PID.
7. Click **OK** to create the server pool qualification policy.
8. Click **OK**, and then click **OK** again.

8.2.28 Creating a server BIOS policy

The following policies are for optimal performance for VMware. Depending on your requirements, you can change the settings as needed. For more information, see your Cisco UCS documentation.

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **BIOS Policies** and select **Create BIOS Policy**.

4. Enter VM-Host-Infra as the BIOS policy name and select the options that are shown in Figure 8-74.

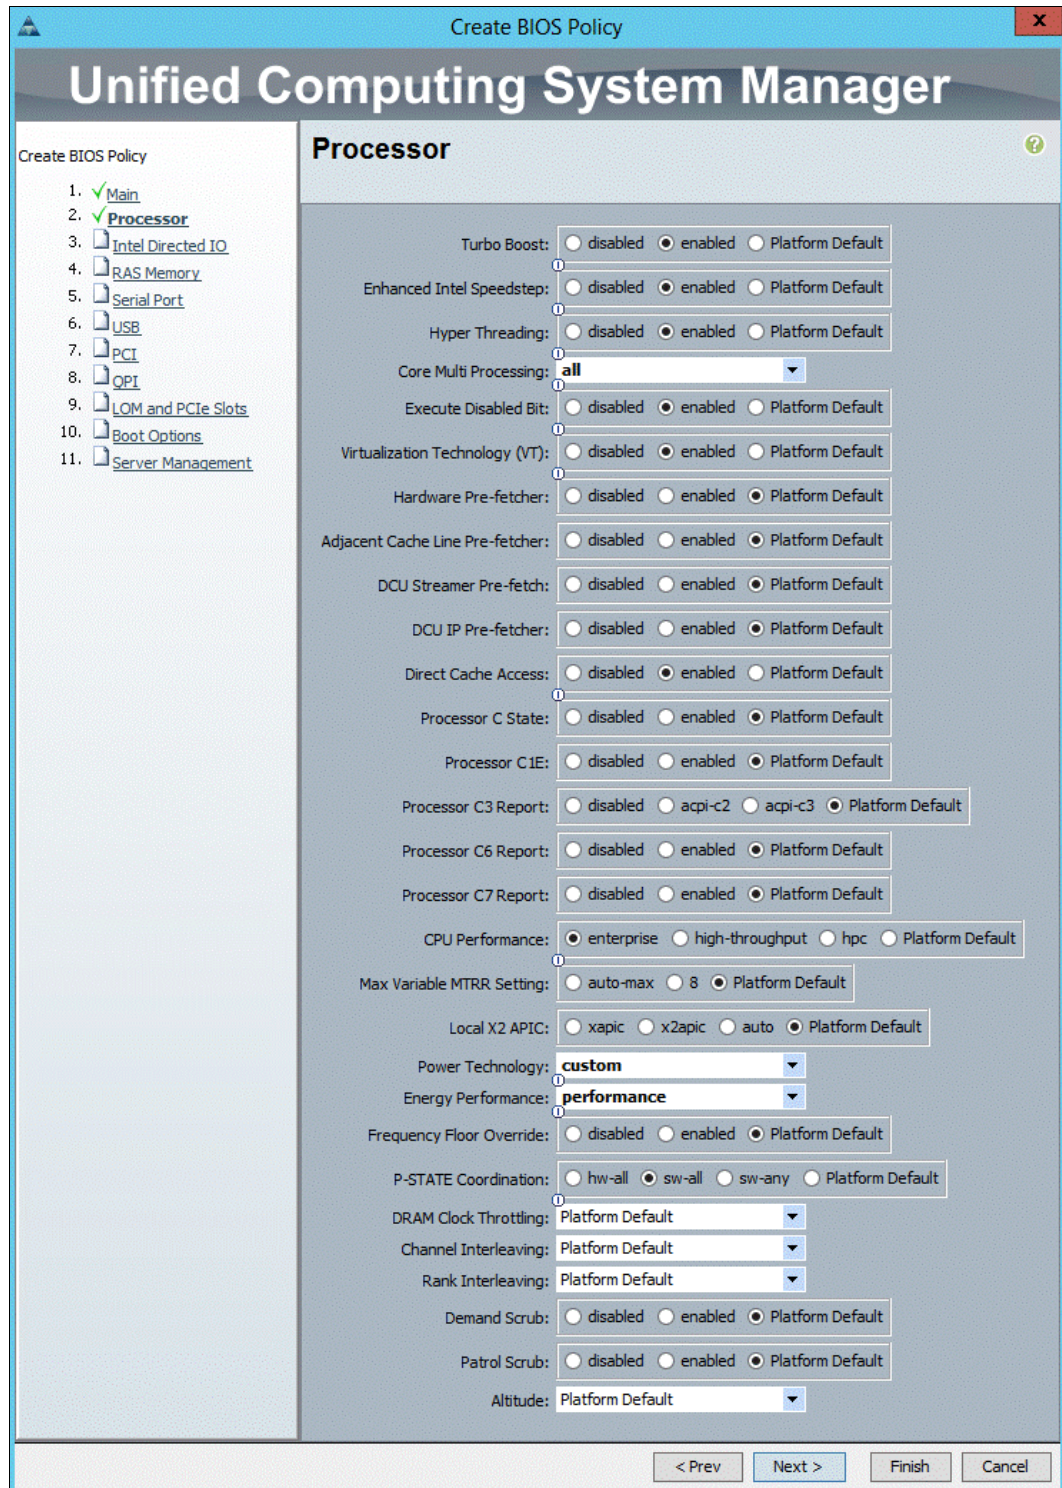


Figure 8-74 Create BIOS Policy

5. Click **Next** to open the Intel Directed IO window and select the options that are shown in Figure 8-75 on page 109.

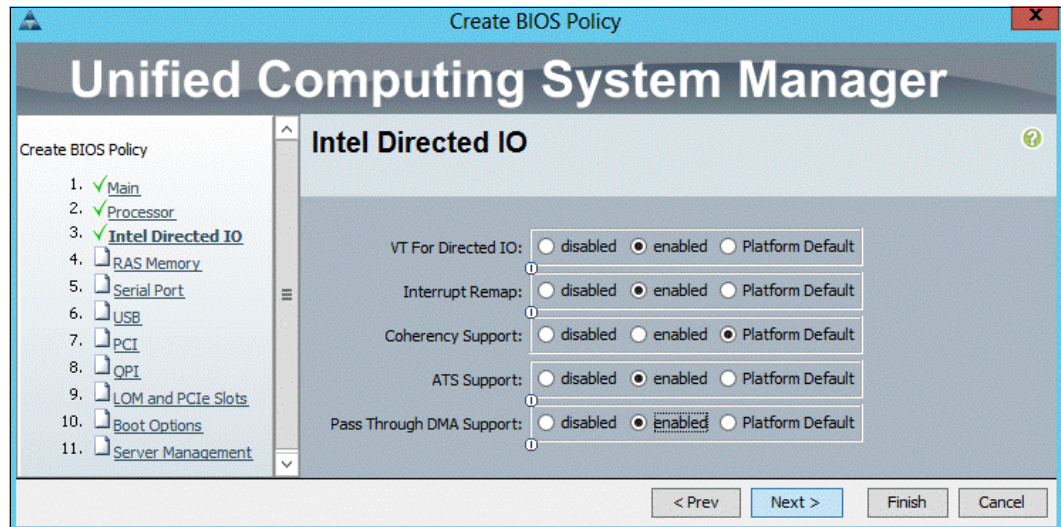


Figure 8-75 Intel Directed IO

6. Click **Next** to open the RAS Memory window and select the options that are shown in Figure 8-76.

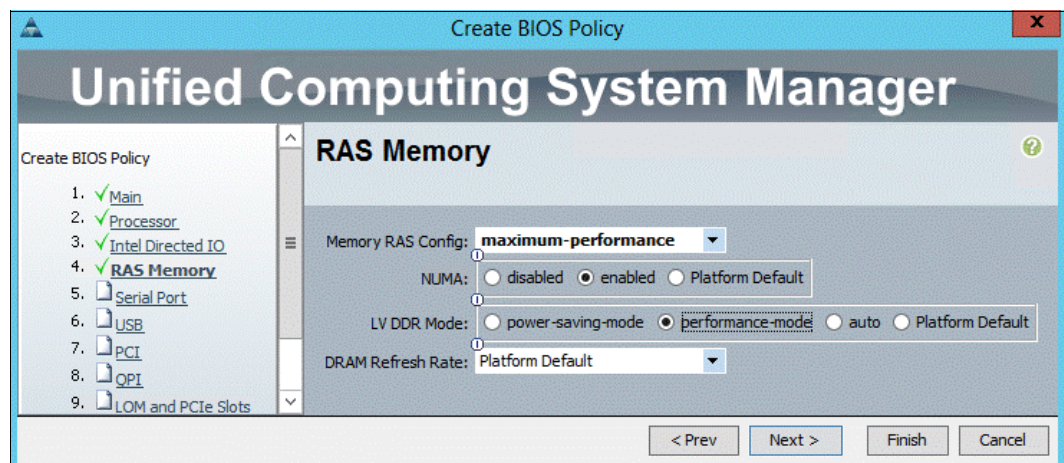


Figure 8-76 RAS Memory

7. Click **Finish** to create the BIOS policy.
8. Click **OK**.

8.2.29 Creating a vNIC/vHBA placement policy for VM infrastructure hosts

To create a vNIC/vHBA placement policy for the infrastructure hosts, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.

3. Right-click **vNIC/vHBA Placement Policies** and select **Create Placement Policy**. The window that is shown in Figure 8-77 opens.

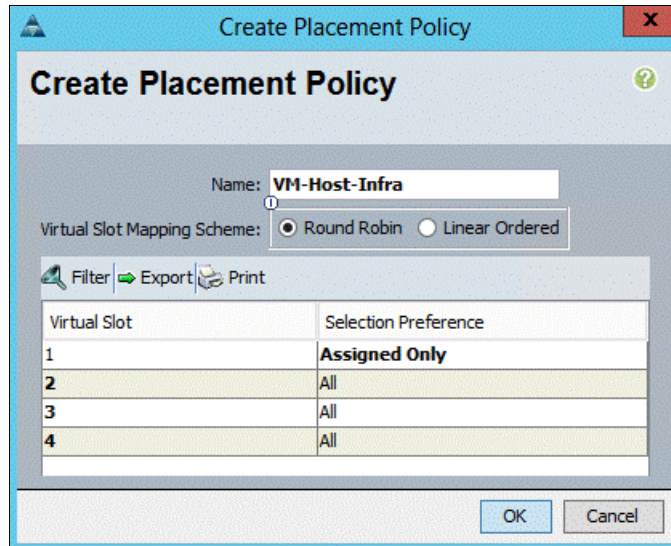


Figure 8-77 Create Placement Policy

4. Enter VM-Host-Infra as the name of the placement policy.
5. Click **1** and select **Assigned Only**.
6. Click **OK**, and then click **OK** again.

8.2.30 Updating the default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane. Figure 8-78 on page 111 shows the Servers tab.
2. Click **Policies** → **root**.
3. Click **Maintenance Policies** → **default**.
4. Change the Reboot Policy to **User Ack**.
5. Click **Save Changes**.
6. Click **OK** to accept the change.

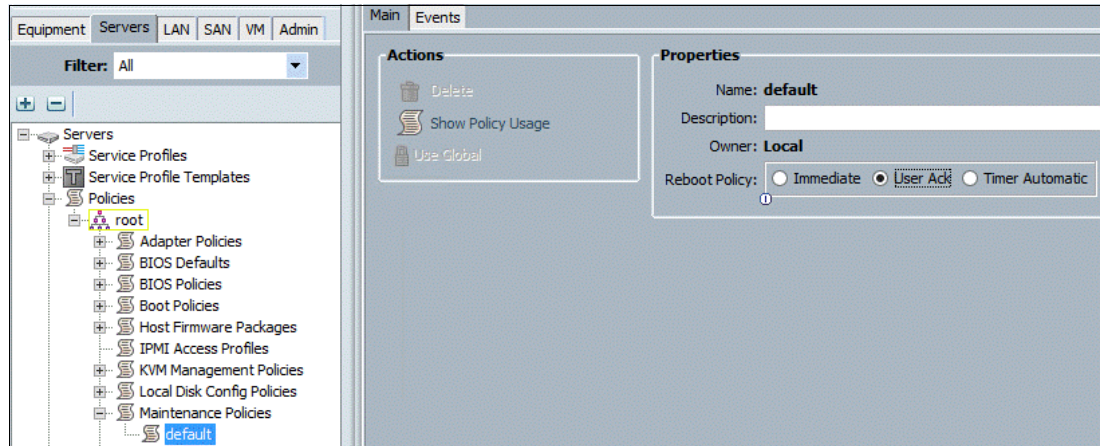


Figure 8-78 Servers tab

8.2.31 Creating vNIC templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps.

Note: Do not select the Enable Failover option if the network adapters will be teamed up later in the OS/hypervisor. In this example, because we are teaming the vNICs in this VersaStack environment, the Enable Failover option is left clear.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Click **Policies** → **root**.

- Right-click **vNIC Templates** and select **Create vNIC Template**. The window that is shown in Figure 8-79 opens.

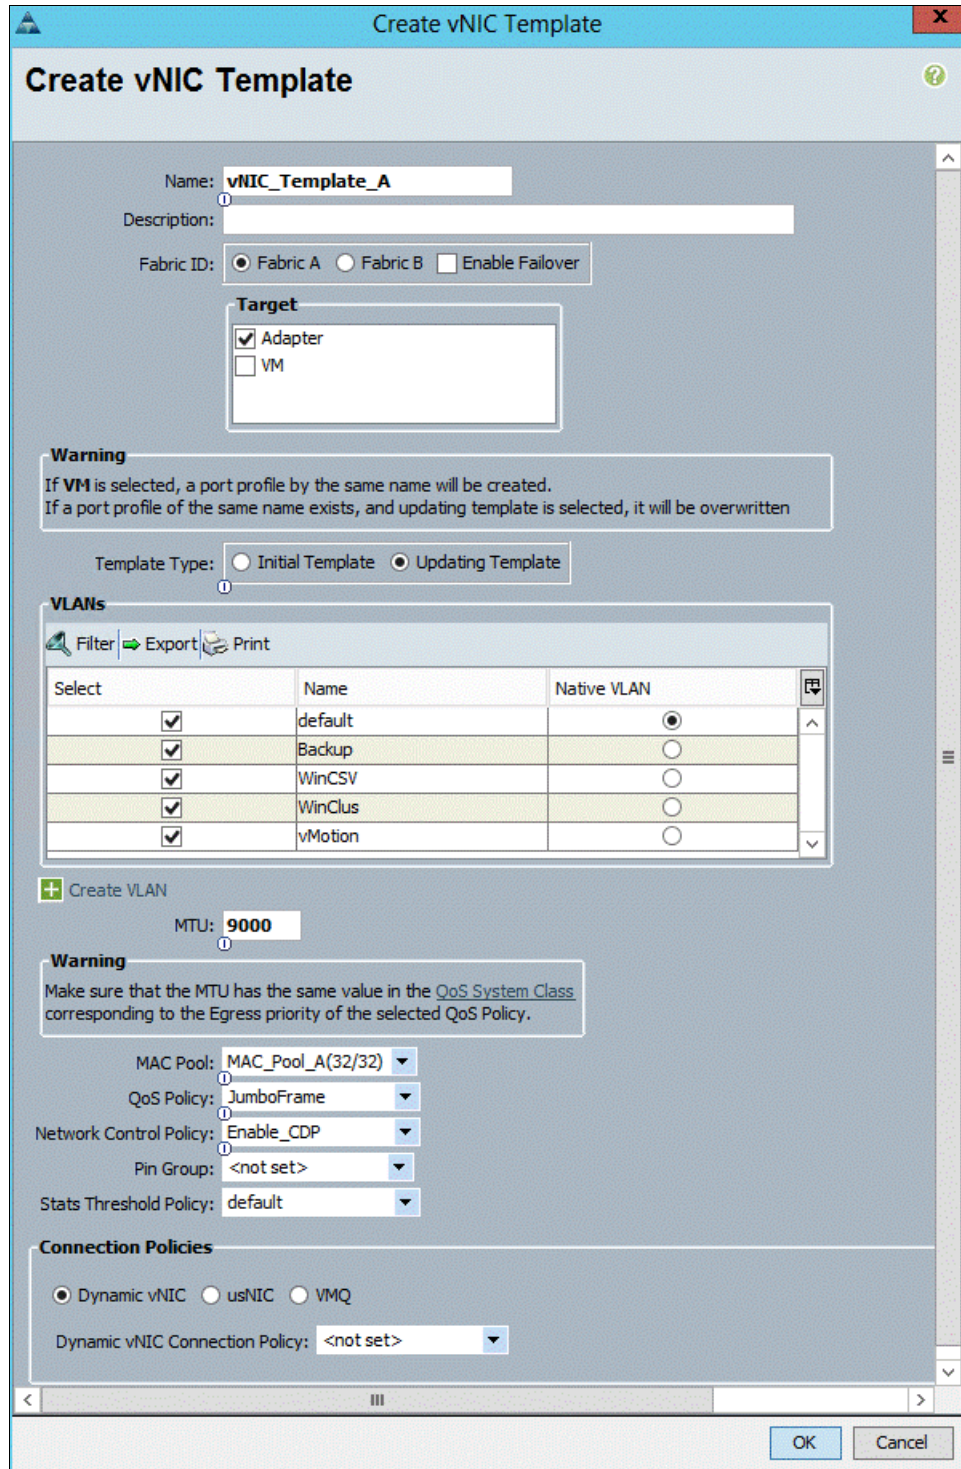


Figure 8-79 Create the vNIC template

- Enter vNIC_Template_A as the vNIC template name.
- Keep **Fabric A** selected.
- Select the **Enable Failover** check box.

7. Under Target, make sure that the VM check box is not selected.
8. Select **Updating Template** as the Template Type.
9. Under VLANs, select the check boxes for **Default (Mgmt)**, **WinClus**, **WinCSV**, and **Backup**.
10. Set Default as the native VLAN.
11. For MTU, enter 9000.
12. In the MAC Pool list, select **MAC_Pool_A**.
13. In the Network Control Policy list, select **Enable_CDP**.
14. Click **OK** to create the vNIC template, and click OK again.
15. In the navigation pane, select the **LAN** tab.
16. Click **Policies** → **root**.

17. Right-click **vNIC Templates** and select **Create vNIC Template**. The window that is shown in Figure 8-80 opens.

Create vNIC Template

Name:

Description:

Fabric ID: Fabric A Fabric B Enable Failover

Target

Adapter
 VM

Warning
If VM is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type: Initial Template Updating Template

VLANs

Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	Backup	<input type="radio"/>
<input checked="" type="checkbox"/>	WinCSV	<input type="radio"/>
<input checked="" type="checkbox"/>	WinClus	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>

+ Create VLAN

MTU:

Warning
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

MAC Pool:

QoS Policy:

Network Control Policy:

Pin Group:

Stats Threshold Policy:

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy:

OK Cancel

Figure 8-80 Create vNIC Template

18. Enter vNIC_Template_B as the vNIC template name.

19. Select **Fabric B**.

20. Select the **Enable Failover** check box.

21. Select **Updating Template** as the template type.
22. Under VLANs, select the check boxes for **Default (Mgmt)**, **WinClus**, **WinCSV**, and **Backup**.
23. Set Default as the native VLAN.
24. For MTU, enter 9000.
25. In the MAC Pool list, select **MAC_Pool_B**.
26. In the Network Control Policy list, select **Enable_CDP**.
27. Click OK to create the vNIC template, and click **OK** again.

8.2.32 Creating boot policies

This procedure applies to a Cisco UCS environment in which two FC interfaces are on cluster node 1 and two FC interfaces are on cluster node 2.

Two boot policies are configured in this procedure. The first policy configures the primary target to be fcp_a and the second boot policy configures the primary target to be fcp_b.

To create boot policies for the Cisco UCS environment, complete the following steps.

Note: You use the WWPN variables that were logged in the storage section WWPN table.

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Policies** → **root**.
3. Right-click **Boot Policies** and select **Create Boot Policy**, as shown in Figure 8-81.

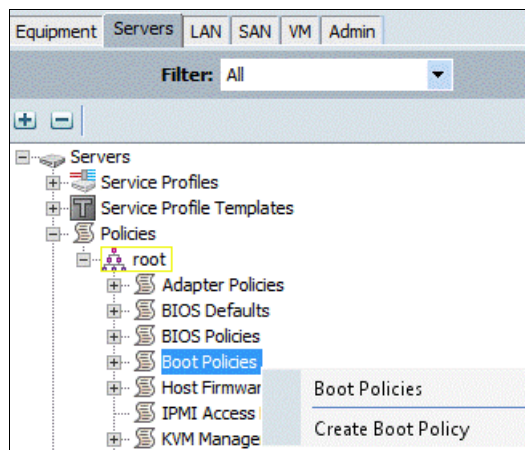


Figure 8-81 Create Boot Policy

The window that is shown in Figure 8-82 opens.

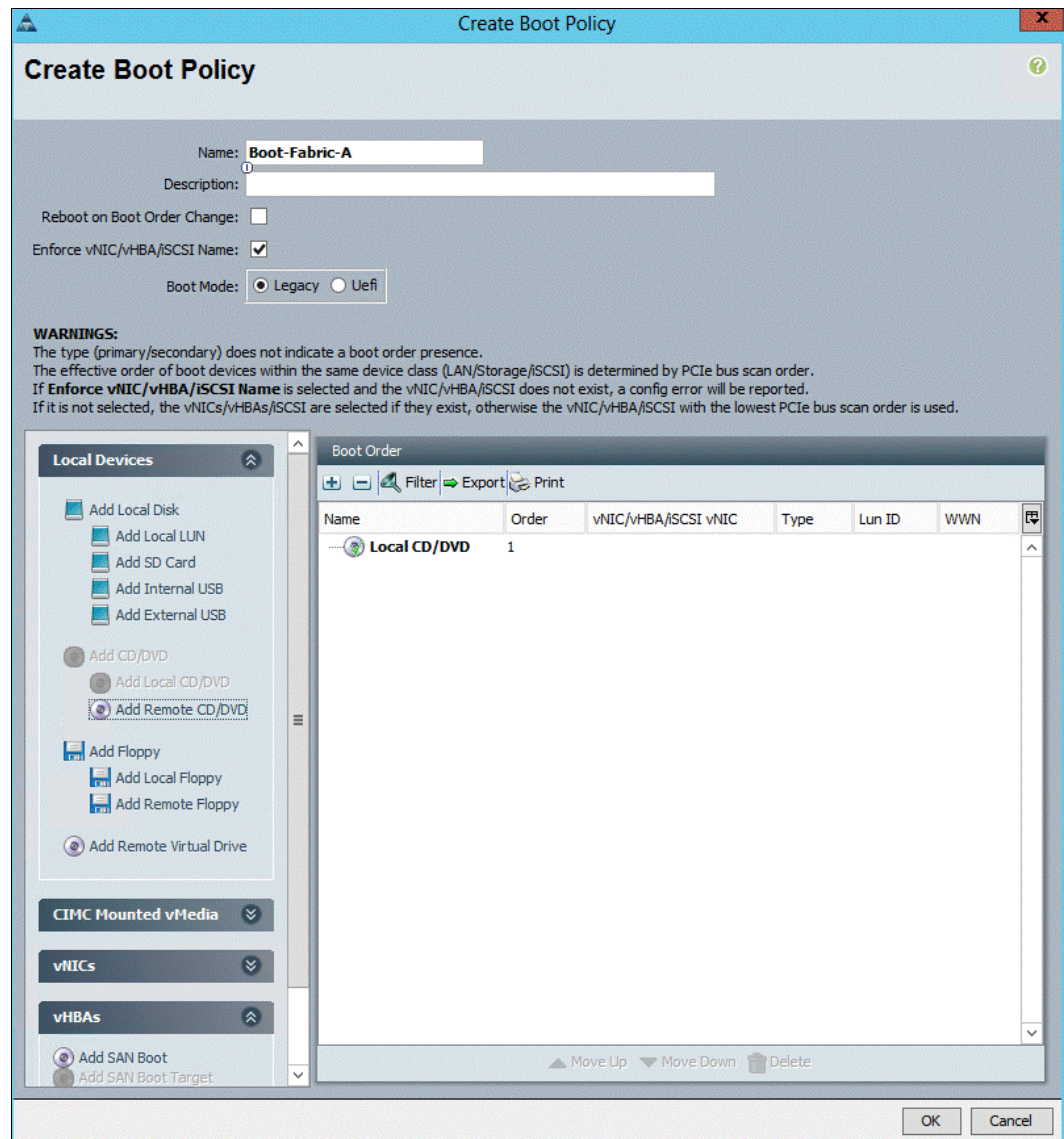


Figure 8-82 Add CD/DVD

4. Enter **Boot-Fabric-A** as the name of the boot policy.
5. (Optional) Enter a description for the boot policy.
6. Keep the **Reboot on Boot Order Change** check box clear.
7. Expand the **Local Devices** drop-down menu and click **Add CD/DVD** (you should see Local and Remote disabled).
8. Scroll down on the left side, expand the **vHBAs** drop-down menu, and click **Add SAN Boot**. The window that is shown in Figure 8-83 on page 117 opens.

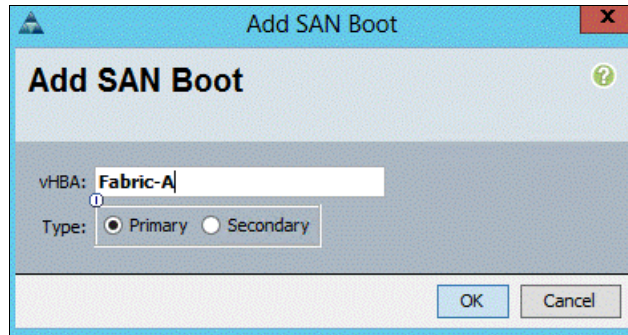


Figure 8-83 Add the SAN boot initiator

9. In the Add SAN Boot dialog box, enter Fabric-A in the vHBA field.
10. Make sure that the **Primary** radio button is selected as the SAN boot type.
11. Click **OK** to add the SAN boot initiator.
12. From the vHBA drop-down menu, select **Add SAN Boot Target**. The window that is shown in Figure 8-84 opens.

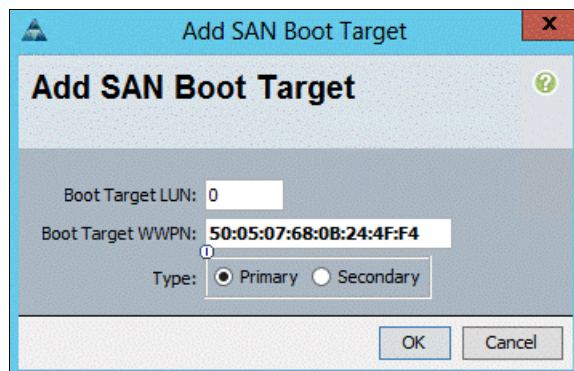


Figure 8-84 Add the primary SAN boot target

13. Keep 0 as the value for Boot Target LUN.
14. Enter the WWPN for node 1 going to switch A (<< var_wwpn_Node1-switch-A >>).
15. Keep the **Primary** radio button selected as the SAN boot target type.

16. Click **OK** to add the SAN boot target.
17. From the vHBA drop-down menu, select **Add SAN Boot Target**. The window that is shown in Figure 8-85 opens.




Figure 8-85 Add the secondary SAN boot target

18. Keep 0 as the value for Boot Target LUN.
19. Enter the WWPN for node 2 going to switch A (<< var_wwpn_Node2-switch-A >>).
20. Click **OK** to add the SAN boot target.
21. From the vHBA drop-down menu, select **Add SAN Boot**. The window that is shown in Figure 8-86 opens.




Figure 8-86 Add SAN boot

22. In the Add SAN Boot dialog box, enter Fabric-B in the vHBA field.
23. The SAN boot type should automatically be set to Secondary.
24. Click **OK** to add the SAN boot initiator..
25. From the vHBA drop-down menu, select **Add SAN Boot Target**. The window that is shown in Figure 8-87 on page 119 opens.

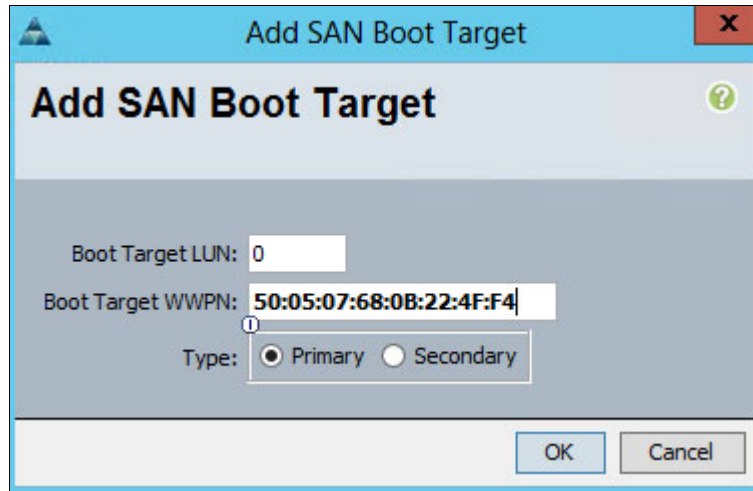


Figure 8-87 Add the primary SAN boot target

26. Keep 0 as the value for Boot Target LUN.
27. Enter the WWPN for node 2 switch B (<<var_wwpn_Node2-switch-B>>).
28. Keep Primary as the SAN boot target type.
29. Click **OK** to add the SAN boot target.
30. From the vHBA drop-down menu, select **Add SAN Boot Target**. The window that is shown in Figure 8-88 opens.

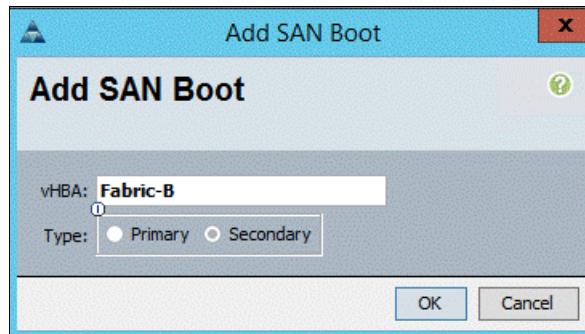


Figure 8-88 Add the secondary SAN boot target

31. Keep 0 as the value for Boot Target LUN.
32. Enter the WWPN for Node 1 switch B (<<var_wwpn_Node1-Switch-B>>).
33. Click **OK** to add the SAN boot target.

34. Click **OK**, and then **OK** again to create the boot policy, as shown in Figure 8-89.

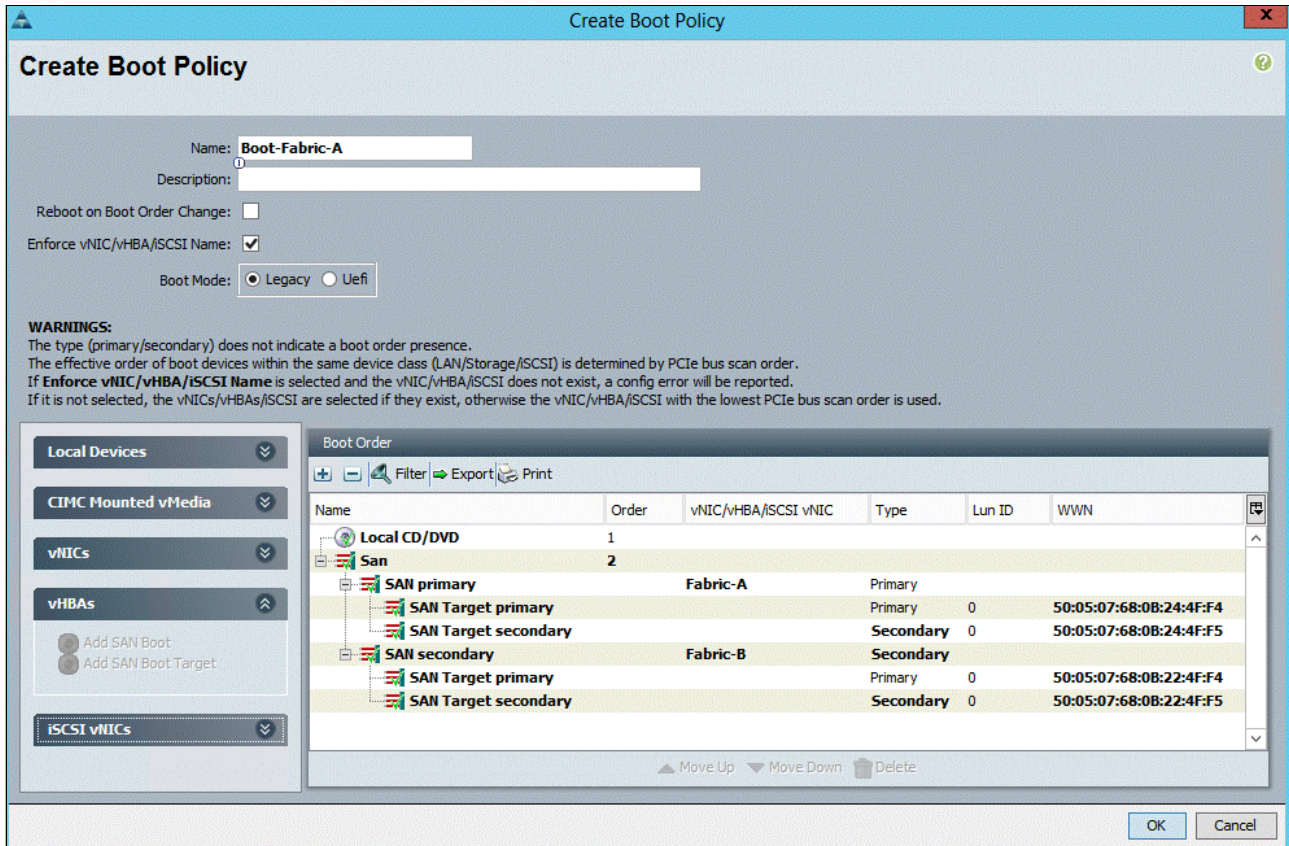


Figure 8-89 Create the boot policy

35. Right-click **Boot Policies** again and select **Create Boot Policy**. The window that is shown in Figure 8-90 on page 121 opens.

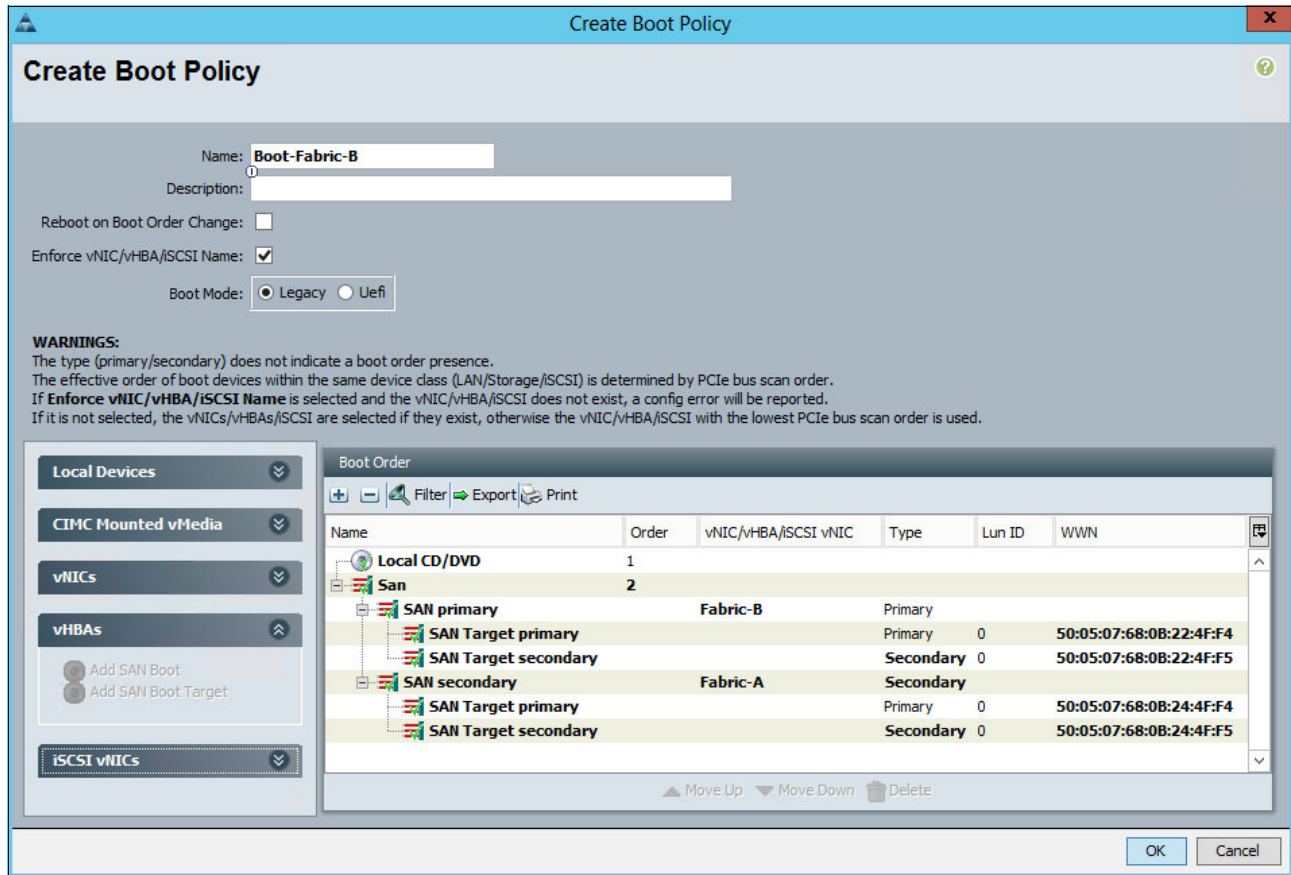


Figure 8-90 Create the boot policy

36. Enter Boot-Fabric-B as the name of the boot policy.
37. (Optional) Enter a description of the boot policy.
38. Keep the **Reboot on Boot Order Change** check box clear.
39. From the Local Devices drop-down menu, select **Add CD/DVD**.
40. From the vHBA drop-down menu, select **Add SAN Boot**.
41. In the Add SAN Boot dialog box, enter Fabric-B in to the vHBA field.
42. Make sure that the **Primary** radio button is selected as the SAN boot type.
43. Click **OK** to add the SAN boot initiator.
44. From the vHBA drop-down menu, select **Add SAN Boot Target**.
45. Keep 0 as the value for Boot Target LUN.
46. Enter the WWPN for <<var_wwpn_Node1-Switch-B>>.
47. Keep Primary as the SAN boot target type.
48. Click **OK** to add the SAN boot target.
49. From the vHBA drop-down menu, select **Add SAN Boot Target**.
50. Keep 0 as the value for Boot Target LUN.
51. Enter the WWPN for <<var_wwpn_Node2-Switch-B>>.
52. Click **OK** to add the SAN boot target.

53. From the vHBA menu, select **Add SAN Boot**.
54. In the Add SAN Boot dialog box, enter Fabric-A into the vHBA field.
55. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.
56. Click **OK** to add the SAN boot initiator.
57. From the vHBA menu, select **Add SAN Boot Target**.
58. Keep 0 as the value for Boot Target LUN.
59. Enter the WWPN for <<var_wwpn_Node2-switch-A >>.
60. Keep Primary as the SAN boot target type.
61. Click **OK** to add the SAN boot target.
62. From the vHBA drop-down menu, select **Add SAN Boot Target**.
63. Keep 0 as the value for Boot Target LUN.
64. Enter the WWPN for <<var_wwpn_Node1-switch-A >>.
65. Click **OK** to add the SAN boot target.
66. Click **OK**, and then click **OK** again to create the boot policy.

8.2.33 Creating service profile templates

In this procedure, two service profile templates are created: one for Fabric A boot and one for Fabric B boot. The first profile is created and then cloned and modified for the second host.

To create service profile templates, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Service Profile Templates** → **root**.
3. Right-click **root** and select **Create Service Profile Template** to open the Create Service Profile Template wizard. The window that is shown in Figure 8-91 on page 123 opens.

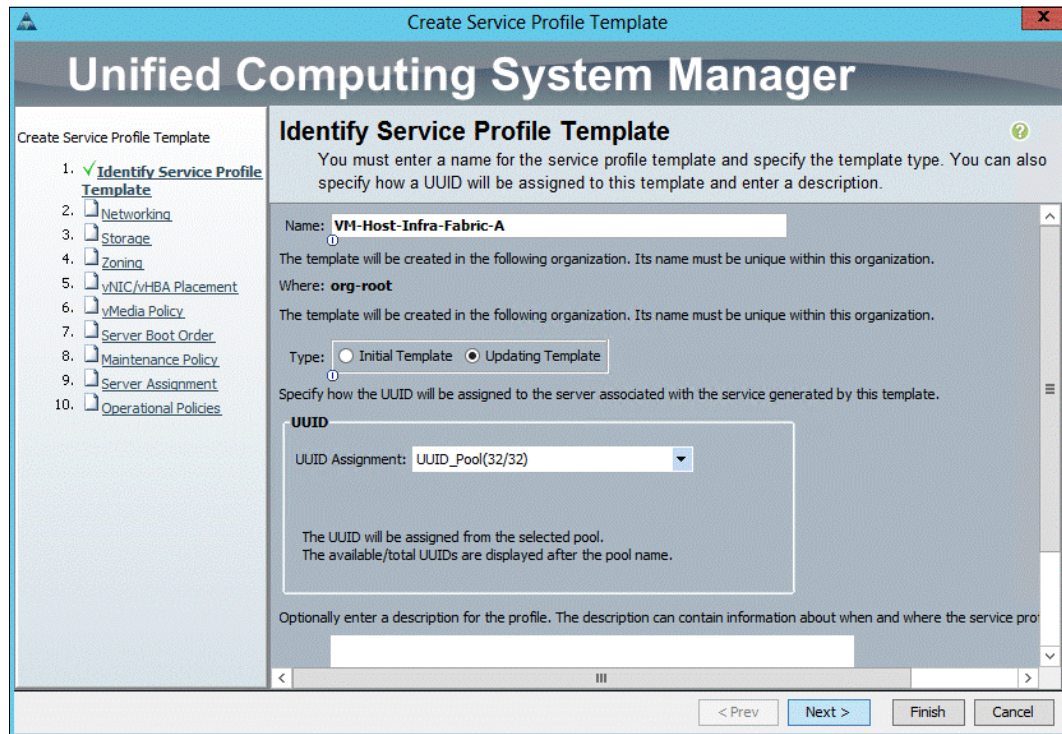


Figure 8-91 Identify the service profile template

4. Identify the Service Profile Template:
 - a. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from node 1 on fabric A.
 - b. Select the **Updating Template** radio button.
 - c. Under UUID, select **UUID_Pool** as the UUID pool.
 - d. Click **Next**.

5. Configure the Networking options:
 - a. Keep the default setting for Dynamic vNIC Connection Policy.
 - b. Select the **Expert** radio button to configure the LAN connectivity.
 - c. Click **Add** to add a vNIC to the template. The window that is shown in Figure 8-92 opens.

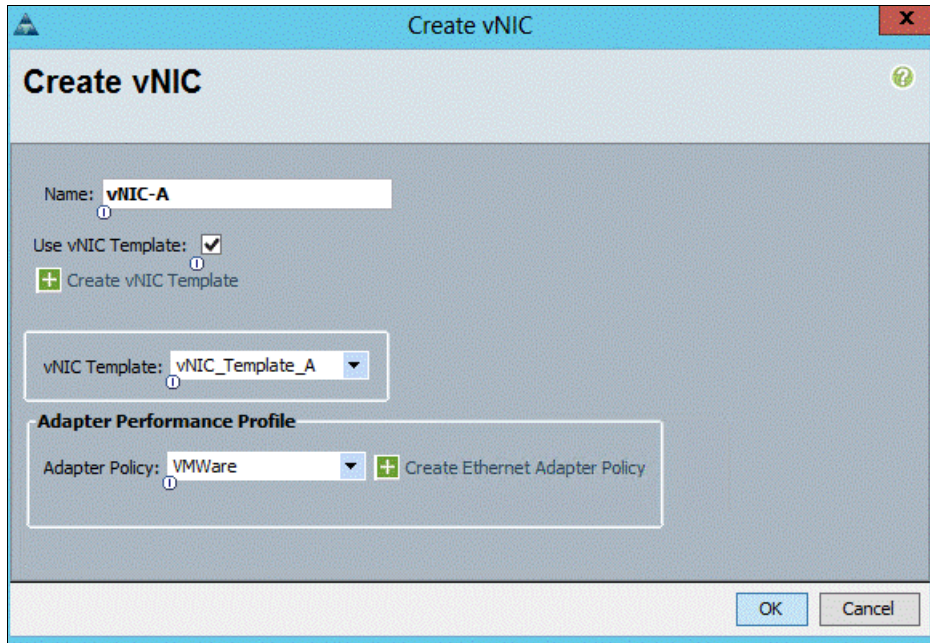


Figure 8-92 Create vNIC

- d. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
- e. Check the **Use vNIC Template** check box.
- f. In the vNIC Template list, select **vNIC_Template_A**.
- g. In the Adapter Policy list, select **VMWare**.
- h. Click **OK** to add this vNIC to the template.
- i. On the Networking window of the wizard, click **Add** to add another vNIC to the template. The window that is shown in Figure 8-93 on page 125 opens.

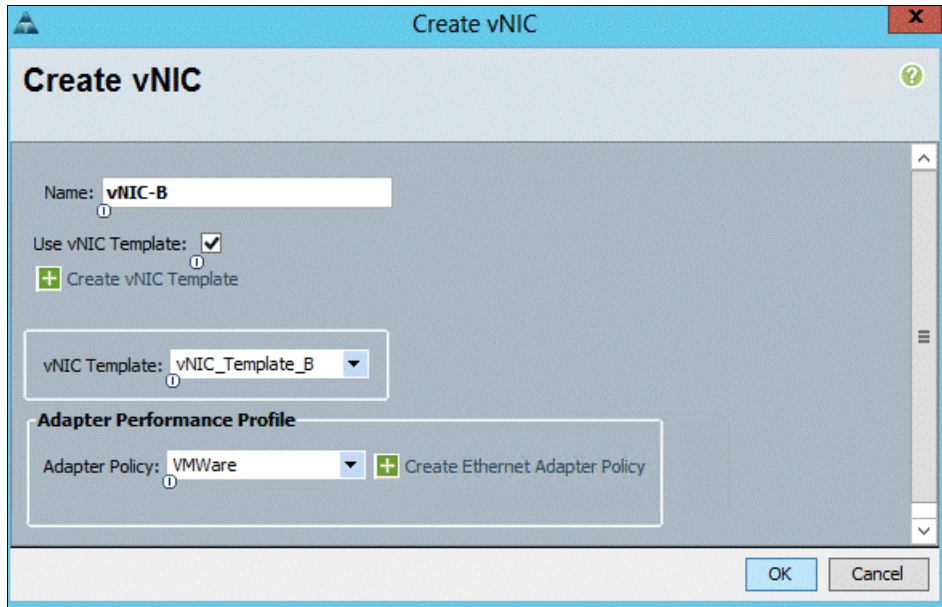


Figure 8-93 Add the vNIC to the template

- j. In the Create vNIC dialog box, enter vNIC-B as the name of the vNIC.
- k. Select the **Use vNIC Template** check box.
- l. In the vNIC Template list, select **vNIC_Template_B**.
- m. In the Adapter Policy list, select **VMWare**.
- n. Click **OK** to add the vNIC to the template. The window that is shown in Figure 8-94

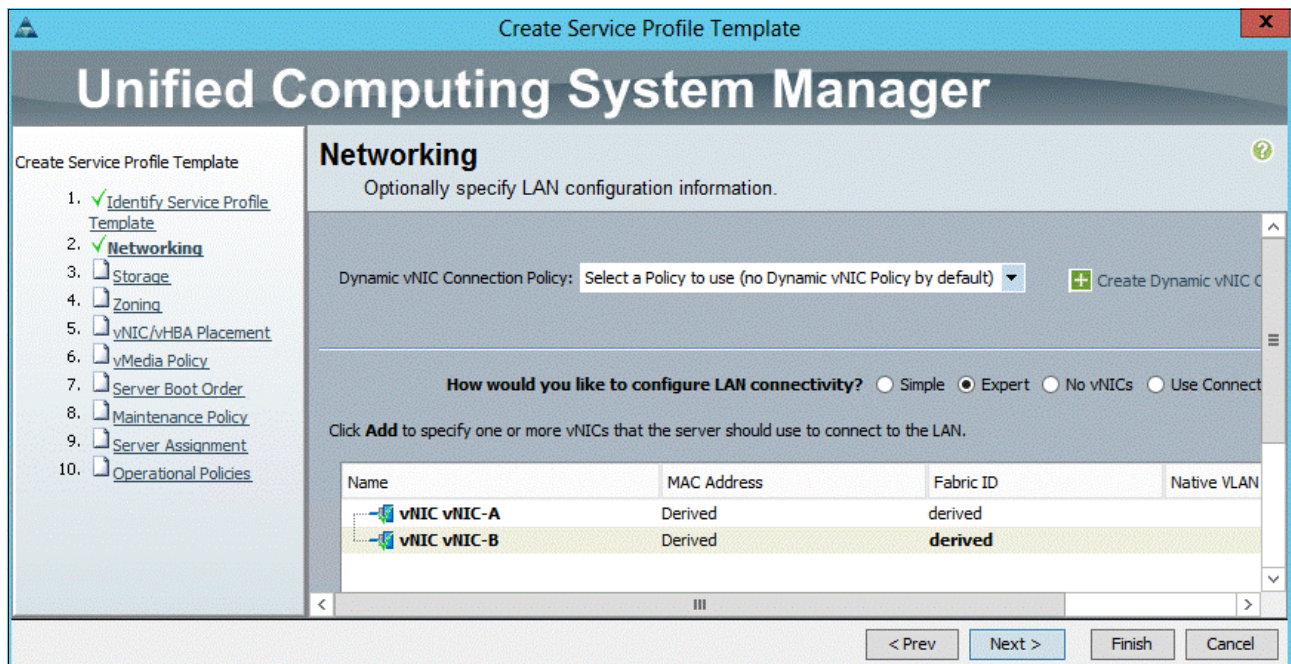


Figure 8-94 Both vNICs created

- o. Review the table in the Networking window to make sure that both vNICs were created.
- p. Click **Next**. The window that is shown in Figure 8-95 opens.

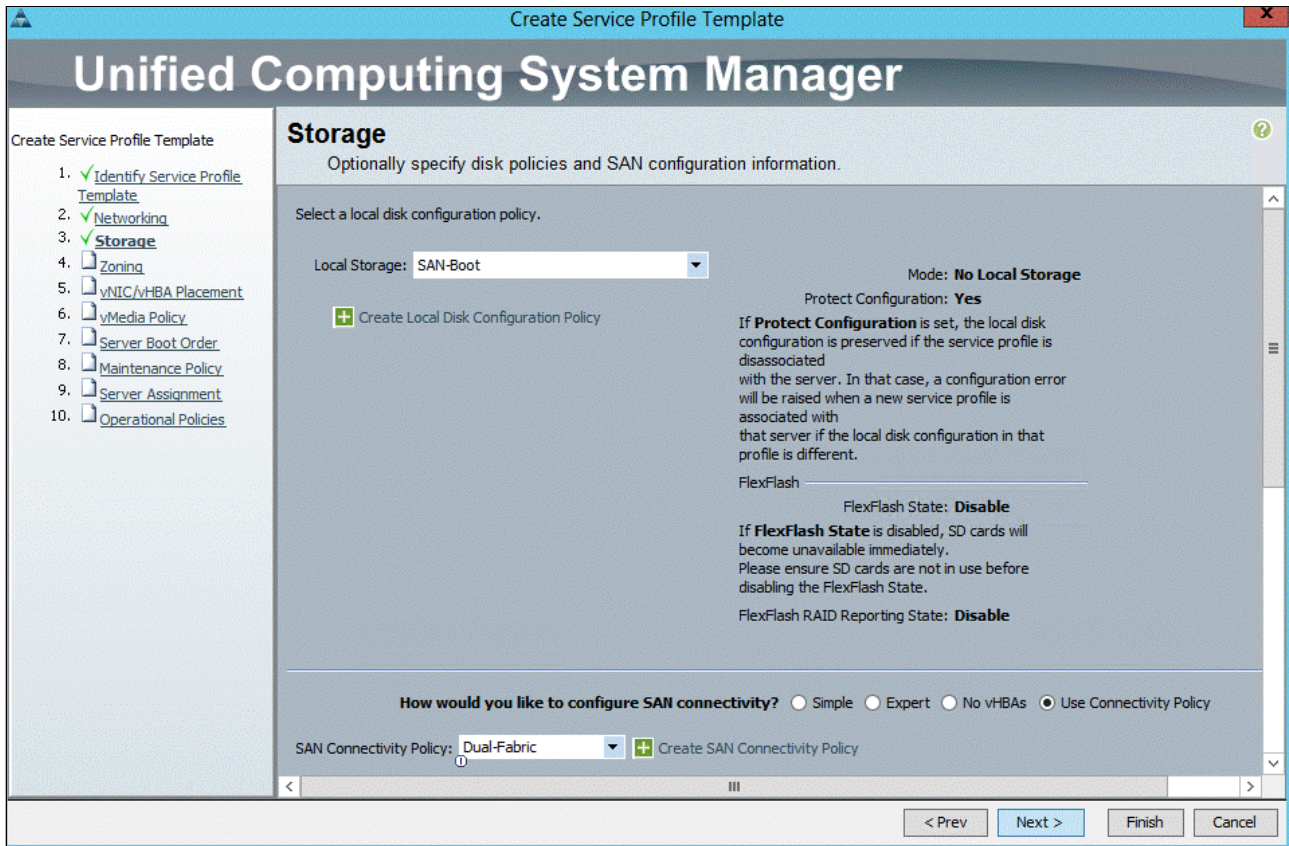


Figure 8-95 Configure the Storage options

6. Configure the Storage options:
 - a. Choose a local disk configuration policy:
 - If the server in question has local disks, choose **Default** from the Local Storage list.
 - If the server in question does not have local disks, select **SAN-Boot**.
 - b. Select the **Use Connectivity Policy** radio button to configure the SAN connectivity.
 - c. For the SAN connectivity Policy, select **Dual-Fabric**.
 - d. Click **Next**.
7. Accept the zoning options and click **Next**, as shown in Figure 8-96.

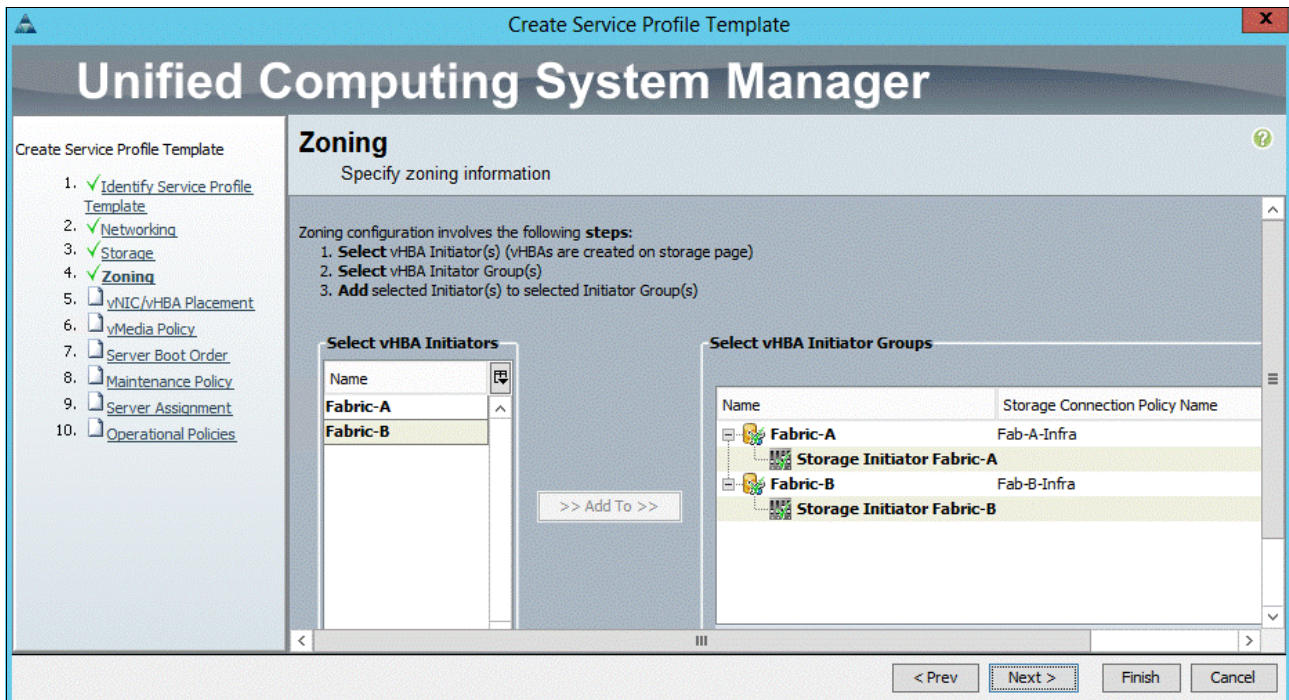


Figure 8-96 Zoning options

The window that is shown in Figure 8-97 opens.

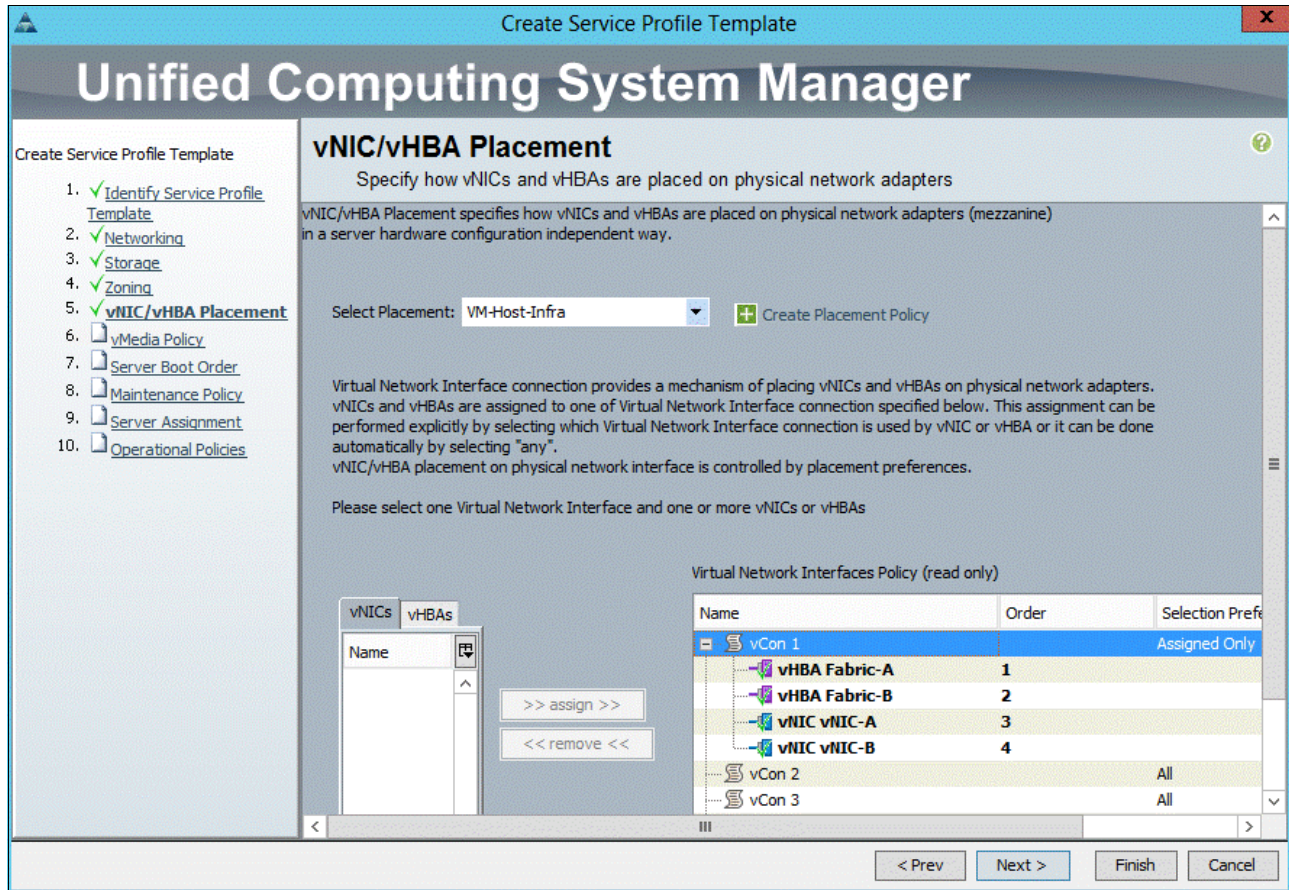


Figure 8-97 Set the vNIC/vHBA placement options

8. Set the vNIC/vHBA placement options:
 - a. In the Select Placement list, choose the VM-Host-Infra placement policy.
 - b. Select **vCon1** and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:
 - i. vHBA Fabric-A
 - ii. vHBA Fabric-B
 - iii. vNIC-A
 - iv. vNIC-B
 - c. Review the table to verify that all vNICs and vHBAs were assigned to the policy in the appropriate order.
 - d. Click **Next**.

9. Click **Next** to bypass the vMedia policy window. The window that is shown in Figure 8-98 opens.

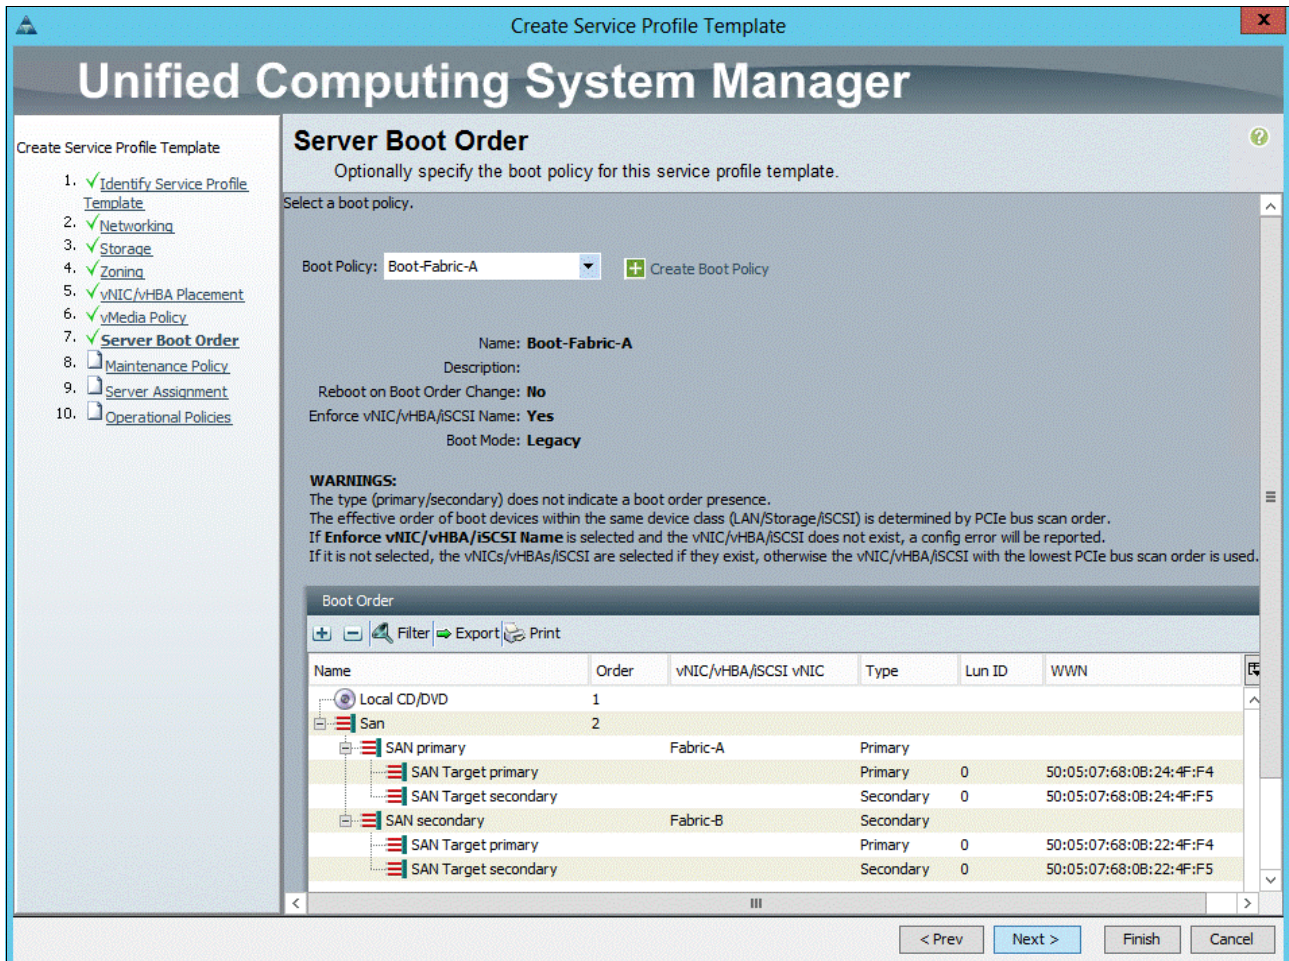


Figure 8-98 Set the server boot order

10. Set the Server Boot Order:
 - a. In the Boot Policy list, select **Boot-Fabric-A**.
 - b. Review the table to verify that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
 - c. Click **Next**.

The window that is shown in Figure 8-99 opens.

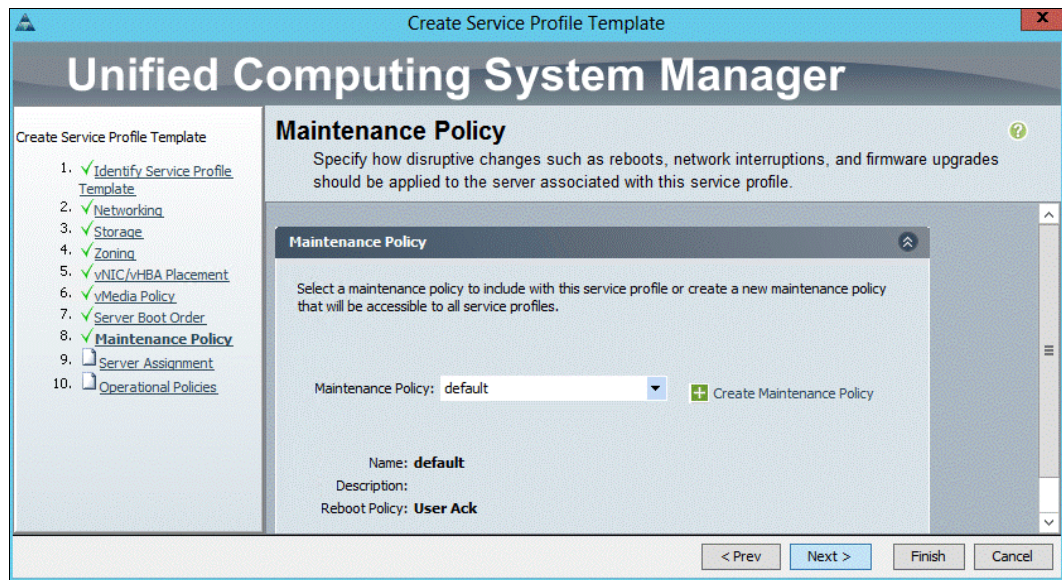


Figure 8-99 Add a maintenance policy

11. Add a Maintenance Policy:

- a. Select the **default** maintenance policy.
- b. Click **Next**. The window that is shown in Figure 8-100 opens.

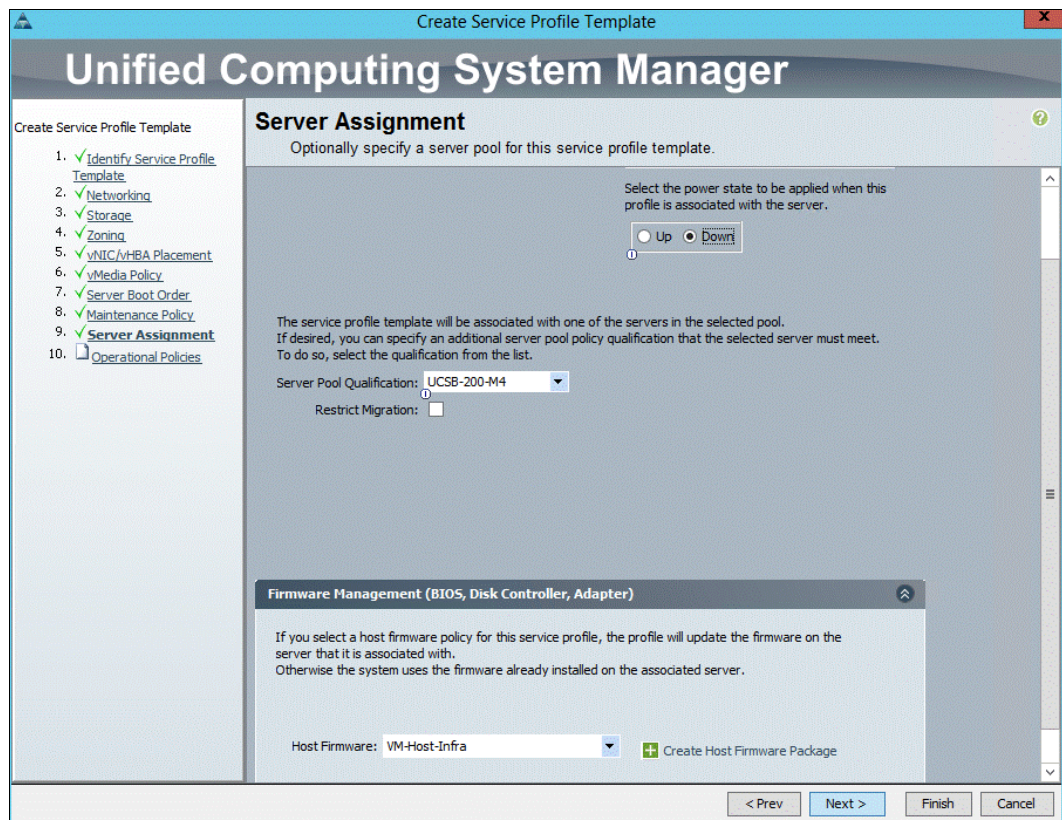


Figure 8-100 Specify the server assignment

12. Specify the Server Assignment:

- a. In the Pool Assignment list, select **Infra_Pool**.
- b. (Optional) Choose a Server Pool Qualification policy.
- c. Select **Down** as the power state to be applied when the profile is associated with the server.
- d. Expand **Firmware Management** at the bottom of the window and select **VM-Host-Infra** from the Host Firmware list.
- e. Click **Next**. The window that is shown in Figure 8-101 opens.

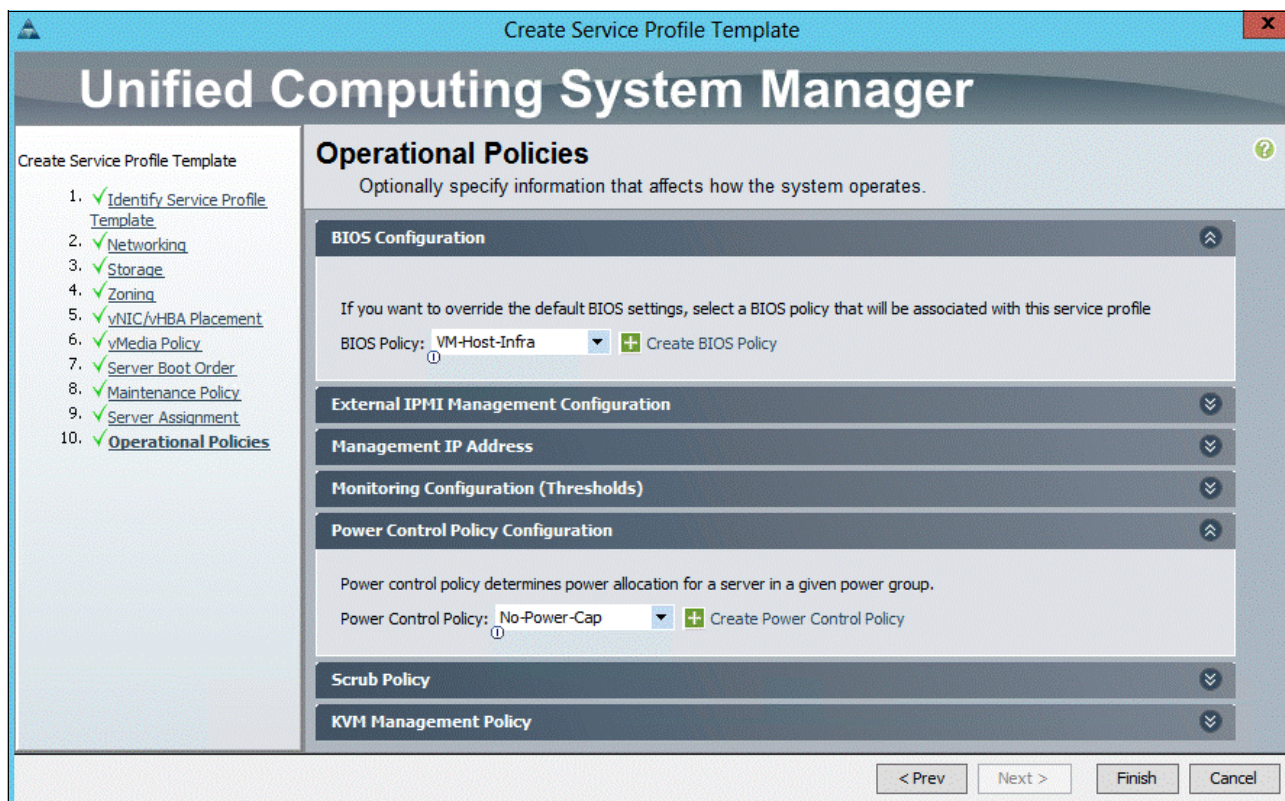


Figure 8-101 Add operational policies

13. Add Operational Policies:

- a. In the BIOS Policy list, select **VM-Host-Infra**.
- b. Expand **Power Control Policy Configuration** and choose **No-Power-Cap** in the Power Control Policy list.

14. Click **Finish** to create the service profile template.

15. Click **OK** in the confirmation message.

16. Click the **Servers** tab in the navigation pane.

17. Click **Service Profile Templates** → **root**.

18. Right-click the previously created VM-Host-Infra-Fabric-A template and select **Create a Clone**. The window that is shown in Figure 8-102 opens.




Figure 8-102 Create a clone

19. In the dialog box, enter VM-Host-Infra-Fabric-B as the name of the clone, select **root** for the Org field, and click **OK**.

20. Click **OK**.

21. Choose the newly cloned service profile template and click the **Boot Order** tab, as shown in Figure 8-103.

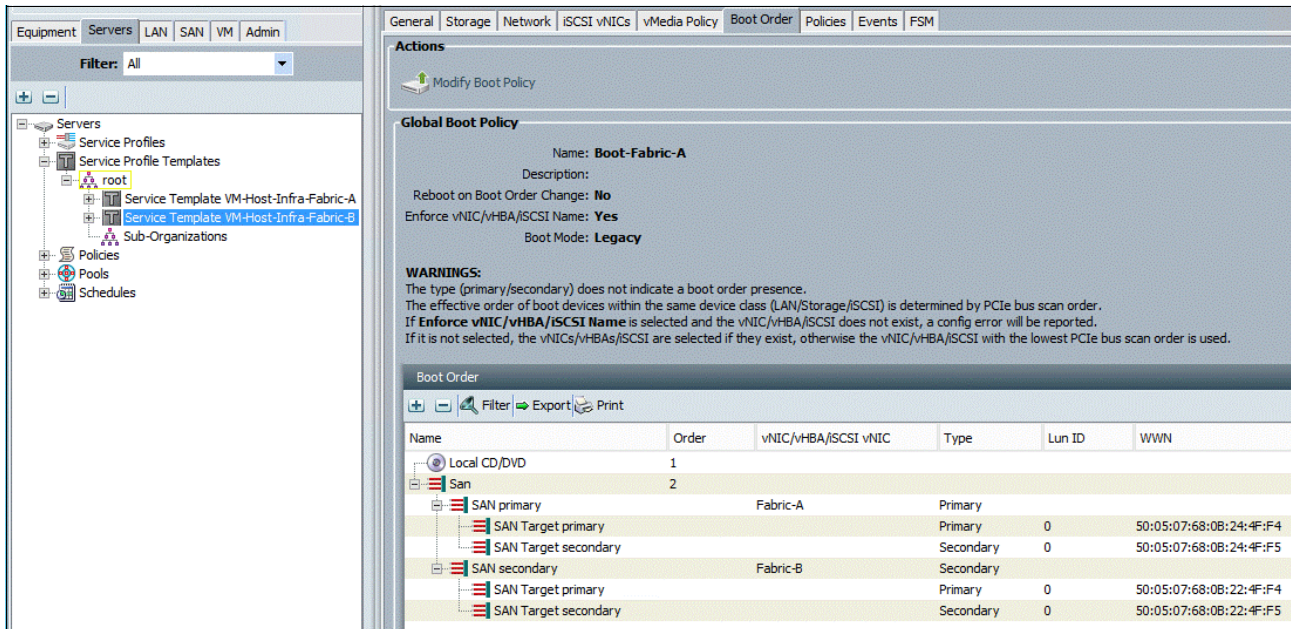


Figure 8-103 Click the Boot Order tab

22. Click **Modify Boot Policy**. The window that is shown in Figure 8-104 on page 133 opens.

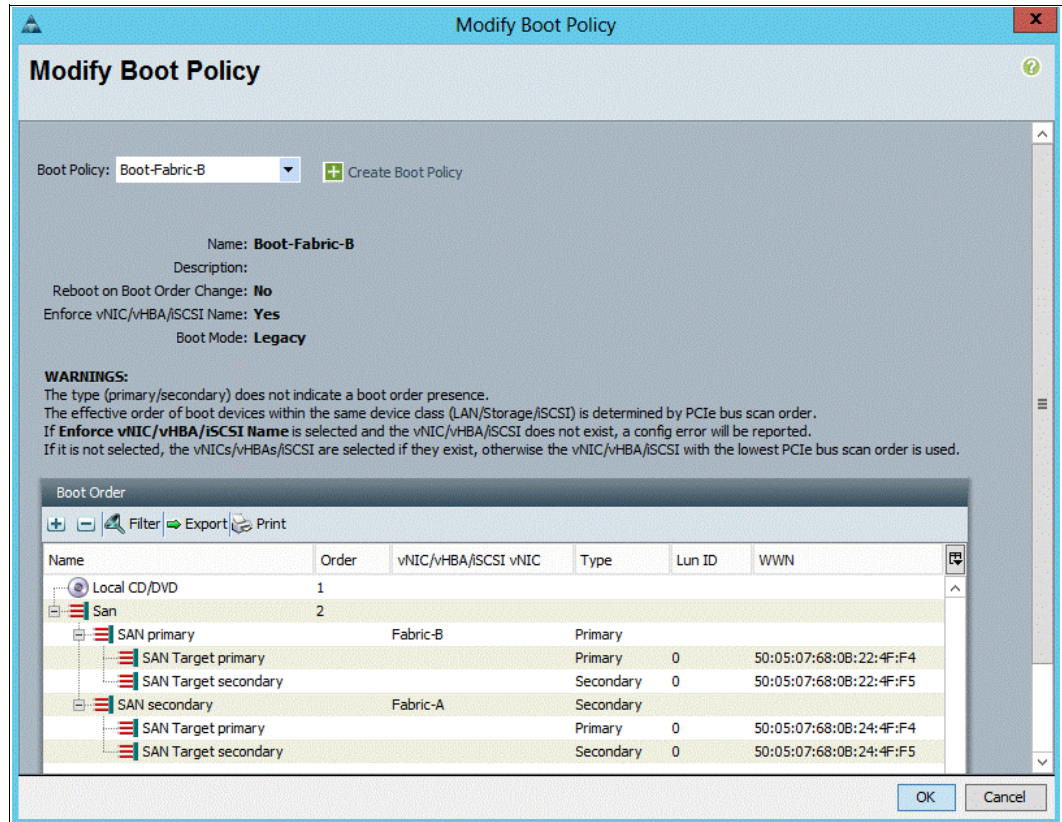


Figure 8-104 Modify Boot Policy

23. In the Boot Policy list, select **Boot-Fabric-B**.
24. Click **OK**, and then click **OK** again.

25. In the right pane, click the **Network** tab and then click **Modify vNIC/HBA Placement**. The window that is shown in Figure 8-105 opens.

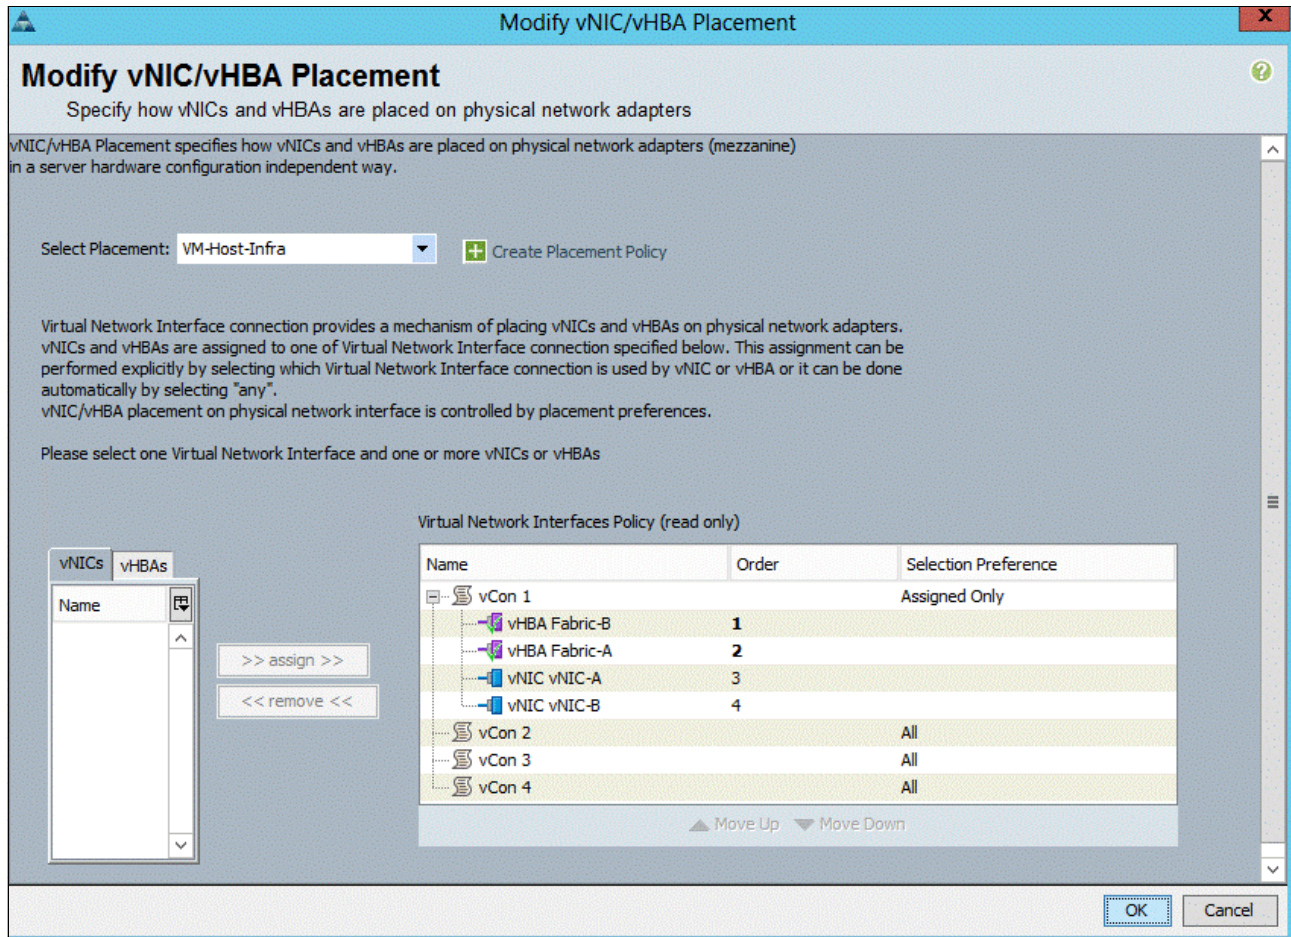


Figure 8-105 Modify vNIC/HBA Placement

26. Select **VM-Host-Infra**, expand **vCon 1**, and move vHBA Fabric-B ahead of vHBA Fabric-A in the placement order.

27. Click **OK**, and then click **OK** again.

8.2.34 Creating service profiles

To create service profiles from the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Click **Service Profile Templates** → **root** → **Service Template VM-Host-Infra-Fabric-A**.
3. Right-click **VM-Host-Infra-Fabric-A** and select **Create Service Profiles from Template**, as shown in Figure 8-106 on page 135. The window that is shown in Figure 8-107 on page 135 opens.

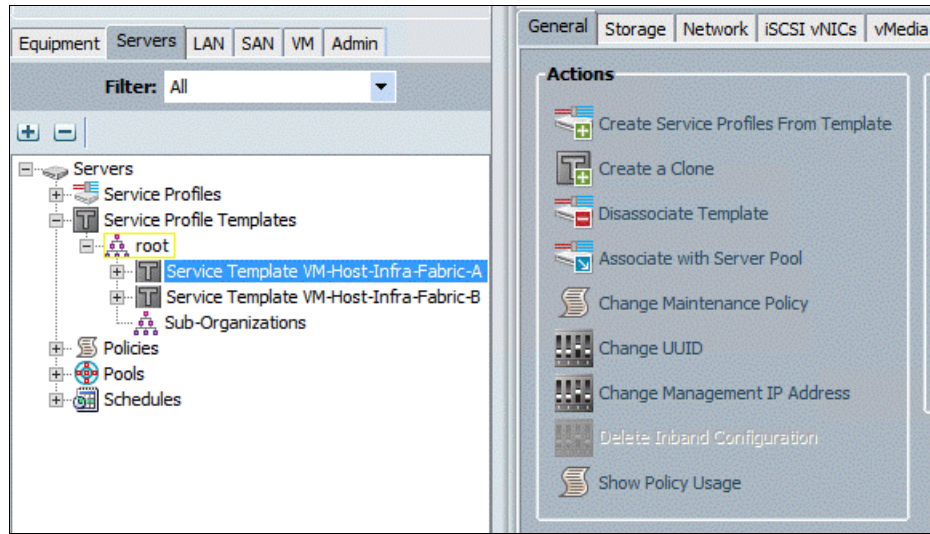


Figure 8-106 Create Service Profiles from Template

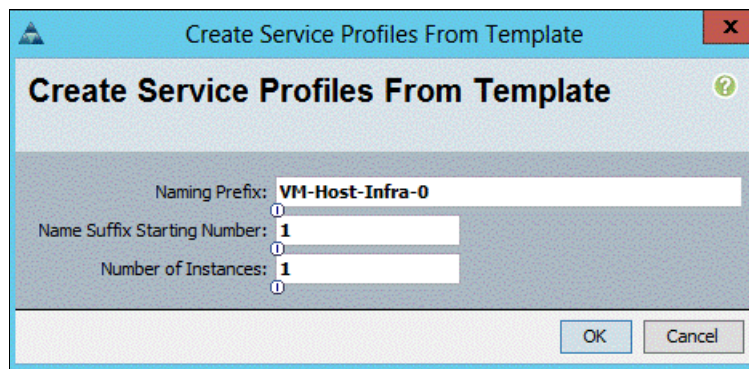


Figure 8-107 Create Service Profiles from Template

4. Enter VM-Host-Infra-0 as the Naming Prefix.
5. Enter 1 as the Name Suffix Starting Number.
6. Enter 1 as the Number of Instances.
7. Click **OK** to create the service profile.
8. Click **OK** in the confirmation message..
9. Click **Service Profile Templates** → **root** → **Service Template VM-Host-Infra-Fabric-B**.

- Right-click **VM-Host-Infra-Fabric-B** and select **Create Service Profiles from Template**. The window that is shown in Figure 8-108 opens.

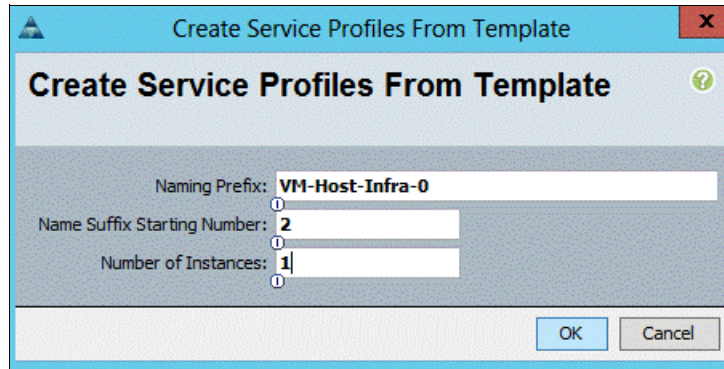


Figure 8-108 Create Service Profiles from Template

- Enter VM-Host-Infra-0 as the service profile prefix.
- Enter 2 as the Name Suffix Starting Number.
- Enter 1 as the Number of Instances.
- Click **OK** to create the service profile.
- Click **OK** in the confirmation message.
- Verify that the service profiles VM-Host-Infra-01 and VM-Host-Infra-02 are created, as shown in Figure 8-109. The service profiles are automatically associated with the servers in their assigned server pools.

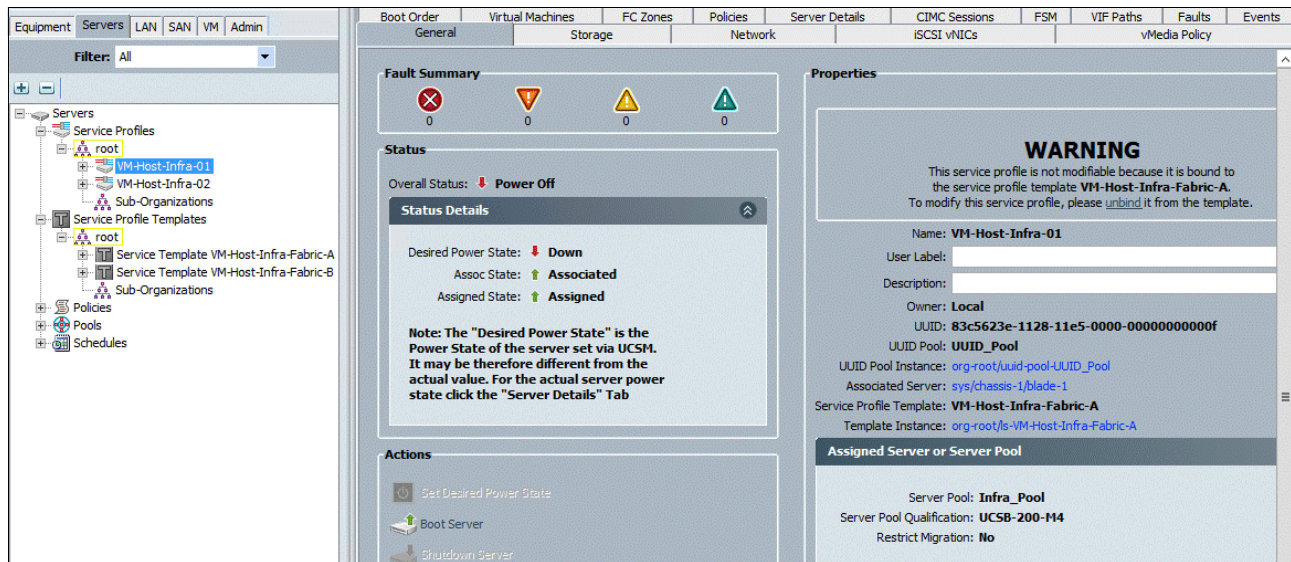


Figure 8-109 Verify that the service profiles are created

- Verify that the FC zones are created after the service profile assignment, as shown in Figure 8-110 on page 137.

Name	Initiator WWPN	Target WWPN	Initia...	Admi...	Ope...	Fabr...	VSA...	Zone...
ucs_Versastack-FI_A_1_VM-Host-Infra-01_Fabric-A	20:00:00:25:B5:01:0A:0F		Fabric-A	Applied	Active	A	101	1
FC Target 50:05:07:68:0B:24:4F:F4		50:05:07:68:0B:24:4F:F4						
ucs_Versastack-FI_A_2_VM-Host-Infra-01_Fabric-A	20:00:00:25:B5:01:0A:0F		Fabric-A	Applied	Active	A	101	2
FC Target 50:05:07:68:0B:24:4F:F5		50:05:07:68:0B:24:4F:F5						
ucs_Versastack-FI_A_3_VM-Host-Infra-01_Fabric-A	20:00:00:25:B5:01:0A:0F		Fabric-A	Applied	Active	A	101	3
FC Target 50:05:07:68:0B:23:4F:F4		50:05:07:68:0B:23:4F:F4						
FC Target 50:05:07:68:0B:23:4F:F5		50:05:07:68:0B:23:4F:F5						
FC Target 50:05:07:68:0B:24:4F:F4		50:05:07:68:0B:24:4F:F4						
FC Target 50:05:07:68:0B:24:4F:F5		50:05:07:68:0B:24:4F:F5						
ucs_Versastack-FI_B_1_VM-Host-Infra-01_Fabric-B	20:00:00:25:B5:01:0B:0F		Fabric-B	Applied	Active	B	102	1
FC Target 50:05:07:68:0B:21:4F:F5		50:05:07:68:0B:21:4F:F5						
FC Target 50:05:07:68:0B:22:4F:F4		50:05:07:68:0B:22:4F:F4						
FC Target 50:05:07:68:0B:22:4F:F5		50:05:07:68:0B:22:4F:F5						
FC Target 50:05:07:68:0B:23:4F:F4		50:05:07:68:0B:23:4F:F4						
ucs_Versastack-FI_B_2_VM-Host-Infra-01_Fabric-B	20:00:00:25:B5:01:0B:0F		Fabric-B	Applied	Active	B	102	2
FC Target 50:05:07:68:0B:22:4F:F4		50:05:07:68:0B:22:4F:F4						
ucs_Versastack-FI_B_3_VM-Host-Infra-01_Fabric-B	20:00:00:25:B5:01:0B:0F		Fabric-B	Applied	Active	B	102	3
FC Target 50:05:07:68:0B:22:4F:F5		50:05:07:68:0B:22:4F:F5						

Figure 8-110 Verify FC zones

18. After completing all the previous steps, power on the servers and you should see the SAN-Boot LUNs during BIOS POST, as shown in Figure 8-111.

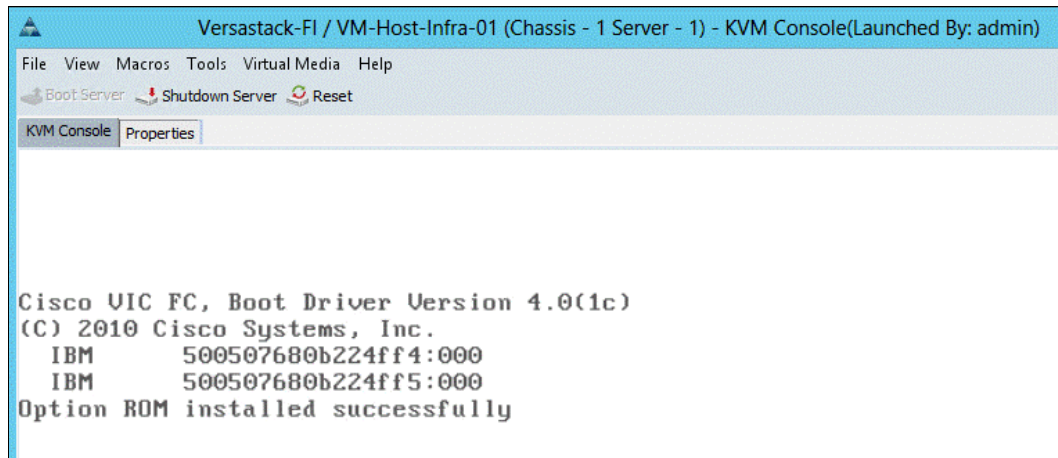


Figure 8-111 SAN-Boot LUNs during BIOS POST

8.3 Backing up the Cisco UCS Manager configuration

Back up your Cisco UCS Manager configuration. For more information about this topic, see the *Cisco UCS Manager GUI Configuration Guide, Release 2.2*, found at:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/2-2/b_UCSM_GUI_Configuration_Guide_2_2/b_UCSM_GUI_Configuration_Guide_2_2_chapter_010_1010.html



SAN boot

This chapter describes how to add the host mappings for the host profiles that are created through Cisco UCS Manager to the Storwize V7000 storage system, connect to the boot LUNs, and perform the initial ESXi installation. The WWPNs for the hosts are to complete the steps in this chapter.

9.1 Adding hosts and mapping the boot volumes on the Storwize V7000 system

To add hosts and map the boot volumes on the Storwize V7000 storage system, complete the following steps:

1. Open the Storwize V7000 management GUI by go to <<var_cluster_mgmt_ip>> and log in with your superuser or admin account.
2. In the left pane, click the Host icon, which is the fourth icon down, and click **Hosts**.
3. Click **Create Host** in the upper left menu to start the Create Host wizard.

Figure 9-1 shows the Add Host window, which shows options for FC and iSCSI hosts.

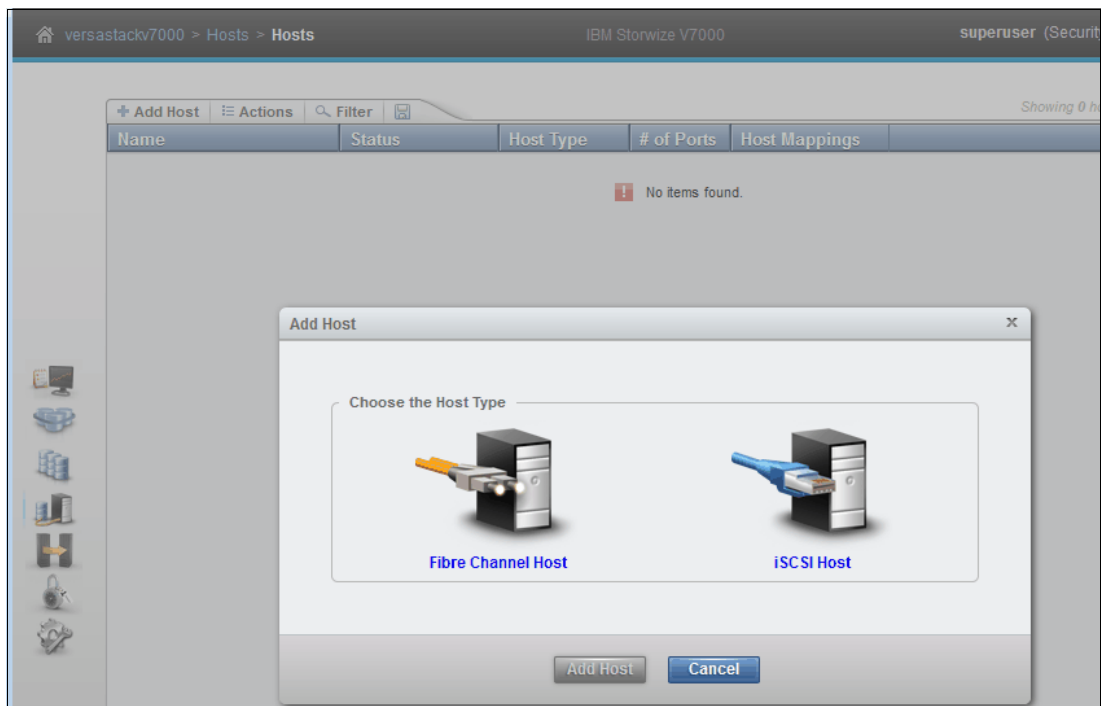


Figure 9-1 Add a host on the Storwize V7000 storage system

4. Select the **Fibre Channel Host** option.
5. For Host Name, enter vm-host-infra-01.
6. For Fibre Channel Ports, click the drop-down menu and select or input the WWPNs for the A path vHBAs (<<var_wwpn_vm-host-infra-01-a>>) and click **Add Port to List**.
7. Click the drop-down menu again, and select or input the host B port (<<wwpn_vm-host-infra-01-b>>) and click **Add Port to List**.
8. Leave Advanced Settings as the default and click **Create Host**.
9. Click **Close**.

Note: If the hosts are powered on and zoned correctly, they appear in the selection dropdown or, if you type in the WWPN, you should see green check marks for each WWPN.

10. Click **Create Host** to create the second host.

11. Select the **Fibre Channel Host** option.
12. For Host Name, enter vm-host-infra-02.
13. For Fibre Channel Ports, select the drop-down menu and select the WWPNs for the A path vHBAs (<<var_wwpn_vm-host-infra-02-a>>) and click **Add Port to List**.
14. Select the B port by selecting the variable for the B path (<<wwpn_vm-host-infra-02-b>>) and click **Add Port To List**.
15. Leave the Advanced Settings as the default and click **Create Host**.
16. Click **Close**.

Figure 9-2 shows creating the host vm-host-infra-02. The FC ports appear in the drop-down menu.

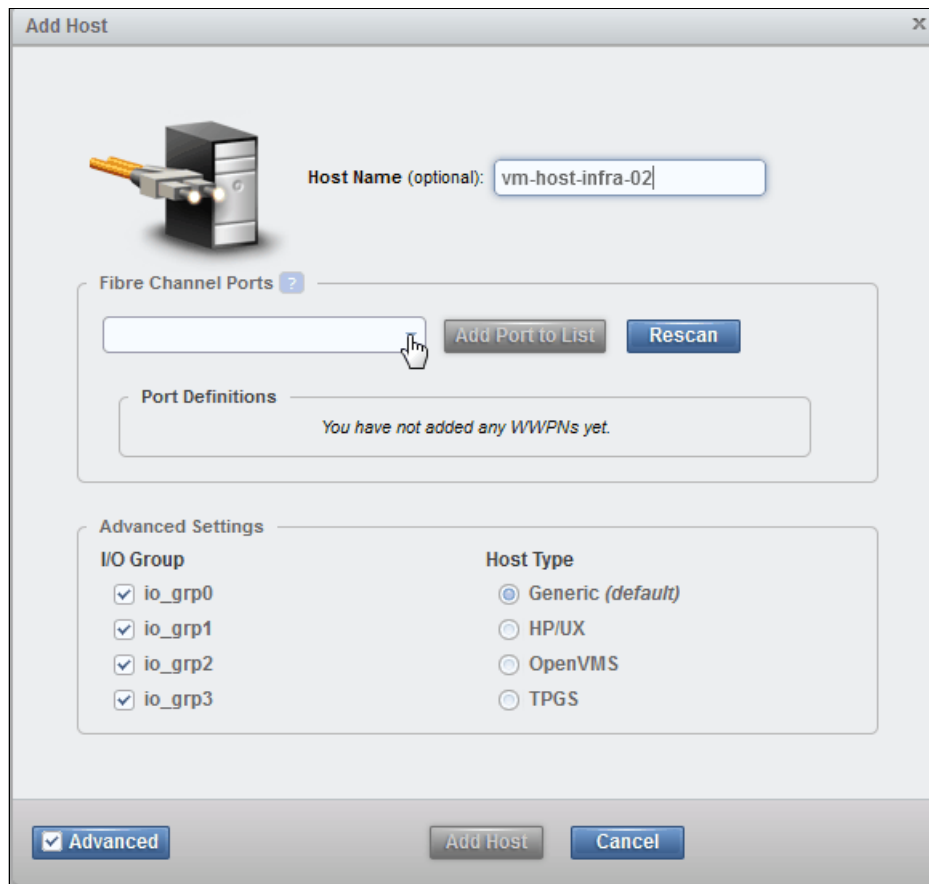


Figure 9-2 Create vm-host-infra-02

17. Click the Volumes icon in the left pane, then click the volumes menu item to display the created volumes.

18. Right-click the volume **vm_host_boot_1** and select **Map to Host**.

Figure 9-3 shows mapping the first boot LUN to the first host.

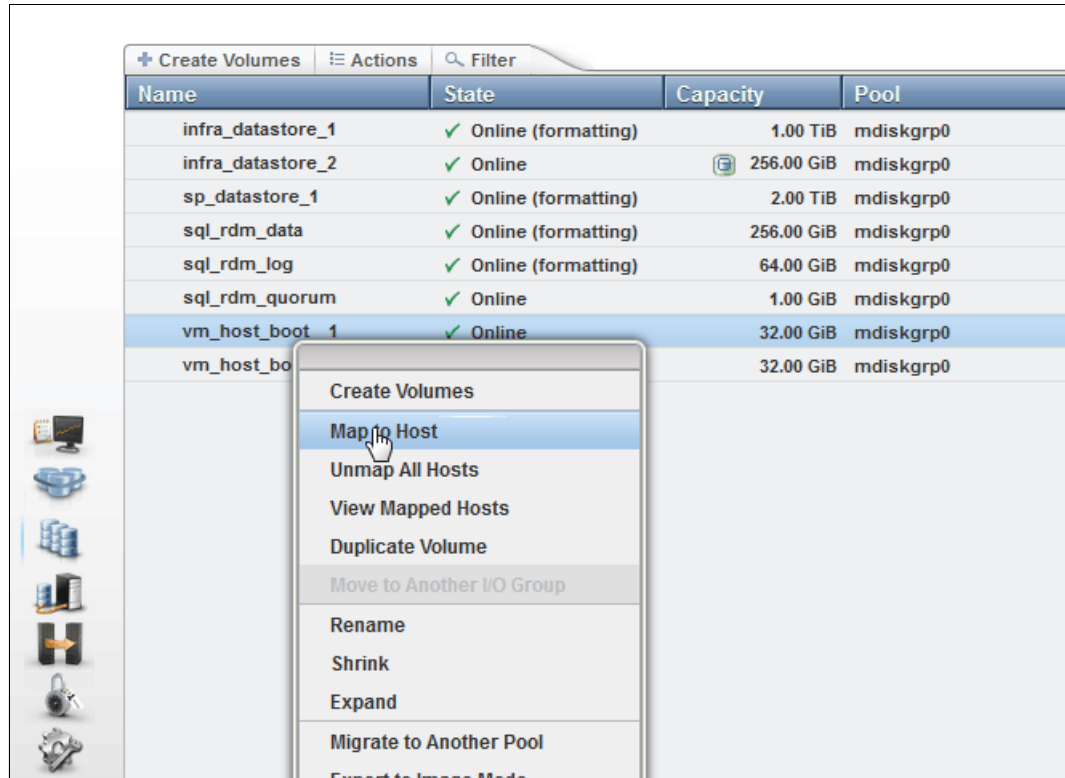


Figure 9-3 Showing mapping a volume to host

19. Right-click **vm-host-infra-01** and click **Map to Host**. Then, in the drop-down menu, select **Map Volumes** and then click **Close**.

20. Right-click **vm_host_boot_02** and click **Map to Host**. Then, in the drop-down menu, select **Map Volumes** and then click **Close**.

21. Power on the servers and verify that the boot LUNs appear during the BIOS POST, as shown in Figure 9-4.

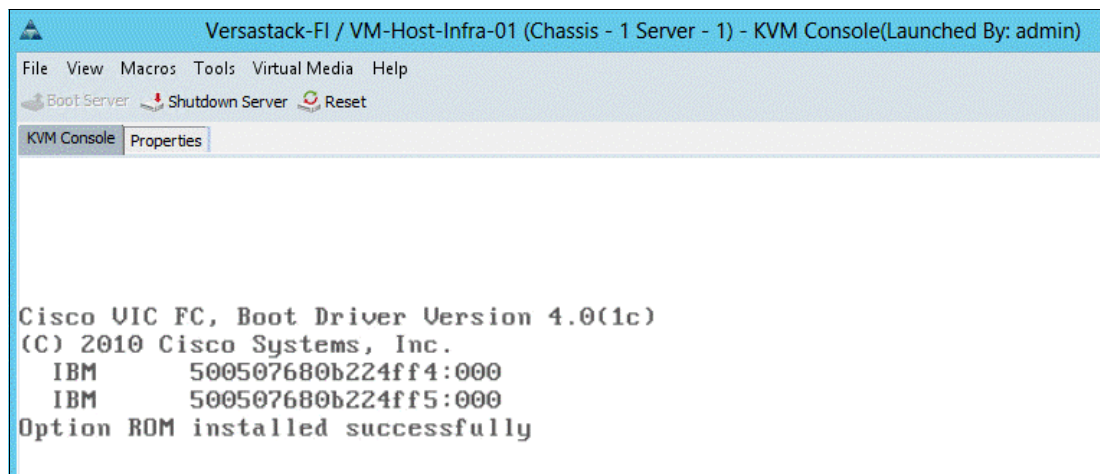


Figure 9-4 Boot LUNs

Note: In this VersaStack environment, there are two paths to the boot LUN, so it appears twice.



VersaStack VMware ESXi 5.5 Update 2 SAN boot installation

This chapter provides detailed instructions for installing VMware ESXi 5.5 Update 2 in a VersaStack environment. After the procedures are completed, two SAN-booted ESXi hosts are provisioned. These deployment procedures are customized to include the environment variables.

Note: Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in Keyboard, Video, Mouse (KVM) console and virtual media features in Cisco Unified Computing System (Cisco UCS) Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs). This method uses the Cisco Custom ESXi 5.5.0 U2 GA ISO file, which is downloaded from the following URL. This file is required for this procedure because it contains custom Cisco drivers, which reduce the number of installation steps.

<https://my.vmware.com/web/vmware/details?downloadGroup=OEM-ESXI55U2-CISCO&productId=353>

10.1 The Cisco UCS 6200 Fabric Interconnect Cisco UCS Manager

The administrator can use KVM to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Download the Cisco Custom ISO for ESXi from the VMware website.
2. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step starts the Cisco UCS Manager application.

Note: You need Java Runtime Environment 1.6 or higher to run this application.

Figure 10-1 shows the Cisco Unified Computing System Manager start window, which has options to start Cisco UCS Manager and the KVM manager.



Figure 10-1 Cisco Unified Computing System Manager start window

3. Log in to Cisco UCS Manager by using the admin user name and password.
4. From the main menu, click the **Servers** tab.
5. Click **Servers** → **Service Profiles** → **root** → **vm-host-infra-01**.
6. Right-click **vm-host-infra-01** and select **KVM Console**.
7. Click **Servers** → **Service Profiles** → **root** → **vm-host-infra-02**.
8. Right-click **vm-host-infra-02** and select **KVM Console Actions** → **KVM Console**.

Figure 10-2 on page 147 shows using UCS manager to start KVM on vm-host-infra-01.

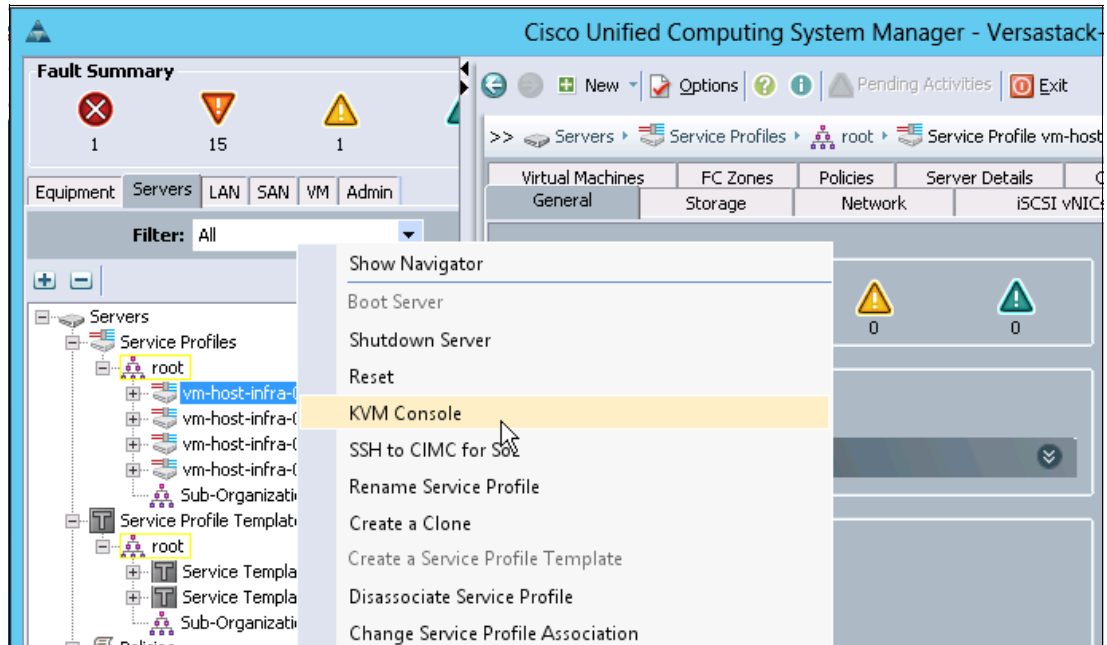


Figure 10-2 Start KVM in Cisco UCS

10.2 Setting up a VMware ESXi installation

This section describes how to complete the VMware ESXi installation.

10.2.1 ESXi hosts vm-host-infra-01 and vm-host-infra-02

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click the **Virtual Media** tab.

Figure 10-3 shows the location of the Activate Virtual Devices option in The Virtual Media menu.




Figure 10-3 Virtual Media menu

2. Click **Activate Virtual Devices**, select **Accept this Session**, and then click **Apply**.
3. Click **Virtual Media** → **Map CD/DVD**, then browse to the ESXi installer ISO image file and click **Open**.
4. Click **Map Device** to map the newly added image.

Figure 10-4 shows mapping the ESXi 5.5.0 u2 custom ISO that was downloaded from the VMWare website.

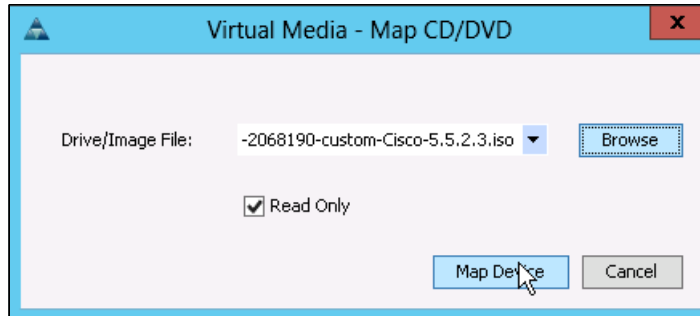


Figure 10-4 Map the ESXi 5.5.0 u2 custom ISO

5. Click the **KVM** tab to monitor the server boot.
6. If the server is powered on, first shut down the server, then start the server by clicking **Boot Server** and clicking **OK**, and then click **OK** again.

10.3 Installing ESXi

This section describes how to install ESXi.

10.3.1 ESXi hosts vm-host-infra-01 and vm-host-infra-02

To install VMware ESXi on to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On start, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the menu that opens.

Note: You might have to press F6 and force the host to boot from the vDVD.

Figure 10-5 on page 149 shows the ESXi Boot device list that is accessed by pressing F6 repeatedly.

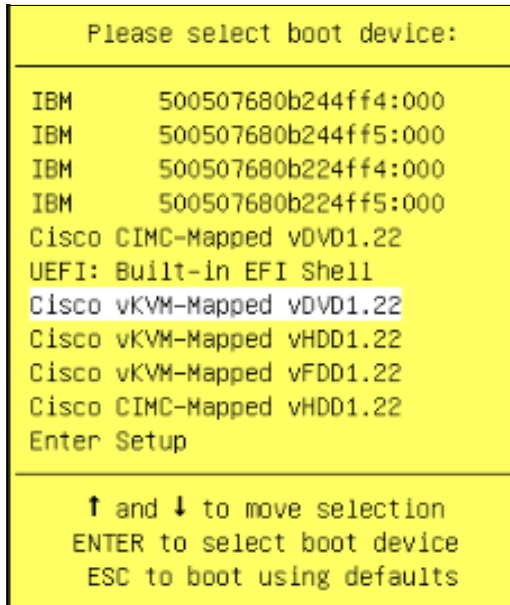


Figure 10-5 Boot device list

2. After the installer finishes loading, press Enter to continue with the installation.
3. Read and accept the user license agreement (EULA). Press F11 to accept and continue.
4. Select the IBM LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

Figure 10-6 shows the available local and remote disks.

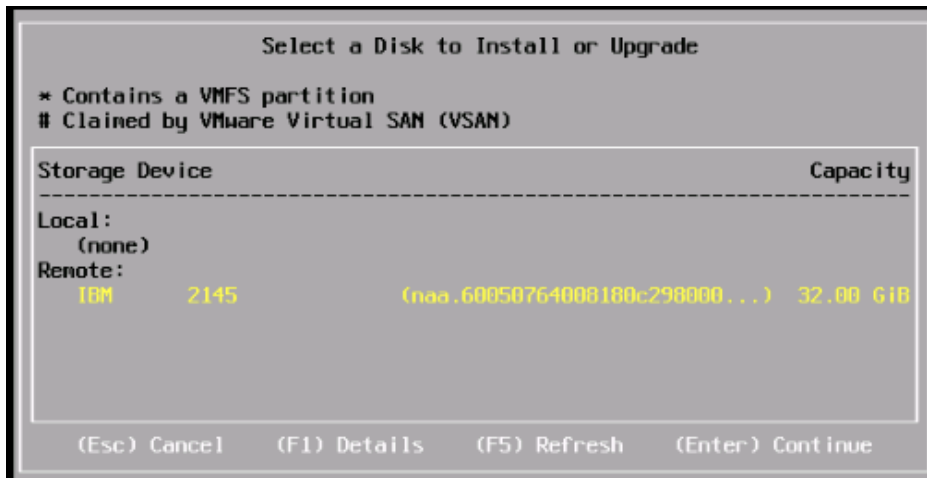


Figure 10-6 The boot LUN that is mapped to vm-host-infra-01

5. Select the appropriate keyboard layout and press Enter.
6. Enter and confirm the root password and press Enter.
7. The installer issues a warning that existing partitions will be removed from the volume. Press F11 to continue with the installation.
8. After the installation is complete, click the check icon to clear the Mapped ISO (in the Virtual Media tab of the KVM console) to unmap the ESXi installation image.

9. The Virtual Media window might issue a warning stating that it is preferable to eject the media from the guest. Because the media cannot be ejected and it is read-only, simply click **Yes** to unmap the image.
10. From the KVM tab, press Enter to restart the server.

10.4 Setting up management networking for ESXi hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host.

10.4.1 ESXi Host vm-host-infra-01

To configure the vm-host-infra-01 ESXi host with access to the management network, complete the following steps:

1. After the server finishes restarting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Click the **Configure the Management Network** option and press Enter.
4. Click the **VLAN (Optional)** option and press Enter.
5. Enter the `<<var_ib-mgmt_vlan_id>>` and press Enter.
6. From the Configure Management Network menu, click **IP Configuration** and press Enter.
7. Select the **Set Static IP Address and Network Configuration** option by using the Spacebar.
8. Enter the IP address for managing the first ESXi host: `<<var_vm_host_infra_01_ip>>`.
9. Enter the subnet mask for the first ESXi host.
10. Enter the default gateway for the first ESXi host.

Figure 10-7 on page 151 shows setting the IP address, subnet mask, and default gateway from the ESXi host.



Figure 10-7 IP configuration on ESXi 5.5.0

11. Press Enter to accept the changes to the IP configuration.
12. Click the **IPv6 Configuration** option and press Enter.
13. Using the spacebar, clear **Enable IPv6** (restart required) and press Enter.
14. Click the **DNS Configuration** option and press Enter.

Note: Because the IP address is assigned manually, the DNS information must also be entered manually.

15. Enter the IP address of the primary DNS server.
16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the fully qualified domain name (FQDN) for the first ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host restarts. After restart, press F2 and log back in as root.
22. Click **Test Management Network** to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

10.4.2 ESXi Host vm-host-infra-02

To configure the vm-host-infra-02 ESXi host with access to the management network, complete the following steps:

1. After the server finishes restarting, press F2 to customize the system.
2. Log in as root and enter the corresponding password.
3. Click the **Configure the Management Network** option and press Enter.
4. Click the **VLAN (Optional)** option and press Enter.
5. Enter the `<<var_ib-mgmt_vlan_id>>` and press Enter.
6. From the Configure Management Network menu, select **IP Configuration** and press Enter.
7. Select the **Set Static IP Address and Network Configuration** option by using the Spacebar.
8. Enter the IP address for managing the second ESXi host: `<<var_vm_host_infra_02_ip>>`.
9. Enter the subnet mask for the second ESXi host.
10. Enter the default gateway for the second ESXi host.
11. Press Enter to accept the changes to the IP configuration.
12. Click the **IPv6 Configuration** option and press Enter.
13. Using the spacebar, clear **Enable IPv6** (restart required) and press Enter.
14. Click the **DNS Configuration** option and press Enter.

Note: Because the IP address is assigned manually, the DNS information must also be entered manually

15. Enter the IP address of the primary DNS server.
16. Optional: Enter the IP address of the secondary DNS server.
17. Enter the FQDN for the second ESXi host.
18. Press Enter to accept the changes to the DNS configuration.
19. Press Esc to exit the Configure Management Network submenu.
20. Press Y to confirm the changes and return to the main menu.
21. The ESXi host restarts. After the restart completes, press F2 and log back in as root.
22. Click **Test Management Network** to verify that the management network is set up correctly and press Enter.
23. Press Enter to run the test.
24. Press Enter to exit the window.
25. Press Esc to log out of the VMware console.

10.5 vSphere setup

In this section, you set up the vSphere environment by using Windows Server 2012 and an SQL Server. The virtual machines that are used in this procedure are installed on a local datastore on VersaStack for any greenfield deployments; however, these VMs can be installed on a different ESX clustered system or physical hardware if you want. This procedure uses the volumes that were created for VMFS Datastores.

10.5.1 Downloading the VMware vSphere Client and vSphere Remote CLI

To download the VMware vSphere Client and install Remote CLI, complete the following steps:

1. Open a web browser on the management workstation and go to the VM-Host-Infra-01 management IP address.
2. Download and install both the vSphere Client and the Windows version of vSphere Remote Command-Line Interface.

Note: These applications are downloaded from the VMware website.

Logging in to VMware ESXi hosts by using the VMware vSphere Client

You must log in to both hosts.

ESXi Host vm-host-infra-01

To log in to the vm-host-infra-01 ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of vm-host-infra-01 as the host you are trying to connect to: <<var_vm_host_infra_01_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click **Login** to connect.

ESXi Host vm-host-infra-02

To log in to the vm-host-infra-02 ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of vm-host-infra-02 as the host you are trying to connect to: <<var_vm_host_infra_02_ip>>.
2. Enter root for the user name.
3. Enter the root password.
4. Click **Login** to connect.

10.6 Setting up VMkernel ports and the virtual switch

For each ESXi host, the steps in the section that follows must be repeated.

10.6.1 ESXi Host vm-host-infra-01

Repeat the steps in this section for all the ESXi hosts.

To set up the VMkernel ports and the virtual switches on the vm-host-infra-01 ESXi host, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the **Configuration** tab.
3. Click **Networking** in the Hardware pane.
4. Click **Properties** on the right side of vSwitch0.
5. Select the vSwitch configuration and click **Edit**.
6. From the General tab, change the MTU to 9000.
7. Click **OK** to close the properties for vSwitch0.
8. Select the Management Network configuration and click **Edit**.
9. Change the network label to VMkernel-MGMT and make sure that the **Management Traffic** check box is checked.
10. Click **OK** to finalize the edits for Management Network.
11. Select the VM Network configuration and click **Edit**.
12. Change the network label to VM-Production and enter `<<var_devmgmt_vlan_id>>` in the VLAN ID (Optional) field.
13. Click **OK** to finalize the edits for VM Network.
14. Click **Add** to add a network element.
15. Select **VMkernel** and click **Next**.
16. Change the network label to VMkernel-vMotion and enter `<<var_vmotion_vlan_id>>` in the VLAN ID (Optional) field.

Important: Whenever you define multiple networks across hosts, the syntax must be the same on those hosts.

17. Select the **Use this port group for vMotion** check box.
18. Click **Next** to continue with the vMotion VMkernel creation.
19. Enter the IP address `<<var_vmotion_vlan_id_ip_host-01>>` and the subnet mask `<<var_vmotion_vlan_id_mask_host-01>>` for the vMotion VLAN interface for VM-Host-Infra-01.
20. Click **Next** to continue with the vMotion VMkernel creation.
21. Click **Finish** to finalize the creation of the vMotion VMkernel interface.
22. Select the VMkernel-vMotion configuration and click **Edit**.
23. Change the MTU to 9000.
24. Click **OK** to finalize the edits for the VMkernel-vMotion network.

25. Click **Add** and select **Virtual Machine Network**, and then click **Next**.
26. Change the network label to VM-WinCSV and enter <<var_vmwincsv_vlan_id>> in the VLAN ID (Optional) field.
27. Click **Next**, and click **Finish** to complete the creation of the VM-WinCSV network.
28. Click **Add** and select **Virtual Machine Network**, and then click **Next**.
29. Change the network label to VM-WinClus and enter <<var_winclus_vlan_id>> in the VLAN ID (Optional) field.
30. Click **Next**, and click **Finish** to complete the creation of the VM-WinClus network.
31. Click **Add** and select **Virtual Machine Network**, and then click **Next**.
32. Change the network label to VM-Backup and enter <<var_vmbackup_vlan_id>> in the VLAN ID (Optional) field.
33. Click **Next**, and click **Finish** to complete the creation of the VM-Backup network.
34. Close the dialog box to finalize the ESXi host networking setup.

Figure 10-8 shows the vSwitch setup on vm-host-infra-01.

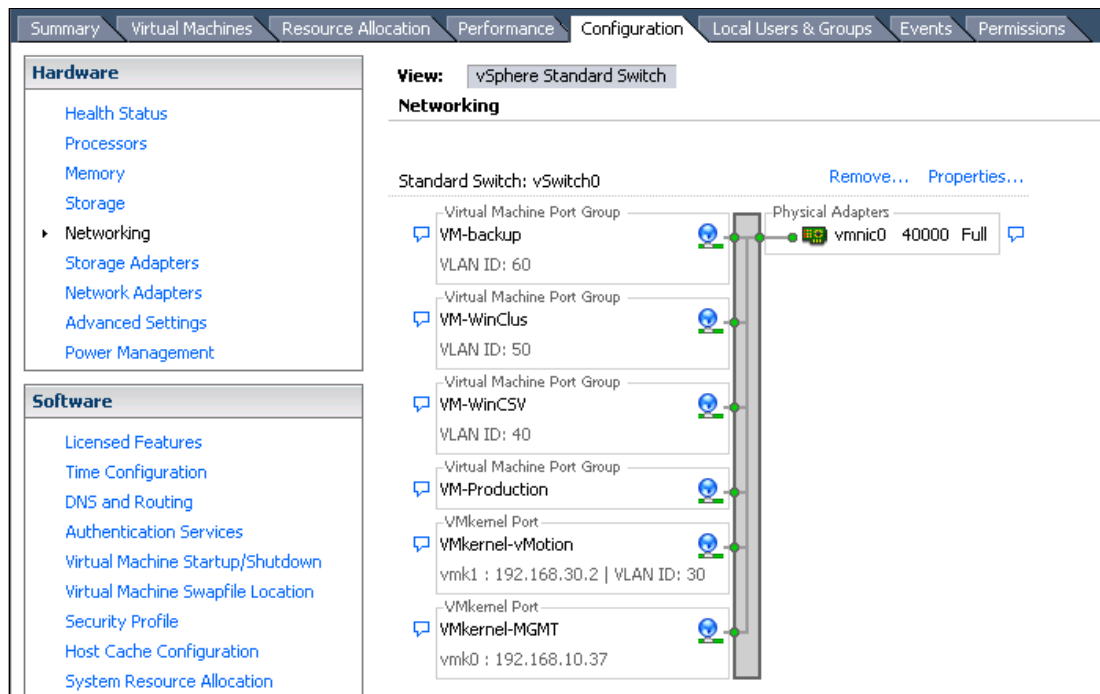


Figure 10-8 vSwitch setup

Figure 10-9 shows adding the second vmNIC on vm-host-infra-01.

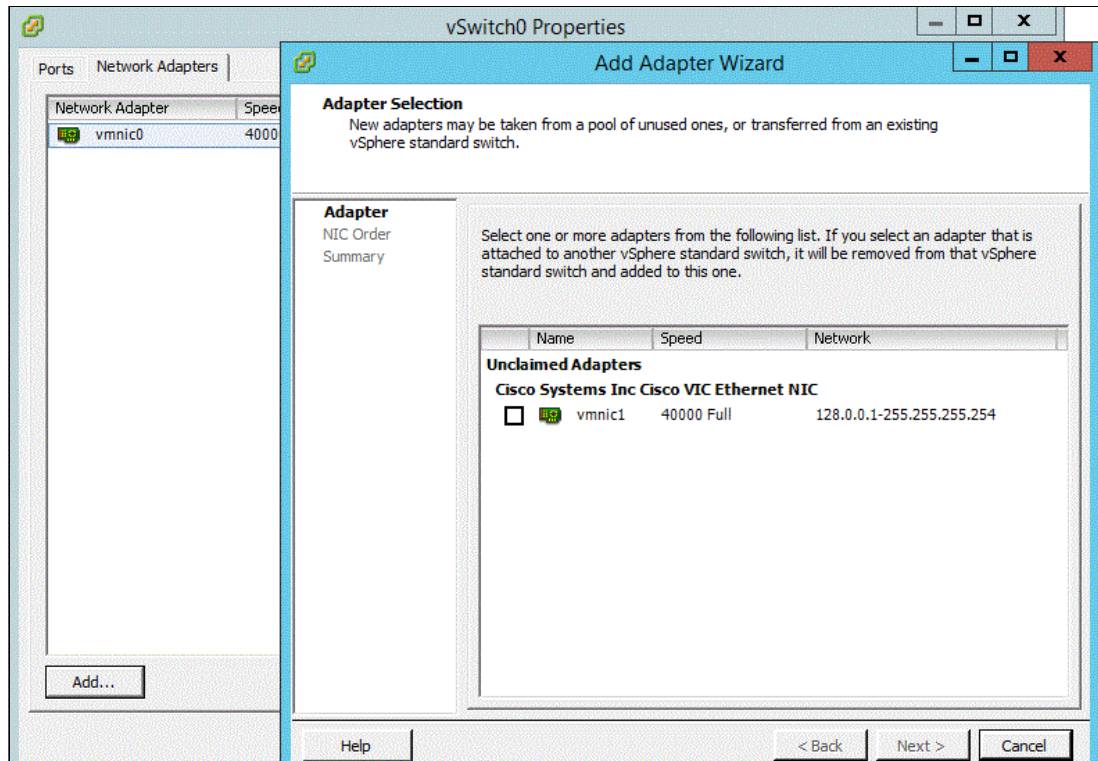


Figure 10-9 Add the second vmNIC

35. You must now assign another physical adapter to the switch to provide redundancy and load balancing features in this environment. To achieve this goal, use the NIC teaming feature that is available in vSwitch.
36. Click the properties of Vswitch0 on the Configuration Networking tab, click the **Network Adapters** tab, click **Add**, select **vmNIC1**, click **Next**, click **Next**, click **Finish**, and click **Close**.
37. Make sure that both vmNICs are in the active/active configuration, as shown in Figure 10-10 on page 157.

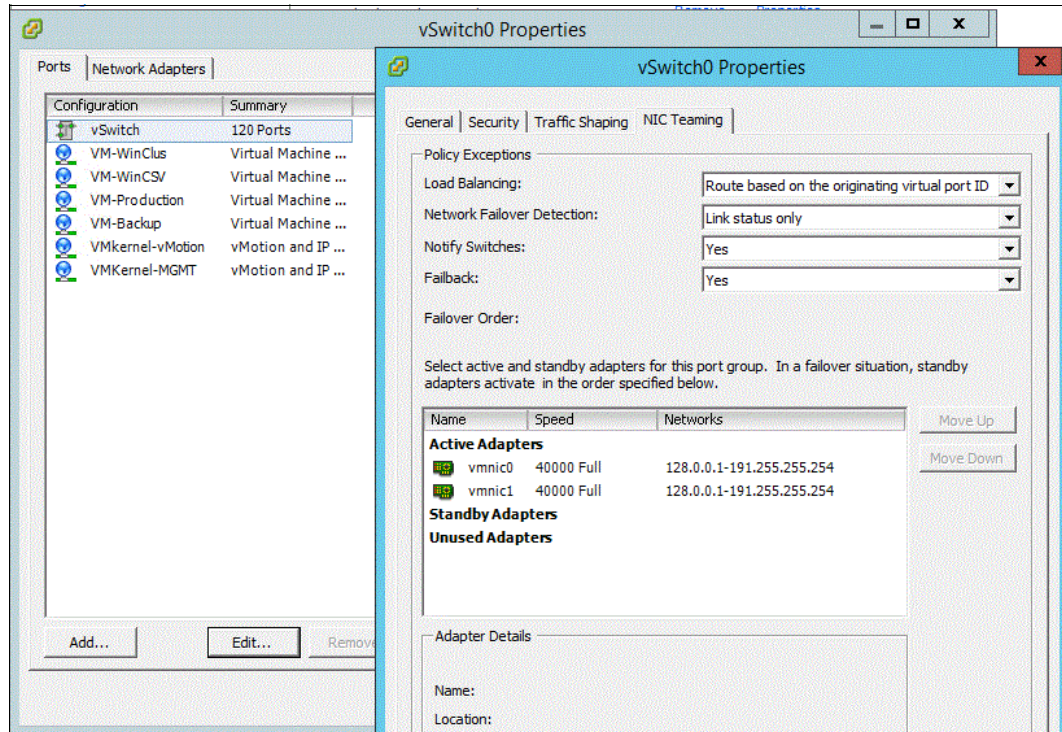


Figure 10-10 vmNICs in the Active Adapters list

10.7 Mapping the required VMFS Datastores

In this section, you map the VMFS Datastores to the hosts.

10.7.1 Mapping the VMFS Datastores to the first host

Note: The second host will be mapped after the cluster is created.

To map the VMFS Datastores to the first host, complete the following steps:

1. Log in to the IBM Storwize V7000 management GUI.
2. Select the volumes icon in the left pane and click the **Volumes** menu item.
3. Right-click the infra_datastore_1 volume, infra_datastore_2, sql_rdm_data, sql_rdm_log, and sql_rdm_quorum, and click **Map to Host**.
4. Select **vm-host-infra-1**, click **Map Volumes**, and then click **Close**.

ESXi Host vm-host-infra-01

To mount the required datastores, complete the following steps on the first ESXi host:

1. From the vSphere Client, select **vm-host-infra-01** in the inventory.
2. Click the **Configuration** tab to enable the configurations.
3. Click **Storage** in the Hardware pane.
4. In the Datastore area, click **Add Storage** to open the Add Storage wizard.

5. Select **Disk/Lun** and click **Next**.
6. Select the **1 TB Datastore** LUN and click **Next**.
7. Accept the default VMFS setting and click **Next**.
8. Click **Next** for the disk layout.
9. Enter `infra_datastore_1` as the datastore name.
10. Click **Next** to retain the maximum available space.
11. Click **Finish**.
12. Click **Add Storage** to open the Add Storage wizard.
13. Select **Disk/Lun** and click **Next**.
14. Select the **256 GB Datastore** LUN and click **Next**.
15. Accept the default VMFS setting and click **Next**.
16. Click **Next** for the disk layout.
17. Enter `infra_datastore_2` as the datastore name.
18. Click **Next** to retain the maximum available space.
19. Click **Finish**.

ESXi Hosts `vm-host-infra-01` and `vm-host-infra-02`

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From each vSphere Client, select the host in the inventory.
2. Click the **Configuration** tab to enable the configurations.
3. Click **Time Configuration** in the Software pane.
4. Click Properties at the upper right side of the window.
5. At the bottom of the Time Configuration dialog box, click **Options**.
6. In the NTP Daemon Options dialog box, complete the following steps:
 - a. Click **General** in the left pane and select **Start and stop with host**.
 - b. Click **NTP Settings** in the left pane and click **Add**.
7. In the Add NTP Server dialog box, enter `<<var_global_ntp_server_ip>>` as the IP address of the NTP server and click **OK**.
8. In the NTP Daemon Options dialog box, select the **Restart NTP Service to Apply Changes** check box and click **OK**.
9. In the Time Configuration dialog box, complete the following steps:
 - a. Select the **NTP Client Enabled** check box and click **OK**.
 - b. Verify that the clock is now set to the correct time.

Note: The NTP server time might vary slightly from the host time.

10.8 Storage I/O Control

Storage I/O Control (SIOC) allows for an increase in the number of VMs per datastore by monitoring datastore latency and adjusting the I/O load that is sent to it.

To configure SIOC, complete the following steps:

1. On vm-host-infra-01, go to **Configuration** and then click **Storage** in the left pane.
2. Right-click the first datastore and select Properties.
3. Select the **Storage I/O Control** check box.
4. Click Close.
5. Repeat steps 1 - 4 for every datastore on both hosts.

Figure 10-11 shows the properties of a boot LUN with SIOC enabled.

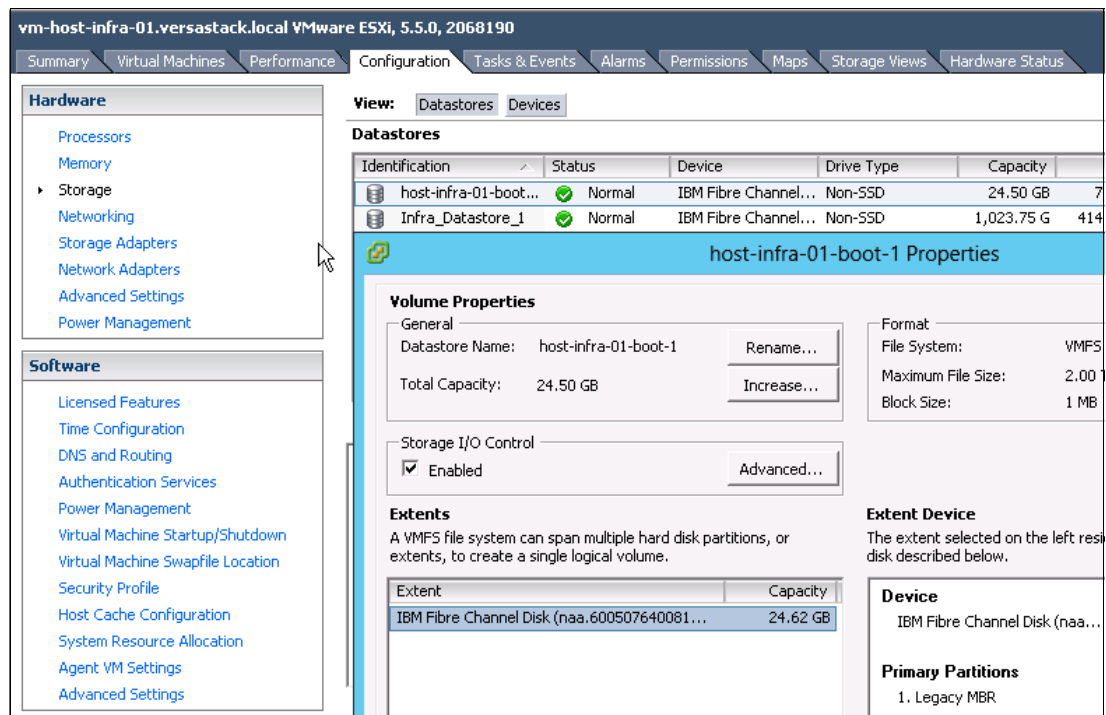


Figure 10-11 Enable SIOC

10.9 VersaStack VMware vCenter 5.5 Update 2

The procedures in the following subsections provide instructions for installing VMware vCenter 5.5 Update 2 in a VersaStack environment. This section focuses on the simple installation of vCenter Server on a Windows virtual machine. This section does not provide the steps or instructions to create and build virtual machines for vCenter Server and Active Directory that are used in this environment. For more information about vCenter Server installation methods and their hardware and software requirements, see the ESXi and vCenter Server 5.5 documentation on the VMware website.

To install VMware vCenter 5.5 Update 2, an accessible Windows Active Directory (AD) Domain is necessary. If an existing AD Domain is not available, an AD virtual machine or AD pair can be set up in this VersaStack environment. For more information, see Appendix A, “Windows Active Directory and running configurations” on page 441.

10.9.1 Installation steps for a simple installation of vCenter Server 5.5

To perform a simple installation of vCenter Server 5.5, complete the following steps:

1. Mount the vSphere 5.5 installation media, go to the VMware vCenter 5.5 Update 2 (VIMSetup) ISO, select it, and click **Open**.
2. In the left pane, click **Simple Install** and then click **Install**.

Figure 10-12 shows the vSphere vCenter installation window.

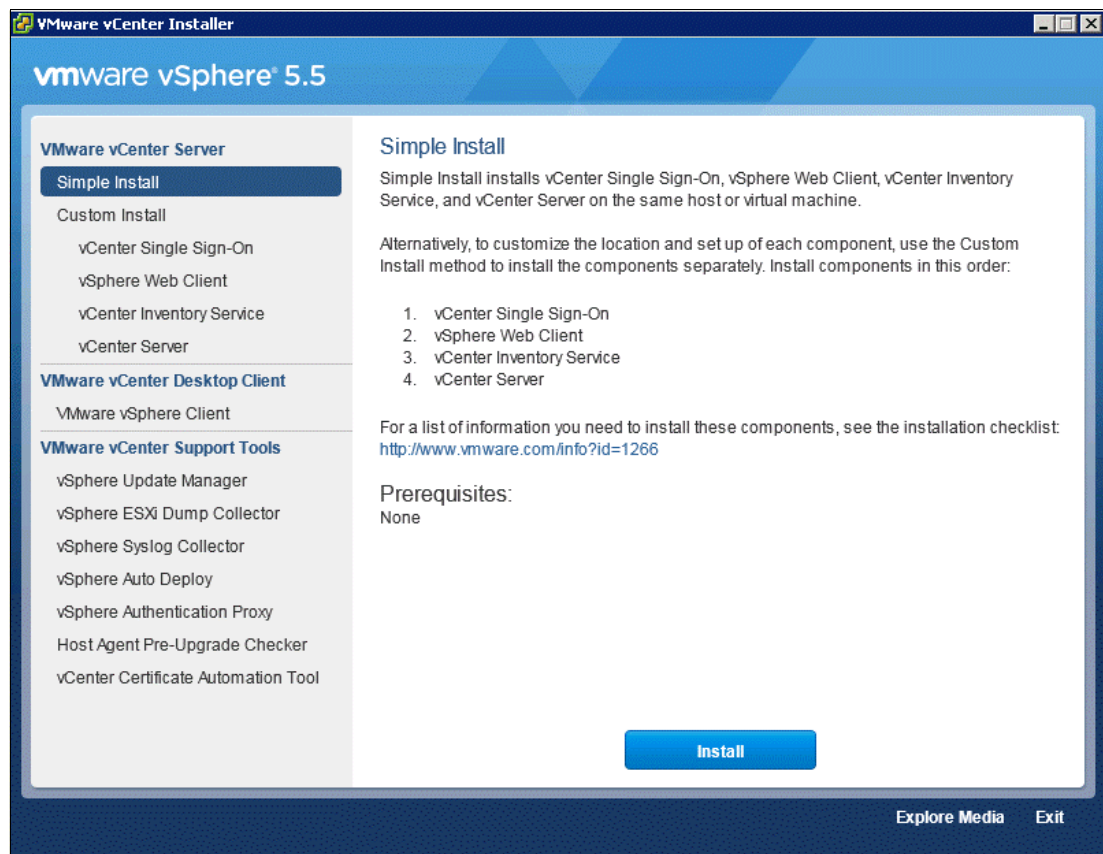


Figure 10-12 vCenter installation window

Note: If any of the prerequisites are not met, they are listed in the right pane under Prerequisites.

3. Click **Yes** if there is a User Account Control warning.

Figure 10-13 on page 161 shows the vCenter Single Sign On installation window

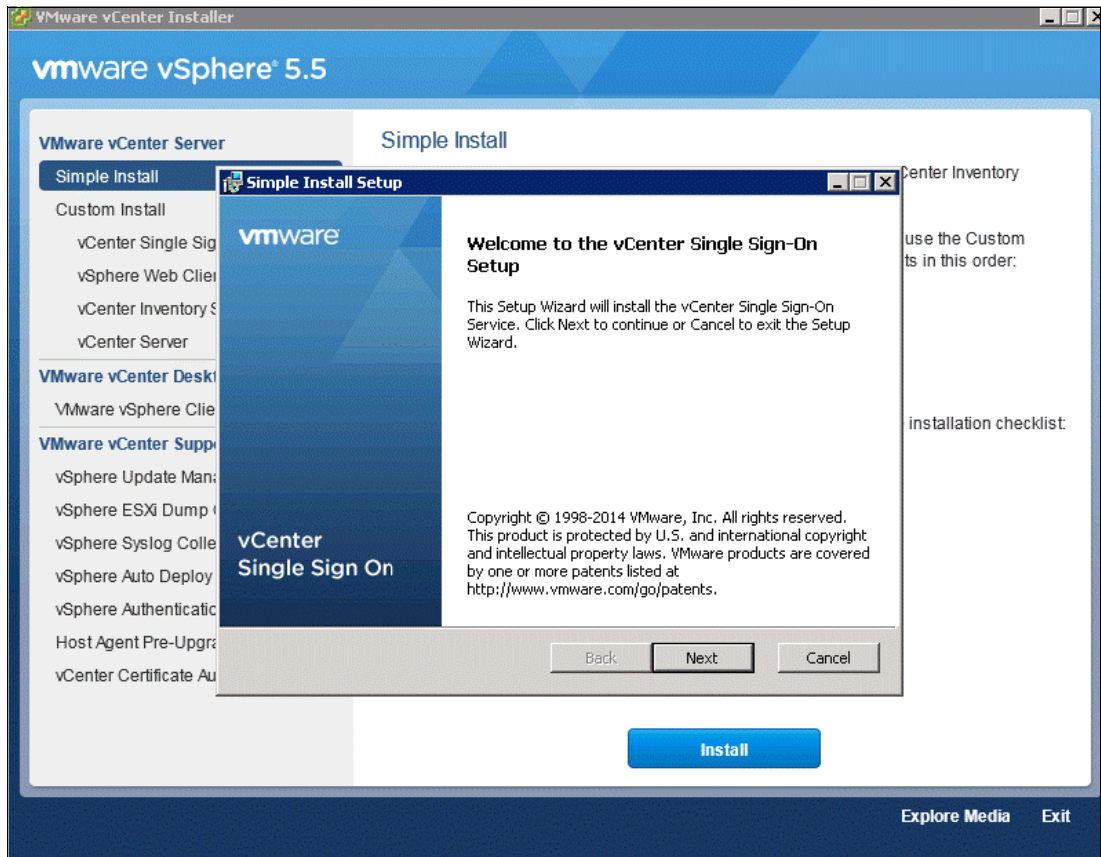


Figure 10-13 vCenter Single Sign On

4. Click **Next** to install vCenter Single Sign On.
5. Accept the terms of the license agreement and click **Next**.
6. In the Prerequisites window, click **Next**.

Figure 10-14 shows the Simple Install Prerequisites Check dialog box.

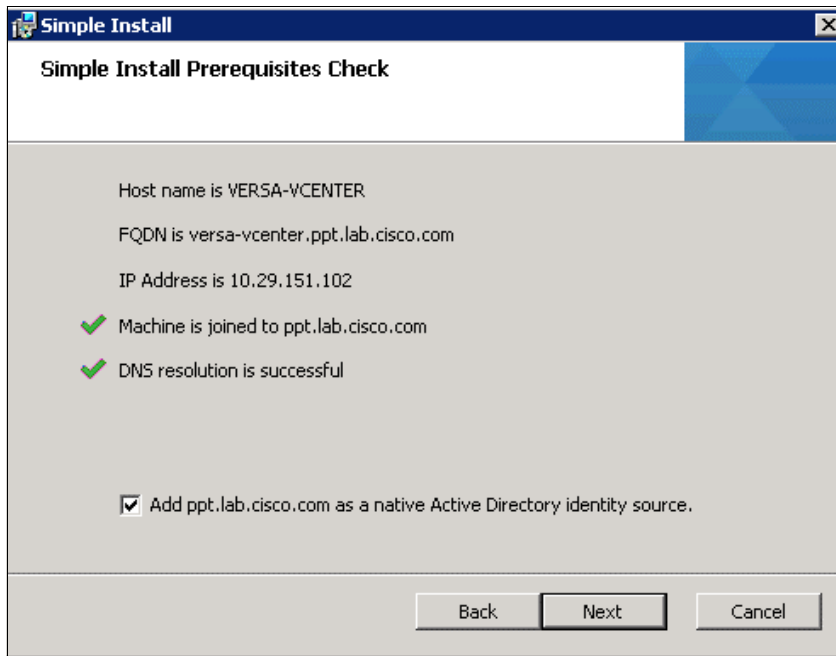


Figure 10-14 vCenter Simple Install Prerequisites Check

7. Enter and confirm <<var_password>> for the administrator user. Click **Next**.

Figure 10-15 shows the vCenter Single Sign-On Information dialog box.

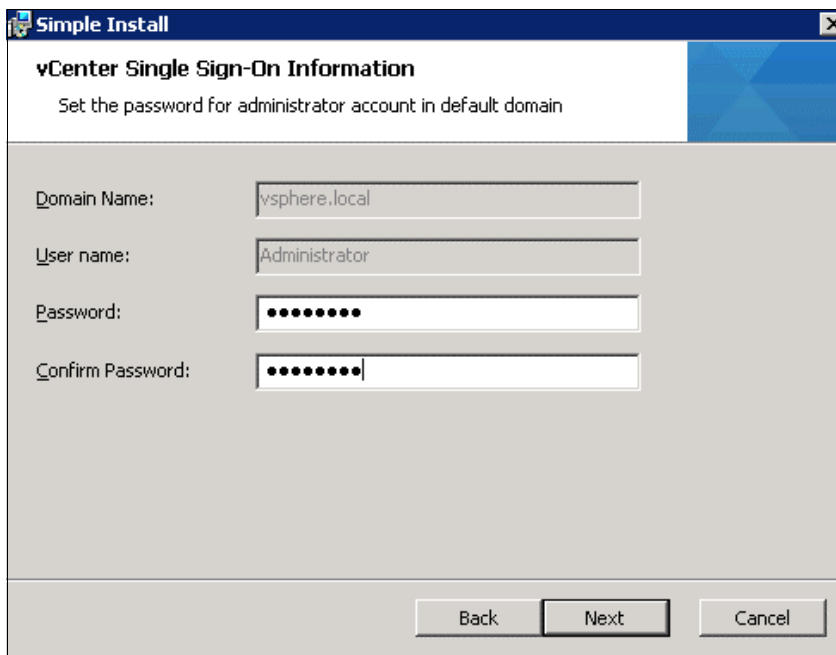


Figure 10-15 vCenter Single Sign-On Information window

Note: This dialog box shows information that is related to a domain with the name vsphere.local. This is not a domain that is auto-detected within the existing environment, but a new domain that is used internally by vSphere. The administrator@vsphere.local account performs the same function as the admin@System-Domain account in previous versions of vSphere.

8. In the Site window, click **Next**.
9. Accept the Default HTTPS port and click **Next**.

Figure 10-16 shows the port settings for a vCenter simple installation.




Figure 10-16 vCenter Simple Install Port Settings

10. Click **Next**.
11. Review the window and click **Install**. This process takes approximately 20 minutes, during which time multiple windows launch.

Figure 10-17 shows the vCenter Simple Install Information window for review.

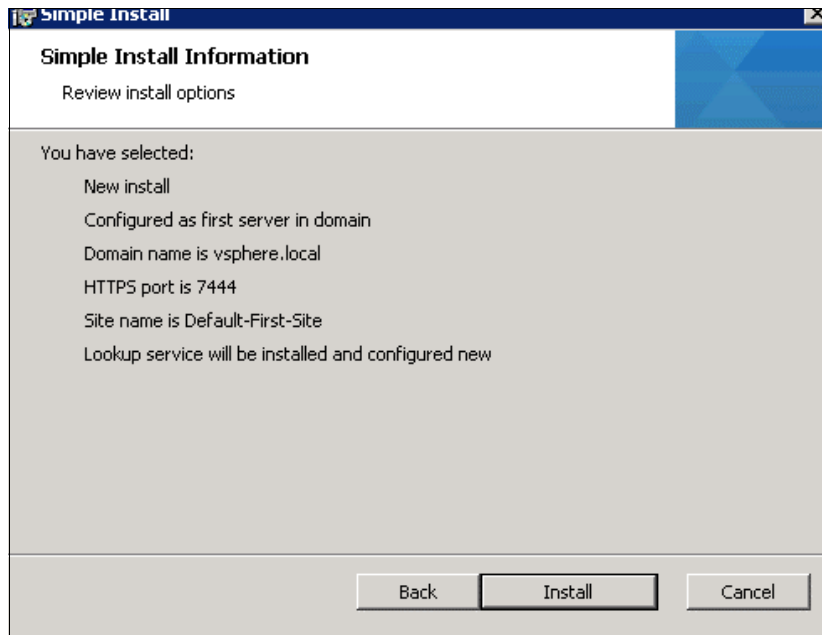


Figure 10-17 vCenter Simple Install Information Review

12. Enter **Yes** in the SSL window that opens.
13. Enter the license key for the vCenter Server.
14. Select **Install a Microsoft SQL Server 2008 Express instance database solution for vCenter Server** and then click **Next**.
15. Click **Next** to use the SYSTEM Account.
16. Click **Next** to accept the default ports.

Figure 10-18 on page 165 shows the window for configuring the default vCenter ports.

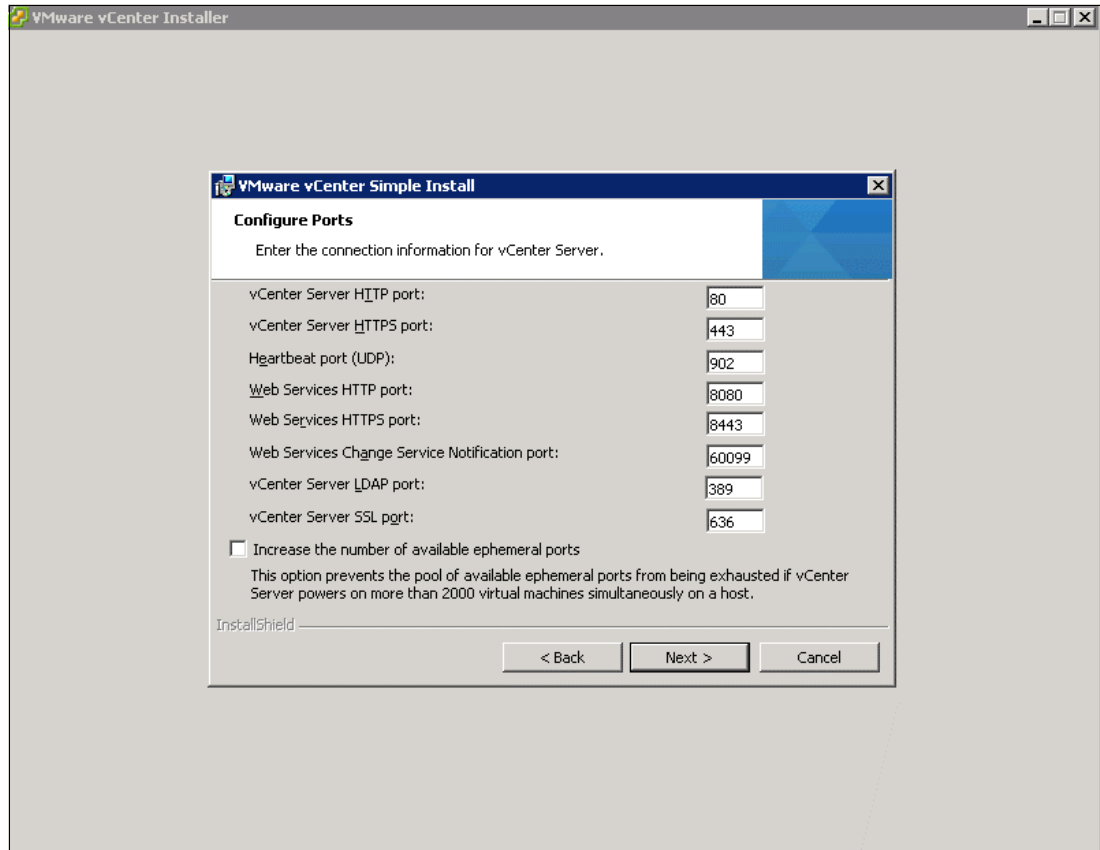


Figure 10-18 Configure the default vCenter Ports

17. Select the appropriate inventory size. Click **Next**.
18. Click **Install**. A new installer window starts and completes in approximately 10 minutes.
19. Click **Finish**, and then **OK**.

10.10 Setting up a vCenter Server

In this section, you learn how to set up a vCenter Server.

10.10.1 vCenter Server VM

To set up vCenter Server on the vCenter Server VM, complete the following steps:

1. Using the vSphere Client, log in to the newly created vCenter Server as the VersaStack admin user or administrator@vsphere.local.
2. Click **File** → **New** → **Datacenter** to create a datacenter.
3. Right-click the datacenter and enter VersaStack_DC_1 as the datacenter name.
4. Right-click the newly created VersaStack_DC_1 datacenter and select **New Cluster**.
5. Name the cluster VersaStack_Management and select the check boxes for **Turn On vSphere HA** and **Turn on vSphere DRS**.
6. Click **Next**.

Figure 10-19 shows creating the VersaStack_Management cluster on the vCenter.

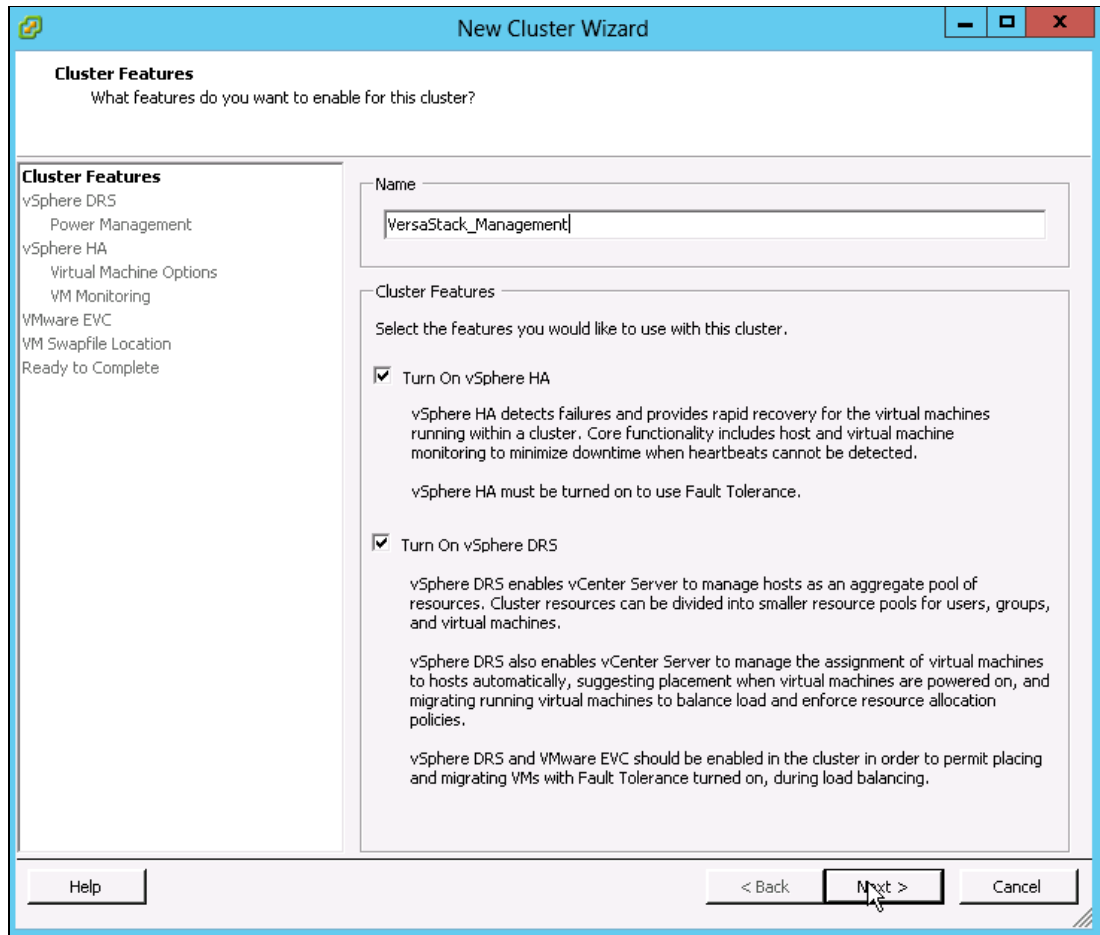


Figure 10-19 Create VersaStack_Management cluster on the vCenter

7. Accept the defaults for vSphere Distributed Resource Scheduler (DRS). Click **Next**.
8. Accept the defaults for Power Management. Click **Next**.
9. Accept the defaults for vSphere HA. Click **Next**.
10. Accept the defaults for Virtual Machine Options. Click **Next**.
11. Accept the defaults for VM Monitoring. Click **Next**.
12. Accept the defaults for VMware EVC. Click **Next**.

Important: If mixing Cisco UCS B or C-Series M3 and M4 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, see Enhanced vMotion Compatibility (EVC) Processor Support, found at the following website:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003212

13. Select **Store the swapfile in the same directory as the virtual machine**. Click **Next**.
14. Click **Finish**.
15. Right-click the newly created VersaStack_Management cluster and select **Add Host**.

16. In the Host field, enter either the IP address or the host name of the vm-host-infra-01 host. Enter root as the user name and the root password for this host. Click **Next**.
17. Click **Yes**.
18. Click **Next**.
19. Select **Assign a New License Key to the Host**. Click **Enter Key** and enter a vSphere license key. Click **OK**, and then click **Next**.
20. Click **Next**.
21. Click **Next**.
22. Click **Finish**. The vm-host-infra-01 host is added to the cluster.
23. Repeat this procedure to add vm-host-infra-02 to the cluster.

10.11 Mapping the datastores on the IBM Storwize V7000 second host after enabling the cluster

To map the datastores on the IBM Storwize V7000 second host after enabling the cluster, complete the following steps:

1. Open the web client for the Storwize V7000 storage system.
2. Click the Volumes icon in the left pane and select **Volume** to open the Volumes window.
3. Right-click the volumes **infra_datastore_1**, **infra_datastore_2**, **sql_rdm_data**, **sql_rdm_log**, and **sql_rdm_quorum**, and select **Map to Host**.
4. Select **vm-host-infra-02** and select **Map Volumes**.
5. Click **Map All Volumes** and click **Close**.
6. Click **Close** again.
7. In vSphere in the left pane, right-click the VersaStack_Management cluster and click **Rescan for Datastores**.
8. Click **OK**.

10.12 Optional: Adding domain account permissions

This section describes how to add a user to provide admin and login permissions in the vSphere web client and vSphere client. Complete the following steps:

1. Open a browser and enter `https://<<vSphere_ip>>:9443/vsphere-client/` to open the vSphere web client.
2. Log in as `administrator@vsphere.local` with the admin password.
3. Click the **Administration** item in the left pane
4. Select the **Configuration and Identity Sources** tab and validate that the domains that you require are listed. You can add other required domain sources by clicking the green +.

Figure 10-20 shows the Identity Sources tab, where you can add more domains.

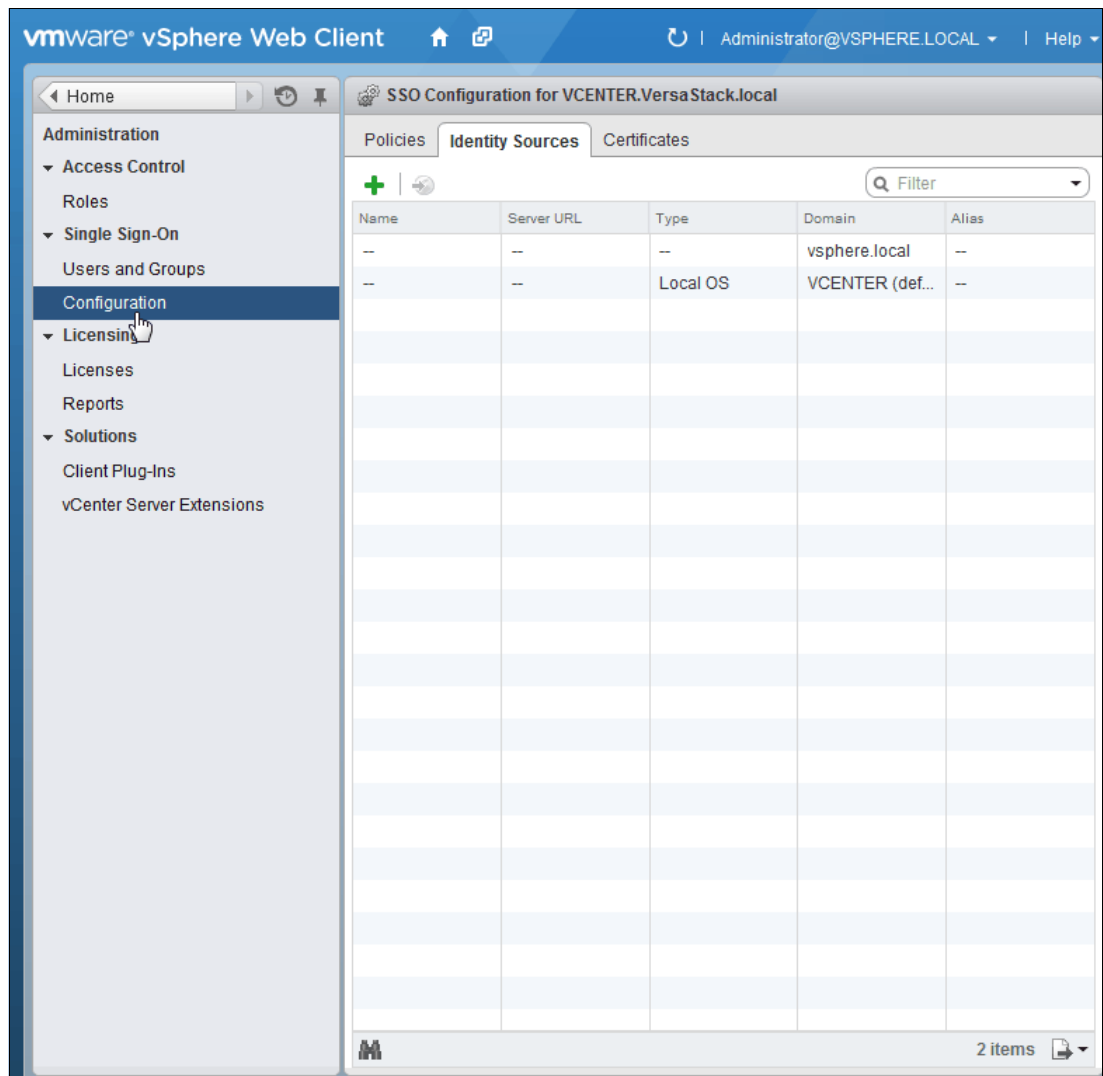


Figure 10-20 Identity Sources tab

5. Select the **Home** button at the upper left.
6. Click **vCenter** to show the vCenter window.
7. Click **vCenter Servers** under the Inventory list.
8. Click the vCenter server name in the left pane, and click the **Manage** tab in the right pane.
9. Click **Permissions**.
10. Click the green + sign to add a user. Select **Add** in the Add Permission window. Select the domain, and highlight the user.
11. Click **Add**, and then click **OK**.
12. For an assigned roll, select the administrator and then click **OK**. You may now log in as that user in your vSphere web client.
13. Open the vSphere Client application, log in as the administrator account, and right-click the vCenter name in the upper left and click **Add Permissions**.

Figure 10-21 shows the start of the process for adding a permission to vCenter.

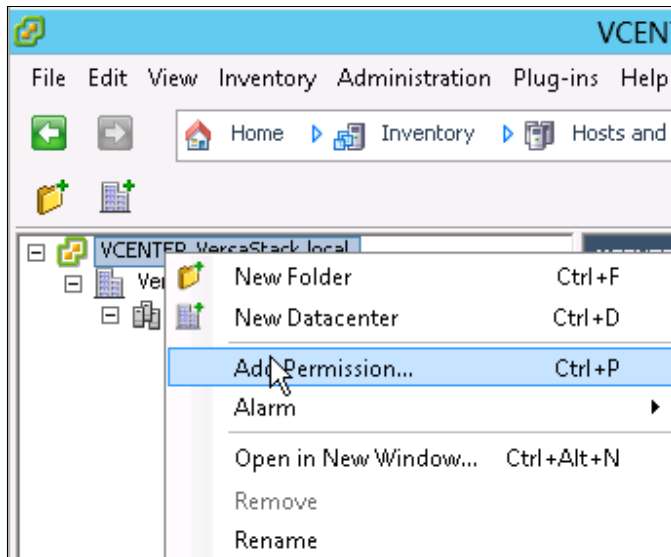


Figure 10-21 Add a permission

14. Click **Add**.
15. Select the correct domain in the drop-down menu.
16. Highlight a user and click **Add**, and then click **OK**.
17. Change the Assigned Role to **Administrator** in the drop-down menu, and click **OK**. You may now log off as administrator and back in as that domain user in the vSphere Client.

In this chapter, you set up two blades with Cisco Custom ESXi 5.5.0 U2 GA, set them up to SAN boot, and then configured the network settings and Storage I/O Control for these hosts. Finally, you set up vCenter Server 5.5 and created a VersaStack management cluster.



IBM DB2 High Availability server and failover cluster implementation

This chapter provides detailed instructions about how to accomplish the following tasks:

- ▶ Creating virtual machines
- ▶ Installing Windows Server 2012 R2
- ▶ Preparing the virtual machines for clustering
- ▶ Installing a Windows Server Failover Cluster
- ▶ Modifying the vSphere High Availability and Distributed Resource Scheduler settings for the Windows Server Failover Cluster virtual machines
- ▶ Installing DB2
- ▶ Installing DB2 Data Studio Client
- ▶ Deploying the DB2 GSDB sample database
- ▶ Configuring DB2 High Availability

For more information, see the PDF about setting up failover clustering and Microsoft cluster service that is found at the following website:

<https://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vc-enterprise-server-55-setup-mscs.pdf>

11.1 Creating virtual machines

In this section, you create two VMs, one on each ESXi host. Using the information that is shown in Table 11-1, you create a VM named SQLVM01 on the vm-host-infra-01 ESXi host and a second VM named SQLVM02 on the vm-host-infra-02 ESXi host.

Table 11-1 Virtual machines

VM name	ESXi hosting the VM	vCPU	Memory	Boot disk size	Boot disk store location	No. of network adapters	No. of shared Raw Device Mapping disks (RDMs)
SQL VM01	vm-host-infra-01	4	16 GB	100 GB	Infra_Datastore_1	3	3
SQL VM02	vm-host-infra-02	4	16 GB	100 GB	Infra_Datastore_1	3	3

Complete the following steps:

1. Open a browser to the vSphere web client by using the following URL:
https://<<vSphere_ip>>:9443/vsphere-client/
2. Log in as administrator@vsphere.local with the admin password, as shown in Figure 11-1.

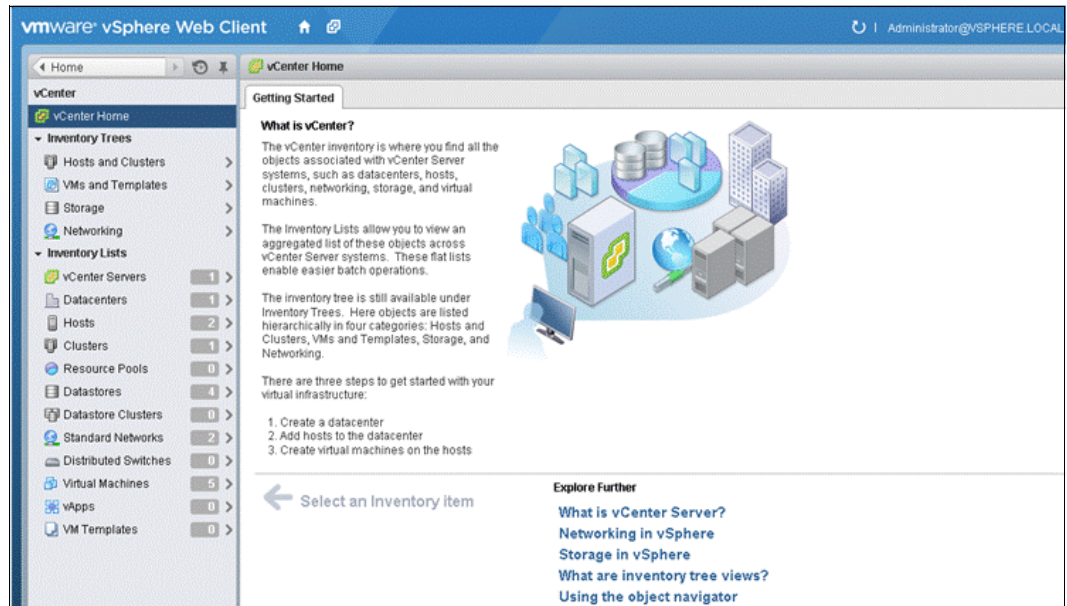


Figure 11-1 vSphere web client

3. Click **Hosts and Clusters** under Inventory Trees.
4. Under the Inventory List, right-click vm-host-infra-01 and select **New Virtual Machine**.
5. In the New Virtual Machine wizard, enter a name for the VM, select a datacenter, and click **Next**, as shown in Figure 11-2 on page 173.

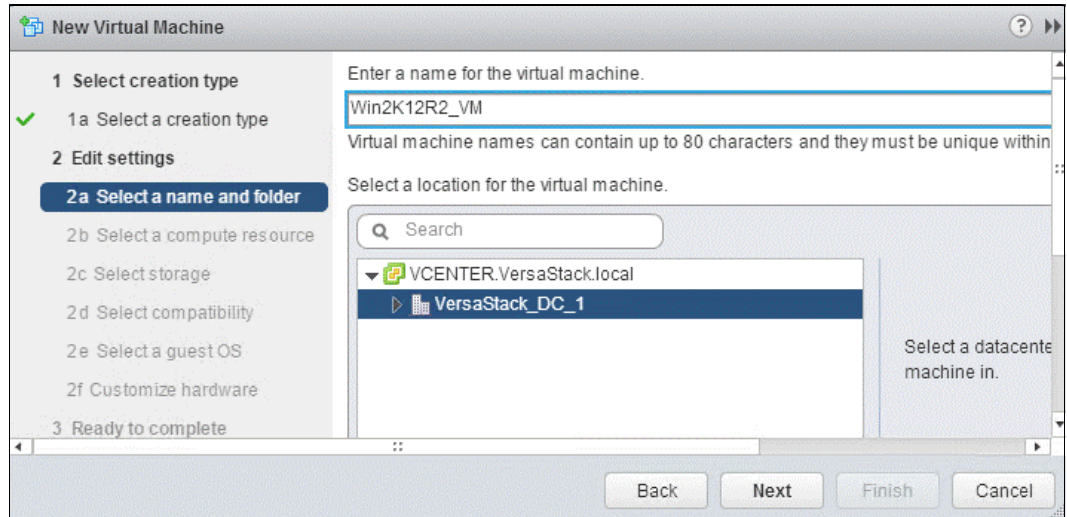


Figure 11-2 New Virtual Machine

6. Select a compute resource to host this VM and click **Next**, as shown in Figure 11-3.

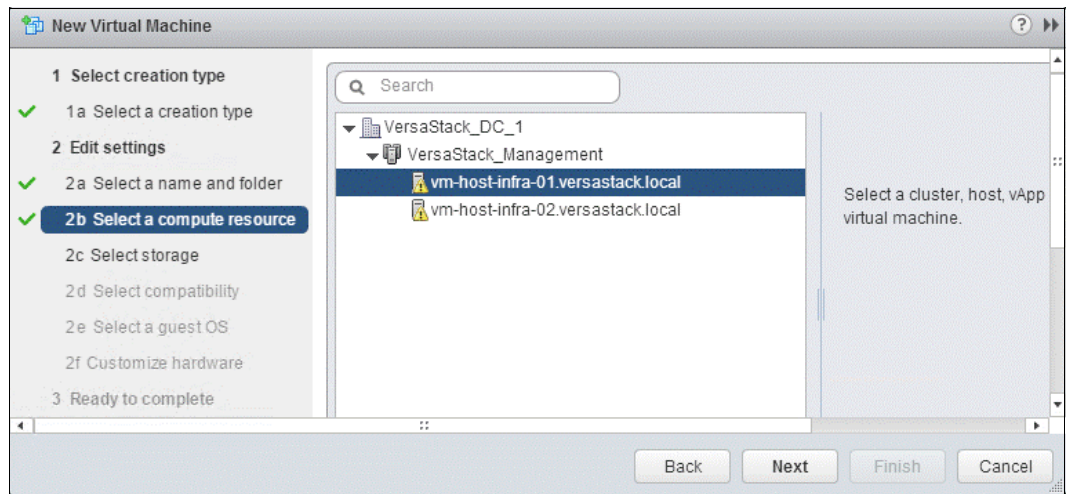


Figure 11-3 Select compute resource

7. In **Select storage**, choose **Infra_Datastore_1** as the storage location for the VM's disk and click **Next**. In this environment, this datastore is used to store the VM's boot disks. (See Figure 11-4.)

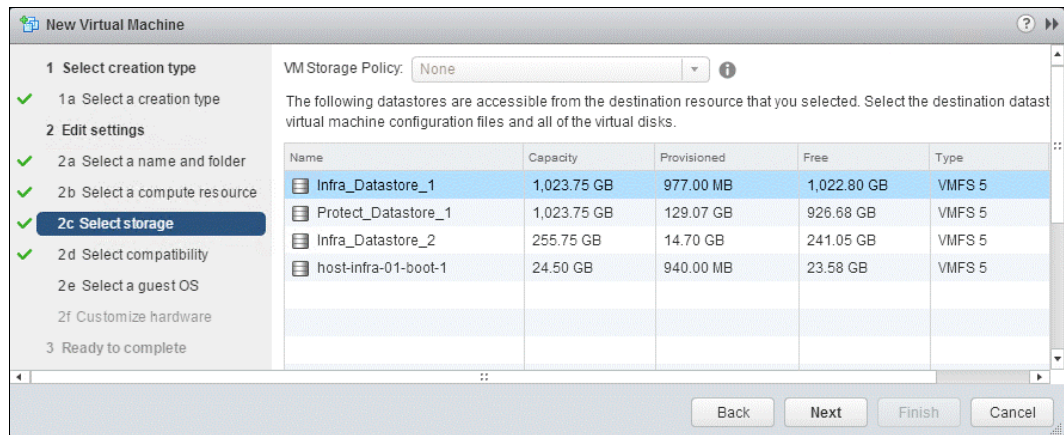


Figure 11-4 Select storage

8. In the **Select Compatibility** window, select **ESXi 5.5 and Later** to create a VM Version 10. Click **Next** (see Figure 11-5).

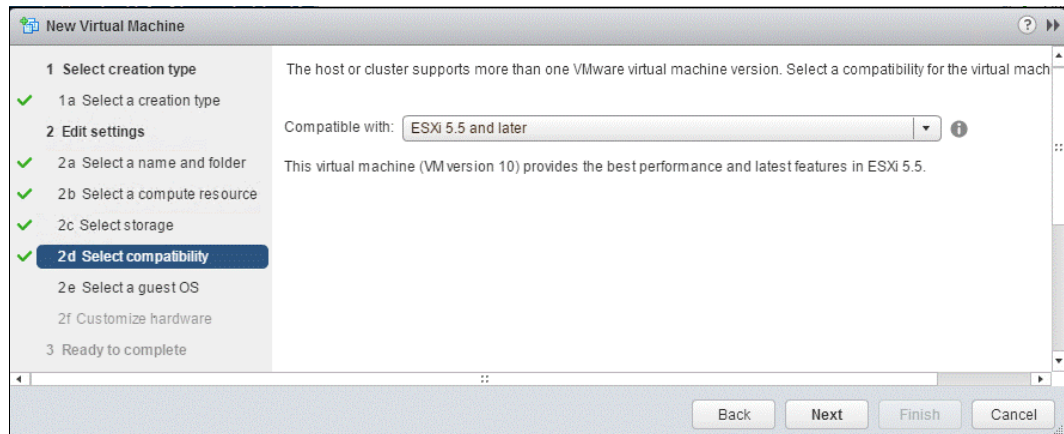


Figure 11-5 Select compatibility

9. In the Select a guest OS window, select **Windows** from the drop-down list next to Guest OS Family and select **Microsoft Windows Server 2012 (64-bit)** as the Guest OS version. Click **Next** (see Figure 11-6).



Figure 11-6 Select a guest OS

10. In the Customize hardware window, complete the following steps:

- a. Assign memory and vCPUs.
- b. Assign a hard disk size for the OS.
- c. Select **Thick Provision Eager Zeroed** for Disk provisioning.
- d. Select **LSI Logic SAS** as the SCSI controller type.
- e. Select **VM-Production** from the drop-down list for Network Adapter 1.
- f. Map and mount the Windows Server 2012 R2 installation ISO file. (See Figure 11-7.)

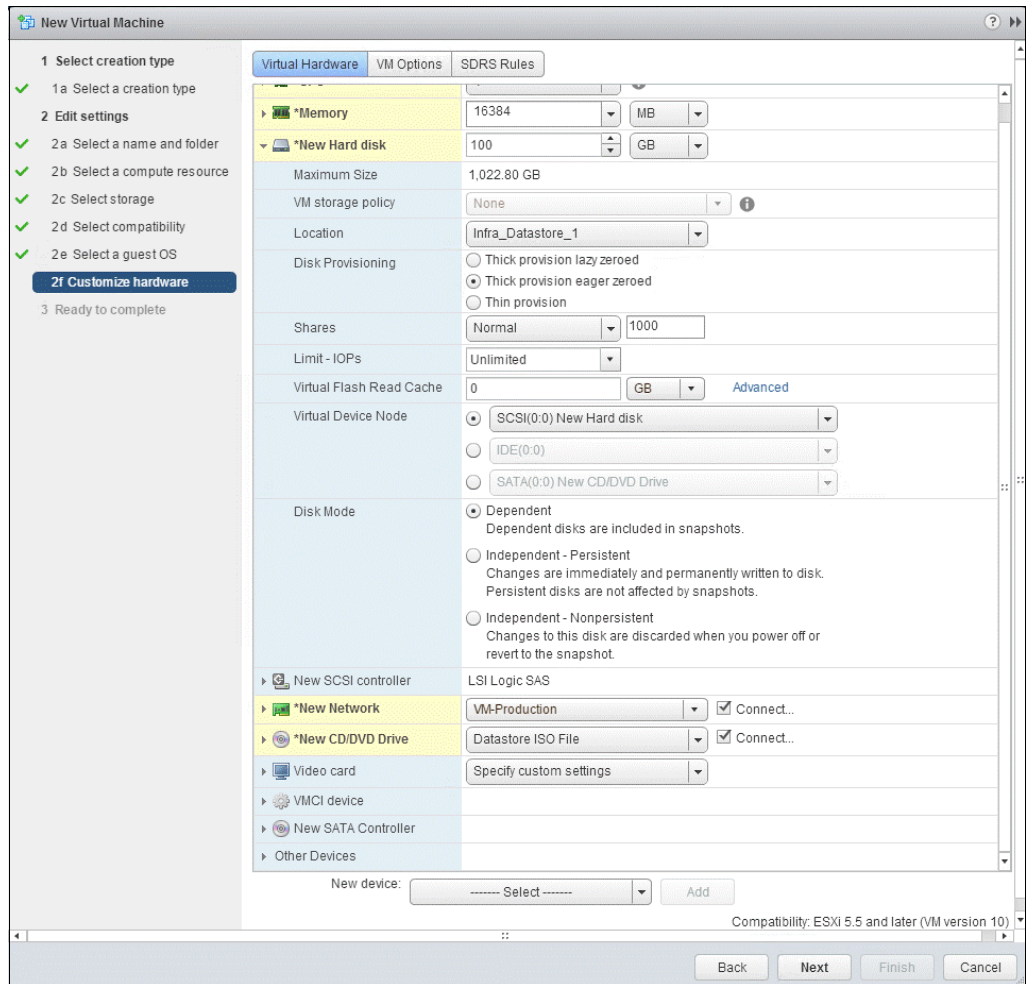


Figure 11-7 Customize hardware

11. In the Ready to complete window, click **Finish** (see Figure 11-8 on page 177).

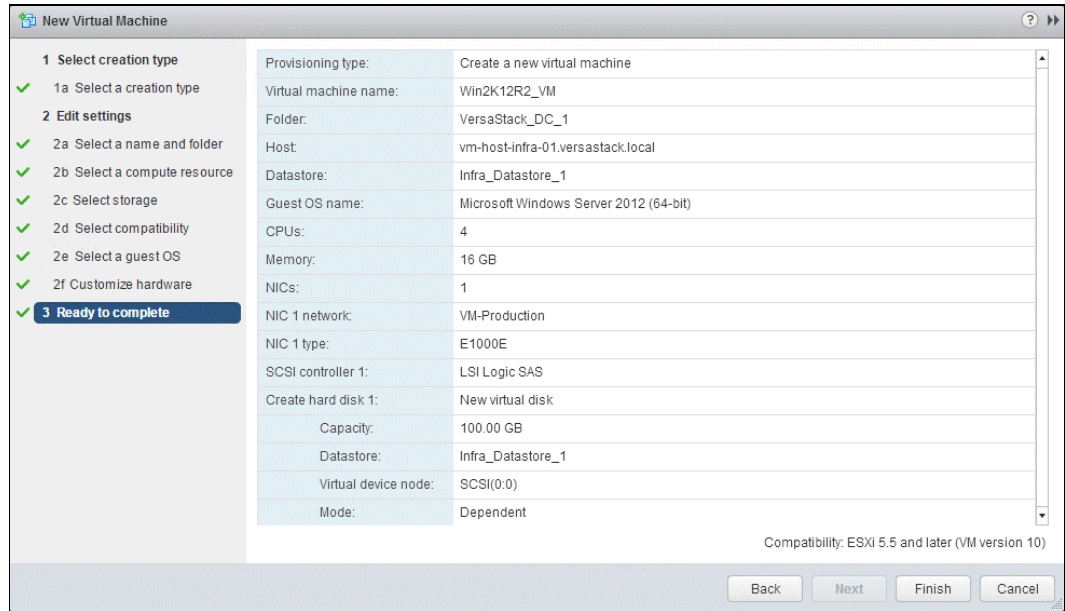


Figure 11-8 Ready to complete

12. Select the new VM, right-click it, and click **Edit Settings**.

13. Select **Network** from the drop-down list next to New Device and click **Add** (Figure 11-9).

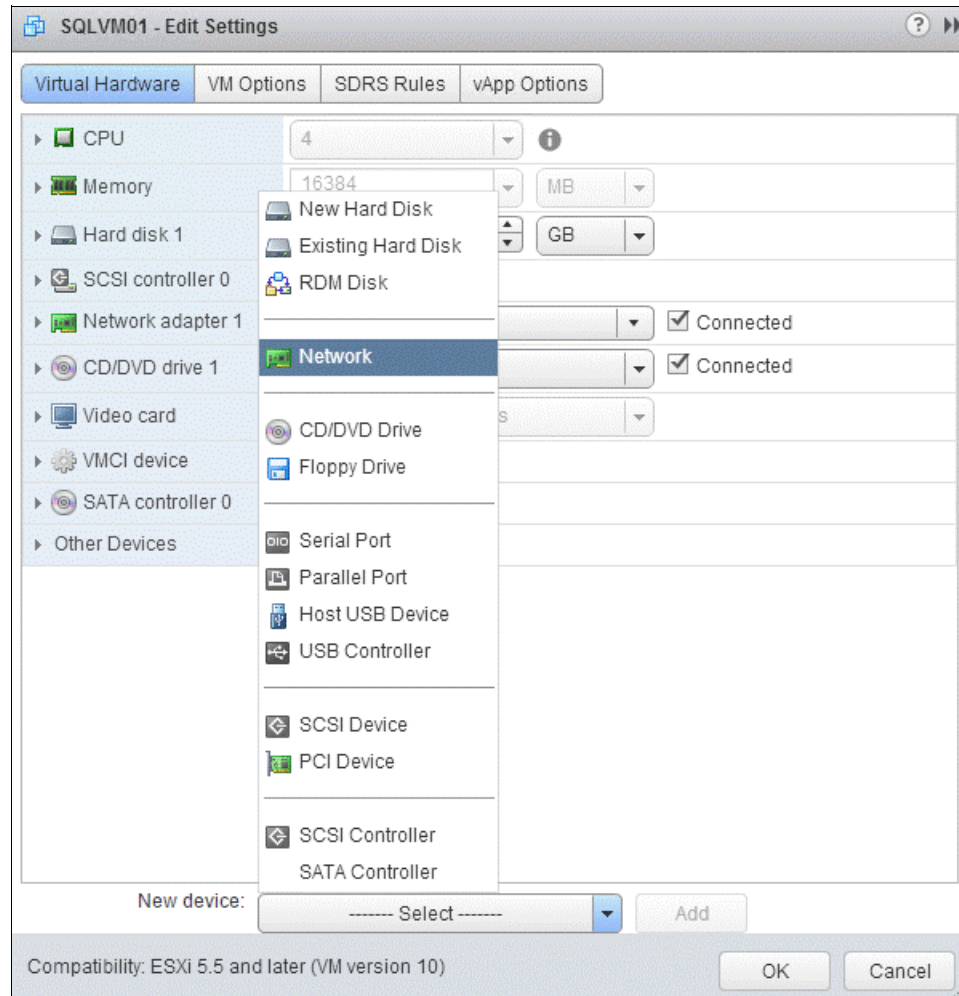


Figure 11-9 Edit Settings

14. Add the other two network adapters that will be used for the Microsoft Windows Server Failover Cluster (WSFC), as shown in Figure 11-10 on page 179.

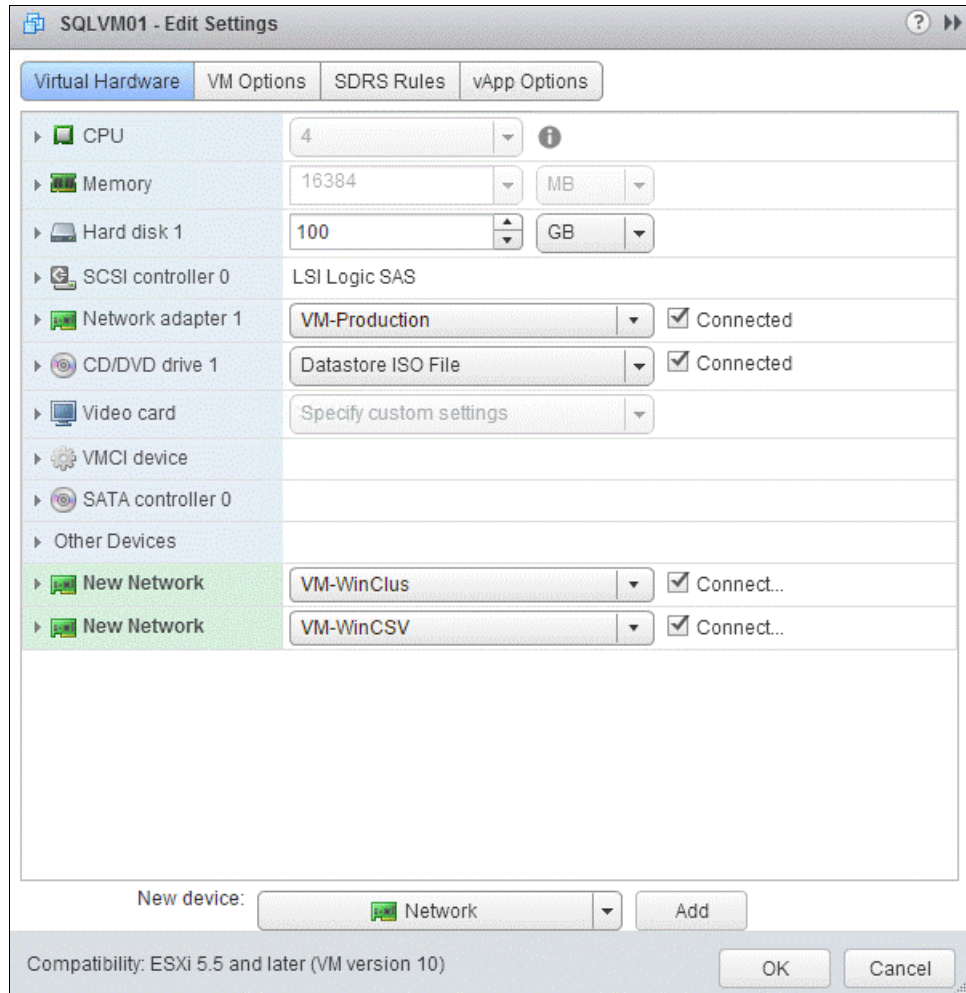


Figure 11-10 Add network adapters

15. Repeat steps 3 on page 172 to 14 on page 178 to create the second VM on vm-host-infra-02.

11.2 Installing Windows Server 2012 R2

This section provides instructions about how to install Windows Server 2012 R2 on the new VMs. To install the OS on both VMs, complete the following steps:

1. Go to the **Edit Settings** of the VM that was created in 11.1, “Creating virtual machines ” on page 172, mount the Windows Server ISO image, and power on the VM to begin the installation.

2. Select the appropriate language and other preferences and click **Next** (see Figure 11-11).

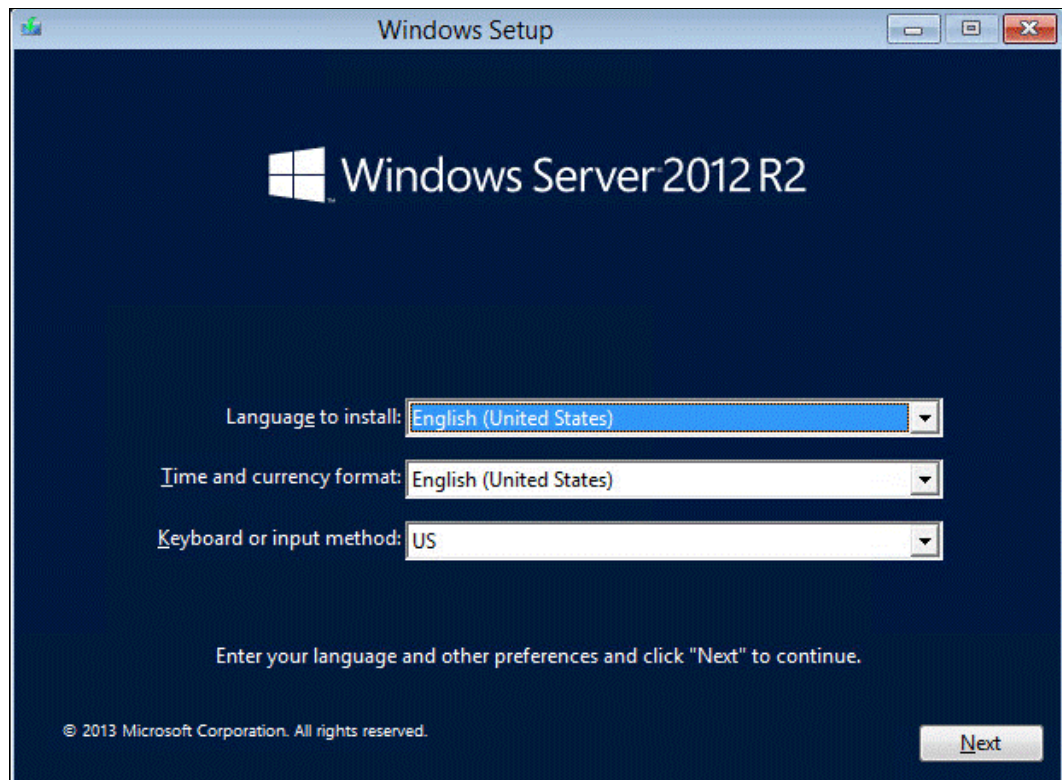


Figure 11-11 Select preferences

3. In the next window, click **Install now**.
4. In the next window, select the operating system to install and click **Next**.
5. Accept the license terms and click **Next**.
6. Select **Custom: Install Windows only (advanced)** and click **Next**.
7. Select the disk to install Windows and click **Next** (see Figure 11-12 on page 181).

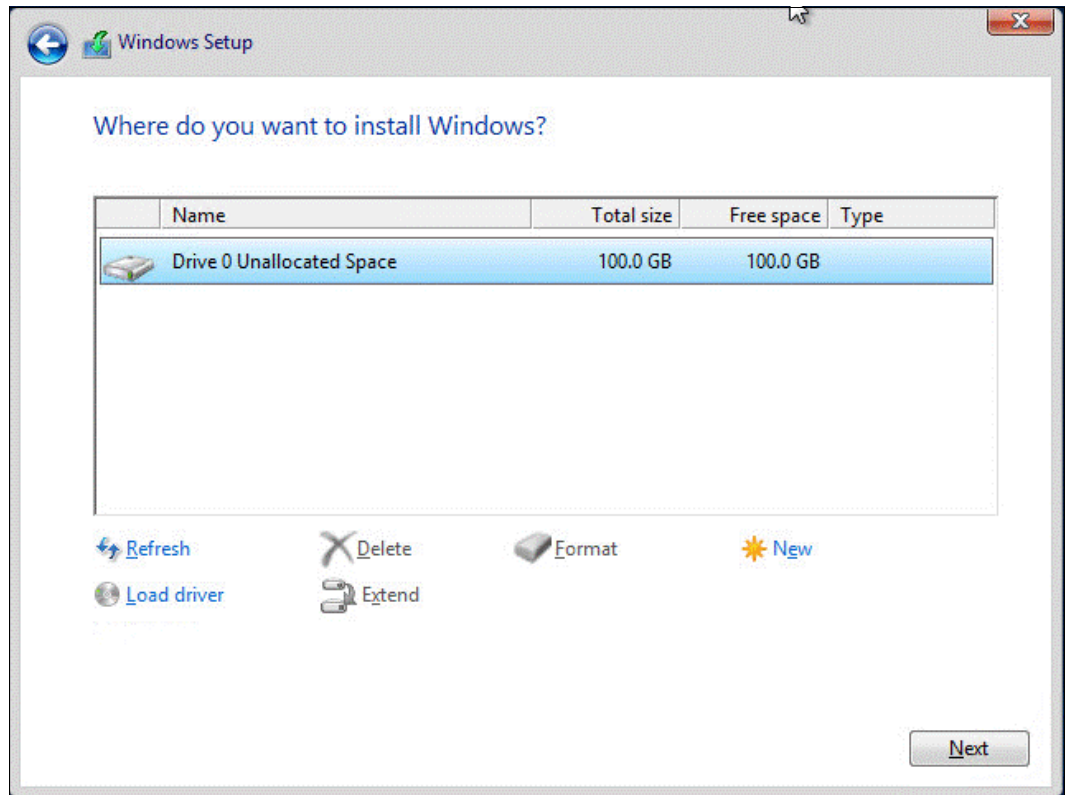


Figure 11-12 Choose where to install

8. The installation begins and restarts the VM upon completion.

9. Provide a password for the VM's built-in administrator account (see Figure 11-13).

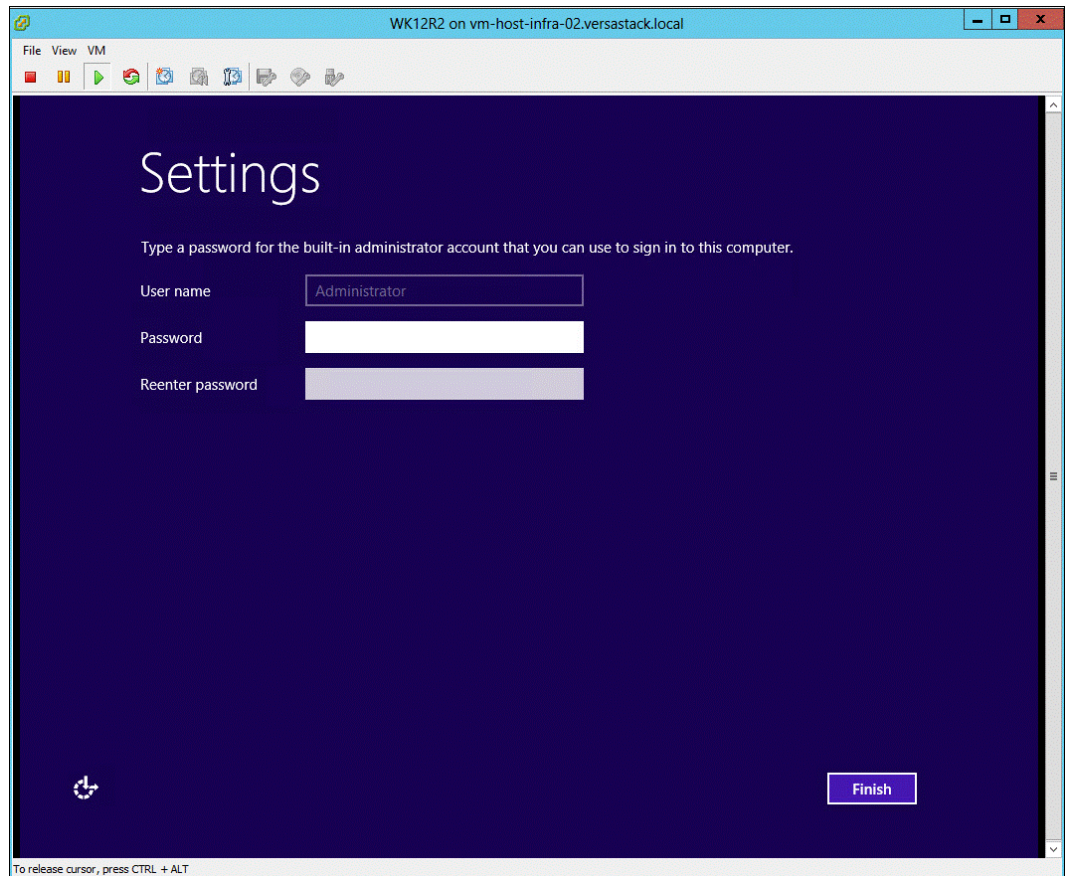


Figure 11-13 Password

11.3 Preparing the virtual machines for clustering

This section provides instructions about preparing the VMs for setting up WSFC later. This section describes the following topics:

- ▶ Renaming and assigning IP addresses to network adapters
- ▶ Enabling jumbo frames for CSV traffic
- ▶ Configuring the network adapters binding order
- ▶ Installing Windows updates and adding roles and features
- ▶ Adding Raw Device Mapping disks to the first virtual machine node
- ▶ Adding Raw Device Mapping disks to the second virtual machine node
- ▶ Preparing the disks for cluster use

11.3.1 Renaming and assigning IP addresses to network adapters

To rename and assign IP addresses to the network adapters, complete the following steps on both VMs:

1. Log in to the VM by using the administrator account.
2. Rename the computer host name and restart the VM.

- Optionally, rename the adapters according to their role for easy identification and troubleshooting purposes, as shown in Figure 11-14.

```

PS C:\Users\administrator.VERSASTACK> Get-NetAdapter

Name                           InterfaceDescription           ifIndex Status      MacAddress           LinkSpeed
-----
Ethernet1                       Intel(R) 82574L Gigabit Network Co...#3 31 Up          00-50-56-B4-0B-18   1 Gbps
Ethernet2                       Intel(R) 82574L Gigabit Network Co...#2 27 Up          00-50-56-B4-27-77   1 Gbps
Public                           Intel(R) 82574L Gigabit Network Conn... 13 Up          00-50-56-B4-3E-62   1 Gbps

PS C:\Users\administrator.VERSASTACK> Rename-NetAdapter -Name "Ethernet1" -NewName Private_WinClus
PS C:\Users\administrator.VERSASTACK> Rename-NetAdapter -Name "Ethernet2" -NewName Private_WinCSV
PS C:\Users\administrator.VERSASTACK> Get-NetAdapter

Name                           InterfaceDescription           ifIndex Status      MacAddress           LinkSpeed
-----
Private_WinClus                 Intel(R) 82574L Gigabit Network Co...#3 31 Up          00-50-56-B4-0B-18   1 Gbps
Private_WinCSV                  Intel(R) 82574L Gigabit Network Co...#2 27 Up          00-50-56-B4-27-77   1 Gbps
Public                           Intel(R) 82574L Gigabit Network Conn... 13 Up          00-50-56-B4-3E-62   1 Gbps

```

Figure 11-14 Adapters

- To rename the network adapters to reflect their role, gather network adapter information from the VM settings, such as MAC address and which virtual switch it is connected to, as shown in Figure 11-15.

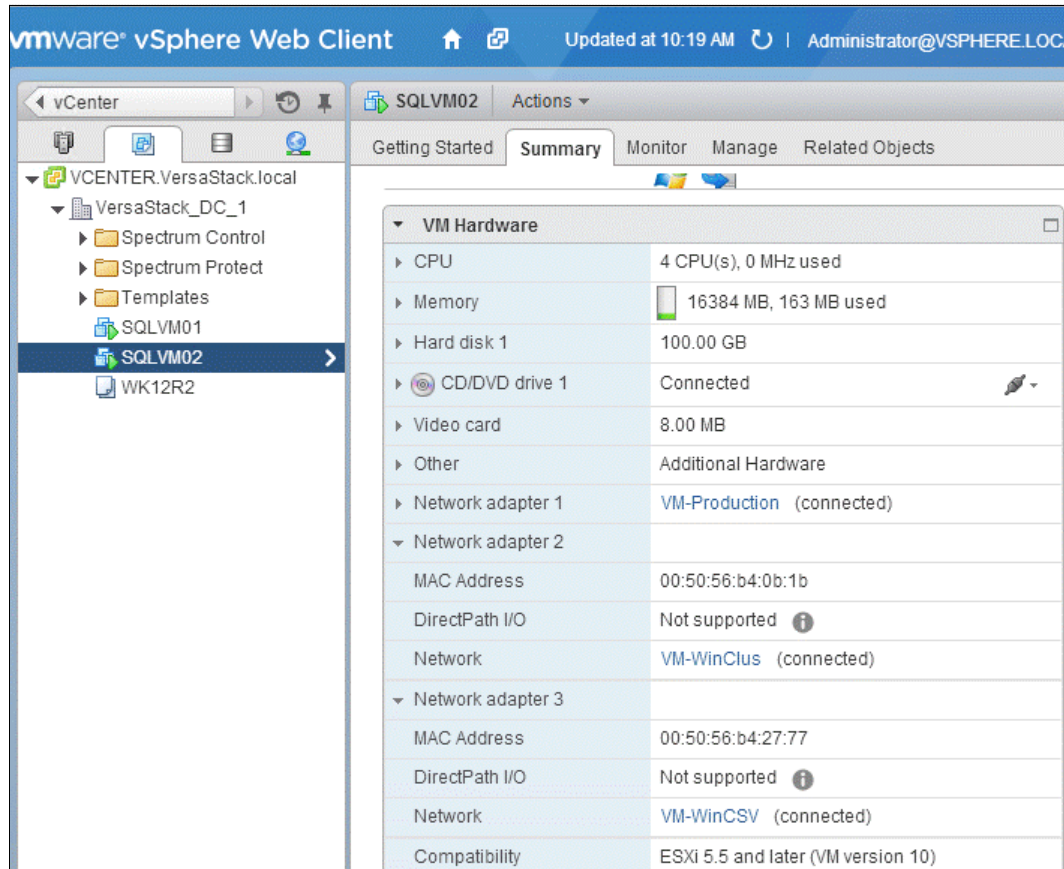


Figure 11-15 Gather network information

5. Assign static IP addresses to the network interfaces. An example of assigning IP addresses by using PowerShell is shown in Figure 11-16.

```
PS C:\Users\administrator.VERSASTACK> New-NetIPAddress -ifIndex 31 -IPAddress 192.168.40.52 -PrefixLength 24

IPAddress           : 192.168.40.52
InterfaceIndex      : 31
InterfaceAlias      : Private_WinClus
AddressFamily       : IPv4
Type                : Unicast
PrefixLength        : 24
PrefixOrigin        : Manual
SuffixOrigin        : Manual
AddressState        : Tentative
ValidLifetime       : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime   : Infinite ([TimeSpan]::MaxValue)
SkipAsSource        : False
PolicyStore         : ActiveStore

IPAddress           : 192.168.40.52
InterfaceIndex      : 31
InterfaceAlias      : Private_WinClus
AddressFamily       : IPv4
Type                : Unicast
PrefixLength        : 24
PrefixOrigin        : Manual
SuffixOrigin        : Manual
AddressState        : Invalid
ValidLifetime       : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime   : Infinite ([TimeSpan]::MaxValue)
SkipAsSource        : False
PolicyStore         : PersistentStore

PS C:\Users\administrator.VERSASTACK> New-NetIPAddress -ifIndex 27 -IPAddress 192.168.50.52 -PrefixLength 24

IPAddress           : 192.168.50.52
InterfaceIndex      : 27
InterfaceAlias      : Private_WinCSV
AddressFamily       : IPv4
Type                : Unicast
PrefixLength        : 24
PrefixOrigin        : Manual
SuffixOrigin        : Manual
AddressState        : Tentative
ValidLifetime       : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime   : Infinite ([TimeSpan]::MaxValue)
SkipAsSource        : False
PolicyStore         : ActiveStore

IPAddress           : 192.168.50.52
InterfaceIndex      : 27
InterfaceAlias      : Private_WinCSV
AddressFamily       : IPv4
Type                : Unicast
PrefixLength        : 24
PrefixOrigin        : Manual
SuffixOrigin        : Manual
AddressState        : Invalid
ValidLifetime       : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime   : Infinite ([TimeSpan]::MaxValue)
SkipAsSource        : False
PolicyStore         : PersistentStore
```

Figure 11-16 Assign IP addresses

11.3.2 Enabling jumbo frames for CSV traffic

To enable jumbo frames for the CSV traffic, complete the following steps on both VMs:

1. Click **Settings** → **Network and Internet** → **Network and Sharing Center** → **Change Adapter Settings**.
2. Right-click the Private_WinCSV network adapter, select **Properties**, and click **Configure**.
3. In the advanced settings, set the Jumbo Packet property value to 9014 bytes, as shown in Figure 11-17 on page 185.

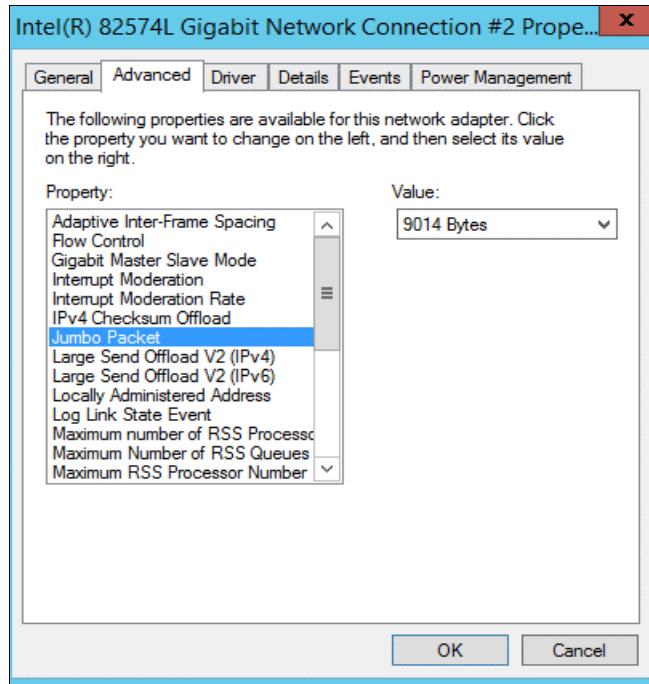


Figure 11-17 Jumbo packet

11.3.3 Configuring the network adapters binding order

To configure the network adapters binding order, complete the following steps on both VMs:

1. Click **Control Panel** → **Network and Internet** → **Network Connections**.
2. Press the Alt key.
3. Click **Advanced** and select **Advanced settings**.

4. Use the Up and Down arrow buttons to configure the adapters binding order, as shown in Figure 11-18.

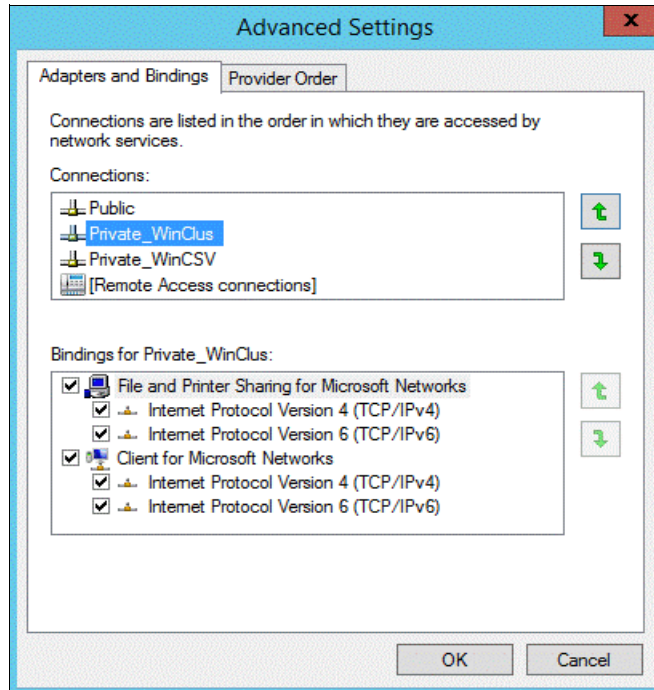


Figure 11-18 Binding order

11.3.4 Installing Windows updates and adding roles and features

To install the latest Windows updates and add the roles and features that are required for WSFC and SQL Server FCI, complete the following steps on both VMs:

1. Install the latest updates and patches from the Microsoft website and make sure that the current version of the VMware tool is running.
2. Join the computer to an Active Directory domain and restart the machine.
3. Click **Server Manager** → **Add Roles and Features** and install the .NET Framework 3.5 and Failover Clustering features, as shown in Figure 11-19 on page 187.

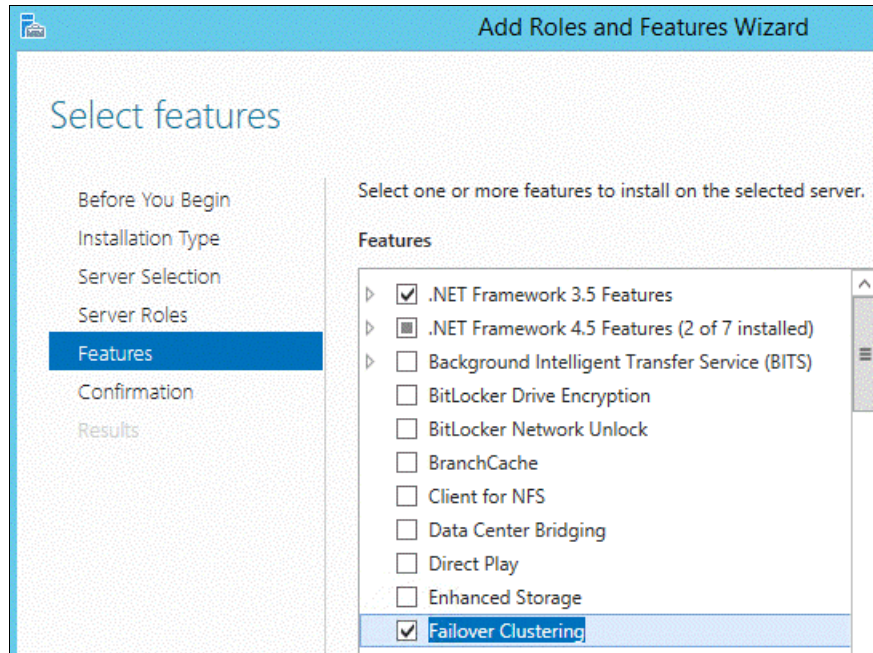


Figure 11-19 Add roles and features

11.3.5 Adding Raw Device Mapping disks to the first virtual machine node

This section provides detailed instructions about how to map the shared LUNs that are presented to ESXi hosts as RDMs to the VMs. An RDM in physical compatible mode is required for clustering VMs running on different ESXi hosts. After the completion of this task, these RDMs are used for creating a WSFC and installing an SQL Server failover cluster instance.

Complete the following steps:

1. In the vSphere Web Client navigator, select the created VM, right-click it, and select **Edit settings**.
2. On the Customize hardware window, click the **Virtual Hardware** tab.
3. Click the **New Device** drop-down menu, select **SCSI Controller**, and click **Add**.
4. Make sure that the SCSI Controller type is LSI Logic SAS, and SCSI Bus Sharing is Physical, as shown in Figure 11-20.

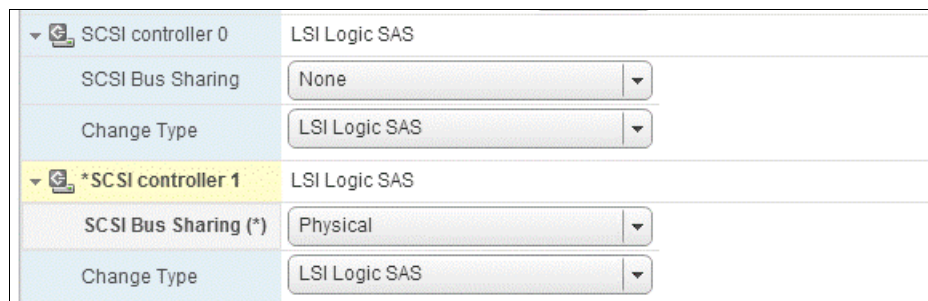


Figure 11-20 Add disks

- Click the New Device drop-down menu, select **RDM Disk**, and click **Add**, as shown in Figure 11-21.




Figure 11-21 Select RDM

- From the list, select an unformatted LUN that will be used as the witness disk for creating the WSFC, as shown in Figure 11-22.

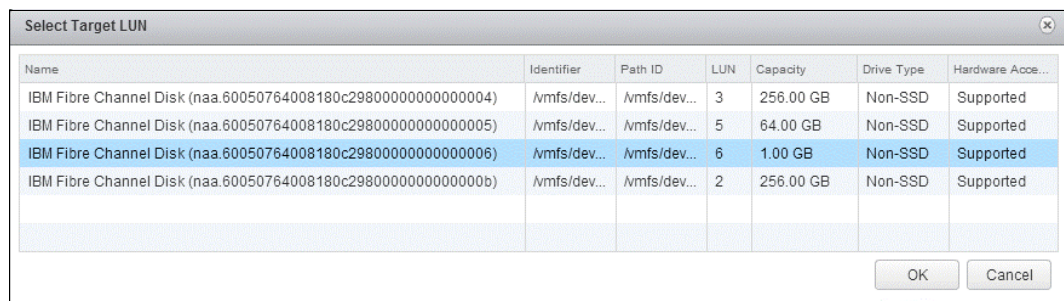


Figure 11-22 Select the failover cluster witness disk

7. Repeat steps 5 on page 188 and 6 on page 188 to add the other unformatted LUNs. The system creates RDM disks that map your VM to the target LUN. The RDM disk is shown in the list of virtual devices as a new hard disk, as shown in Figure 11-23.

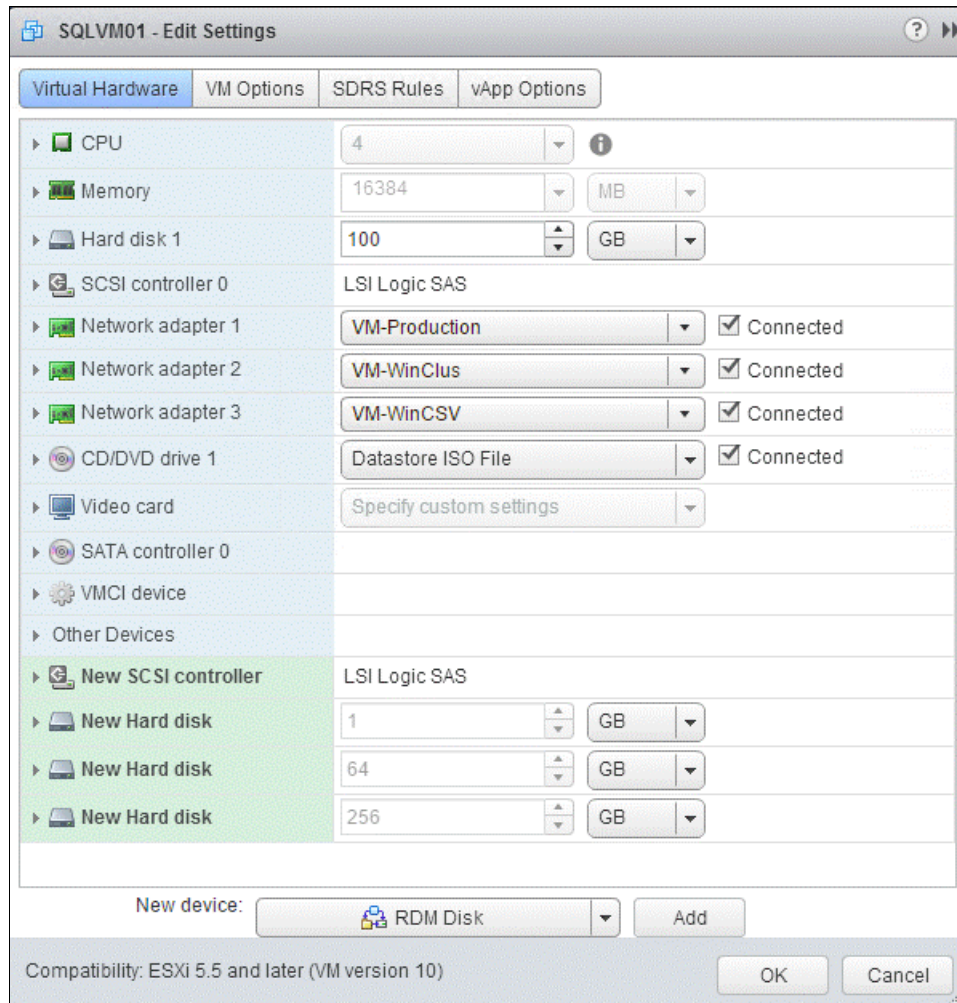


Figure 11-23 New disk

8. Click the arrow next to **New Hard disk** and select the following settings:
 - a. Location: **Store with VM**
 - b. Compatibility Mode: **Physical**
 - c. Virtual Device Node: **SCSI(1:0)**
 Use the created SCSI 1 controller because you cannot use SCSI 0.
9. Click **OK**.

10. Repeat steps 8 on page 189 and 9 on page 189 for the other two new hard disks; for the one for Virtual Device mode, select **SCSI(1:1)** for the second new disk and **SCSI(1:2)** for the third new disk, as shown in Figure 11-24.

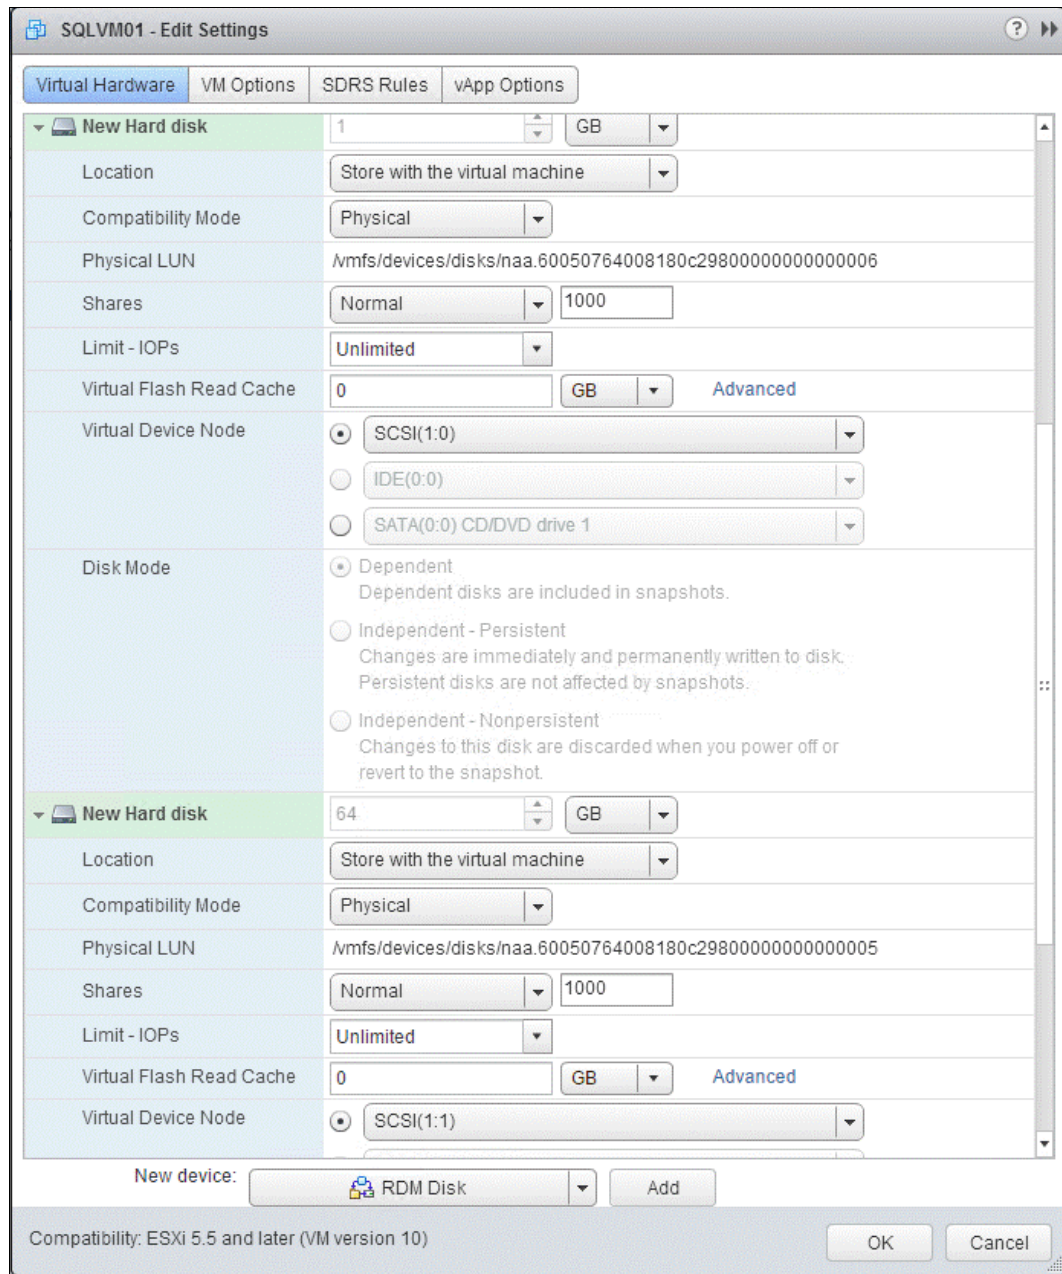


Figure 11-24 Select SCSI

11.3.6 Adding Raw Device Mapping disks to the second virtual machine node

To allow shared access to clustered services and data, point the witness disk of the second node to the same location as the first node's witness disk. Point the additional shared storage disks to the same location as the first node's shared storage disks.

Complete the following steps:

1. In the vSphere Web Client navigator, select the created second VM, right-click it, and select **Edit settings**.
2. On the Customize hardware window, click the **Virtual Hardware** tab.
3. Click the **New Device** drop-down menu, select **SCSI Controller**, and click **Add**.
4. Make sure that the SCSI Controller type is LSI Logic SAS, and that the SCSI Bus Sharing is Physical, as shown in Figure 11-25.

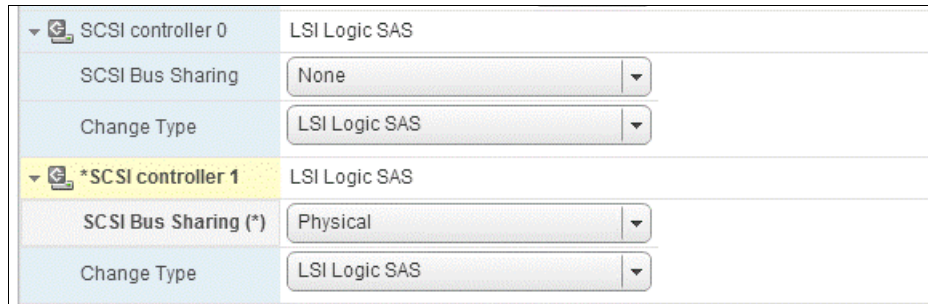


Figure 11-25 Select LSI Logic SAS

5. Click the **New Device** drop-down menu, select **Existing Hard Disk**, and click **Add**, as shown in Figure 11-26.

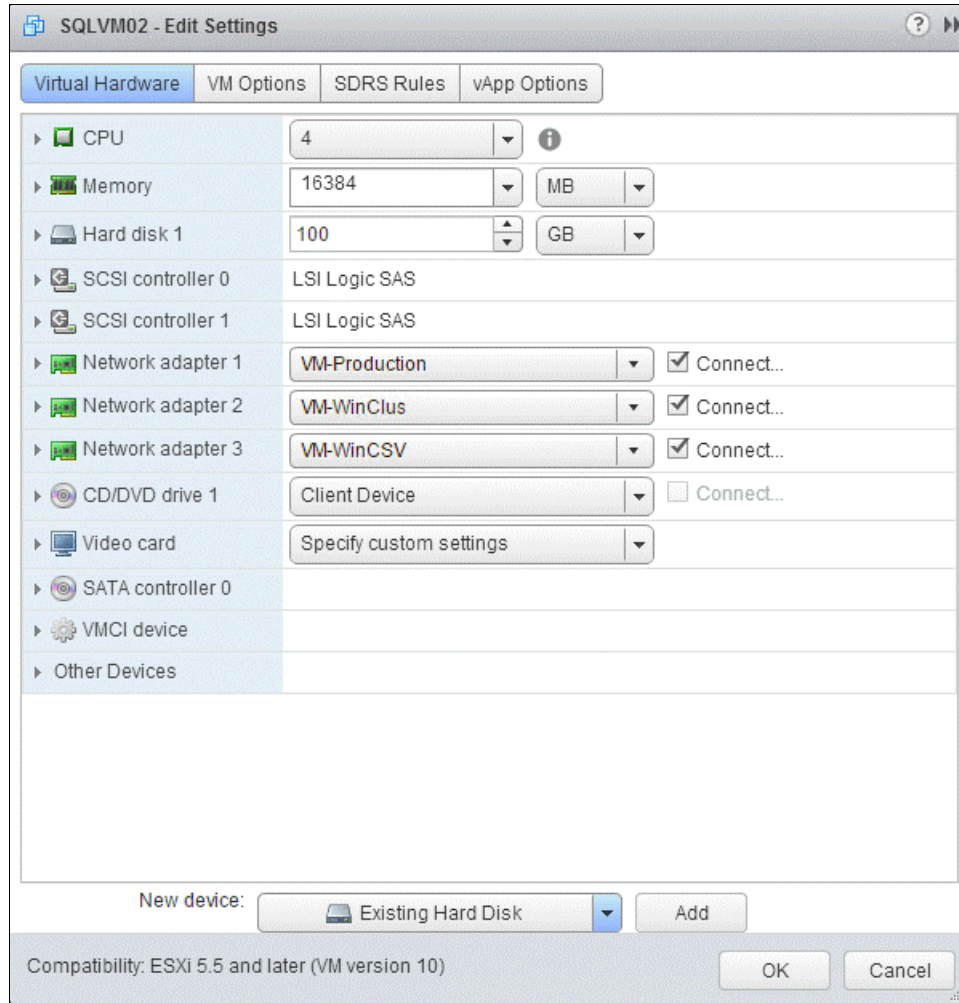


Figure 11-26 Existing hard disk

- In Disk File Path, browse to the location of the witness disk that is specified for the first node, as shown in Figure 11-27.

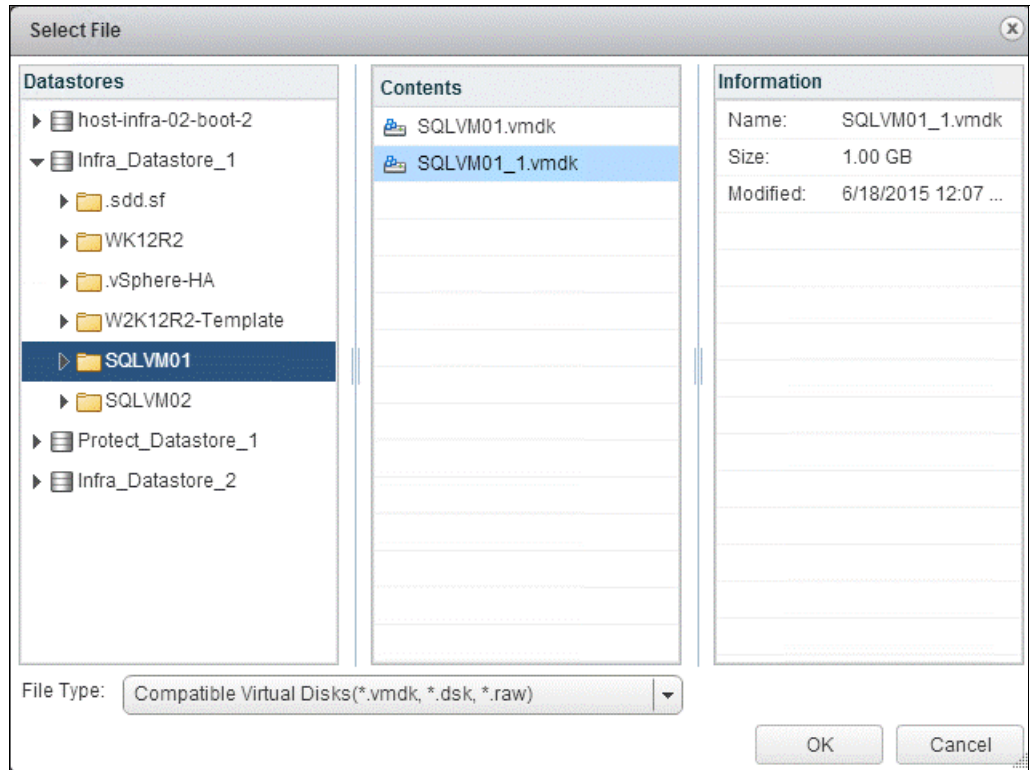


Figure 11-27 Browse to location

7. Select the same **SCSI(1:0)** virtual device node that was selected for the first VM's shared storage disk and click **OK**.

The virtual device node settings for this VM's shared storage must match the corresponding virtual device node for the first VM, as shown in Figure 11-28.

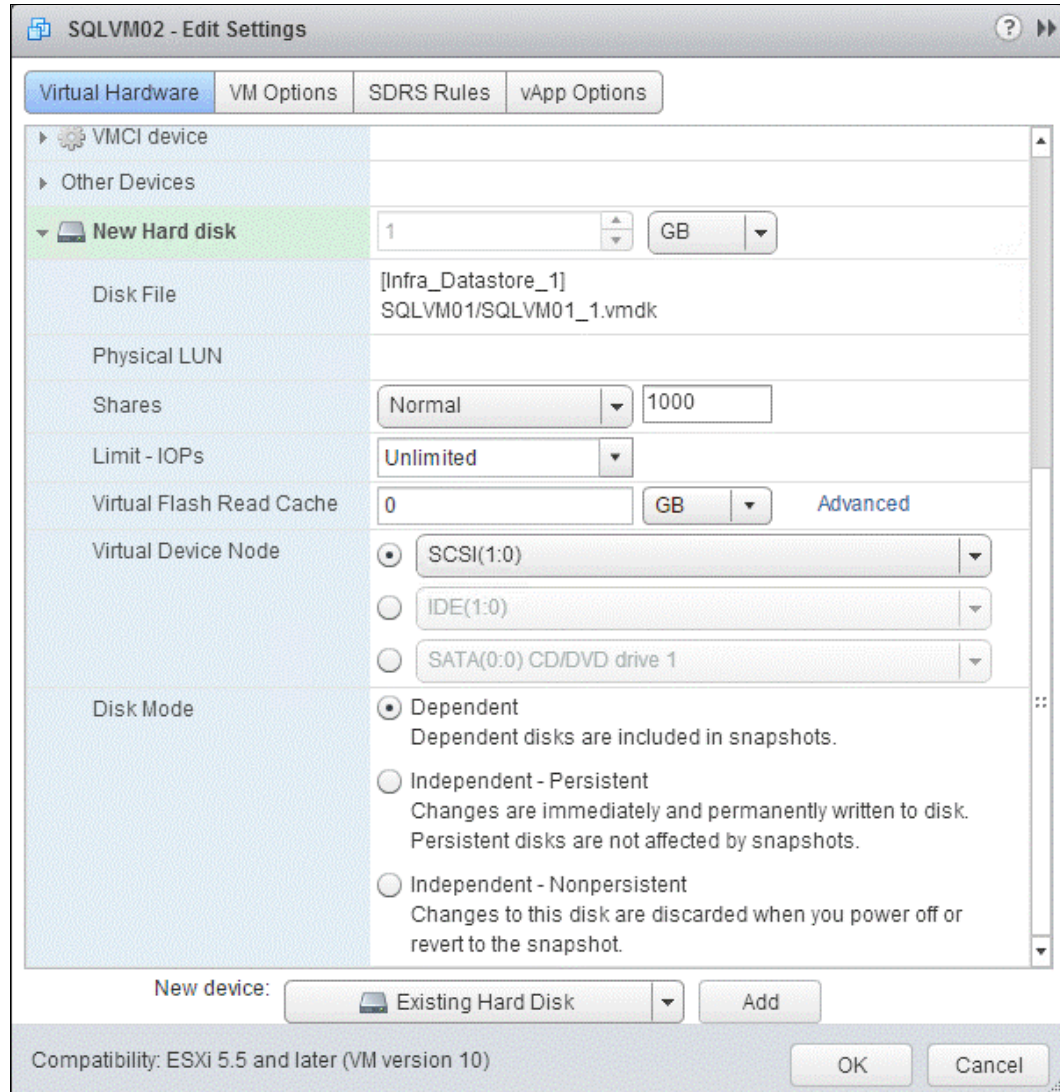


Figure 11-28 New hard disk

8. Repeat steps 5 on page 192 - 7 to add the remaining RDM disks.

11.3.7 Preparing the disks for cluster use

Complete the following steps:

1. Log in to the first VM node.
2. From Server Manager, click **File and Storage Services** → **Volumes** → **Disks**
3. Select an offline disk, right-click it, and select **Bring Online**, as shown in Figure 11-29 on page 195.

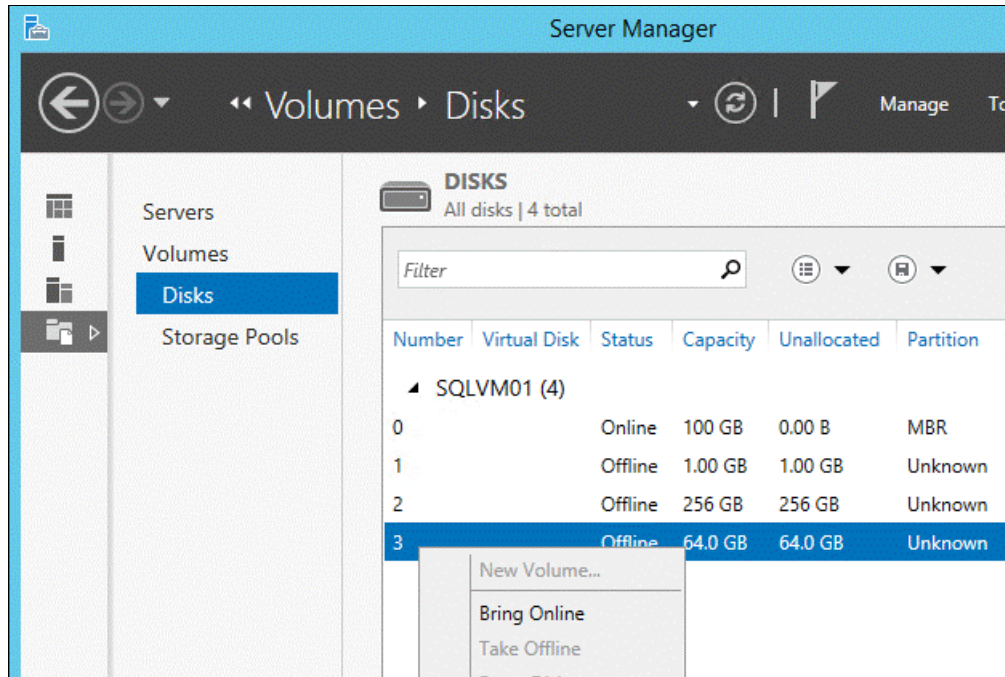


Figure 11-29 Bring disks online

- After the disk comes online, right-click it and select **Initialize**, as shown in Figure 11-30.

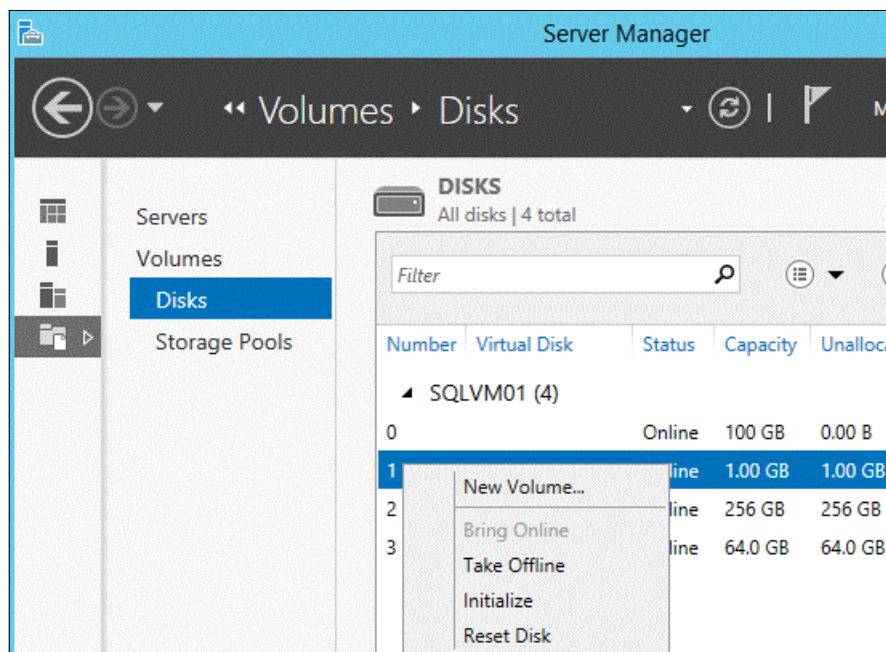


Figure 11-30 Initialize disks

- Repeat steps 3 on page 194 and 4 to bring online the other disks and initialize them.
- Right-click the 1 GB disk and create a simple volume for the witness disk by using the default settings.

7. Right-click the other shared disks that are used later by SQL Server to create a simple volume by using the settings that are shown in Figure 11-31. These disks are used later for the SQL Server FCI.

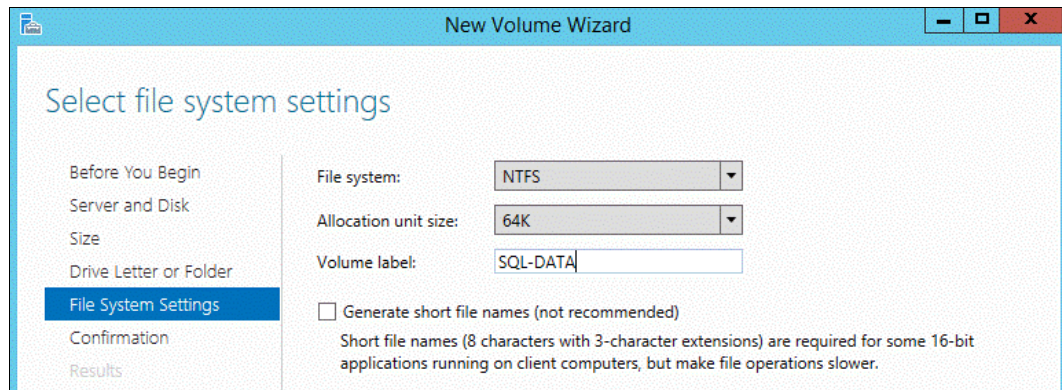


Figure 11-31 Select settings

8. Log in to the second VM node and bring only the disks online because they are shared drives that are already initialized and formatted on the first VM node.

11.4 Installing a Windows Server Failover Cluster

This section provides detailed instructions about how to set up a two-node WSFC on the VMs. This section focuses on validating and setting up failover cluster on VMs. After the completion of this task, a DB2 10.5 HA instance can be installed.

Complete the following steps:

1. Click **Server Manager** → **Tools** and select **Failover Cluster Manager**.
2. In the Failover Cluster Manager window, click **Validate Configuration** under the Management section, as shown in Figure 11-32 on page 197.

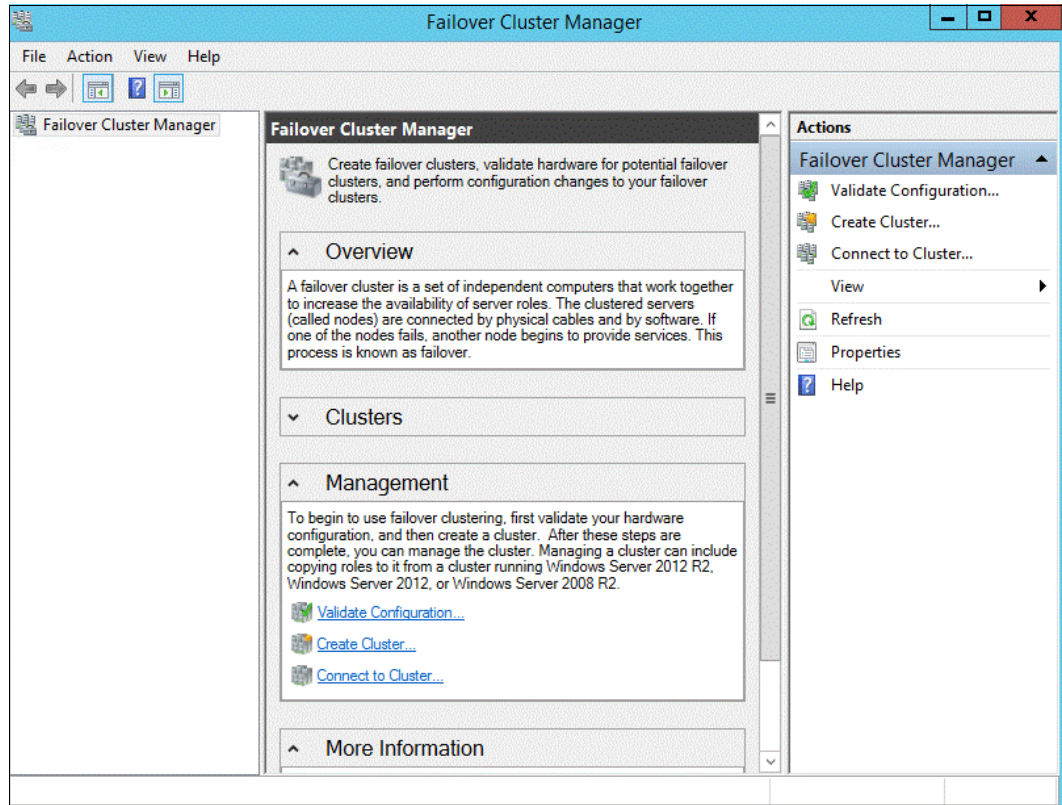


Figure 11-32 Validate configuration

3. In the Before You Begin window, click **Next**.

4. Enter the host names of the nodes or browse and select them and click **Next**, as shown in Figure 11-33.

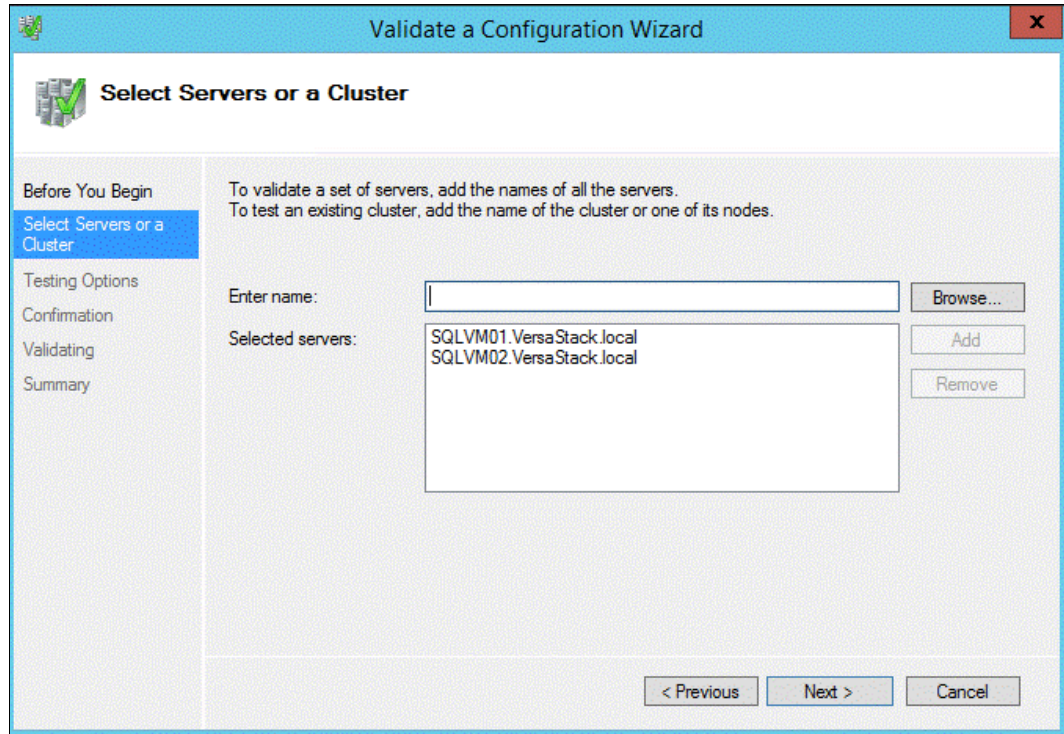


Figure 11-33 Enter host names

5. In the confirmation window, click **Next** to start the validation.
6. After the validation process is complete, review the report and fix any errors, as shown in Figure 11-34 on page 199.

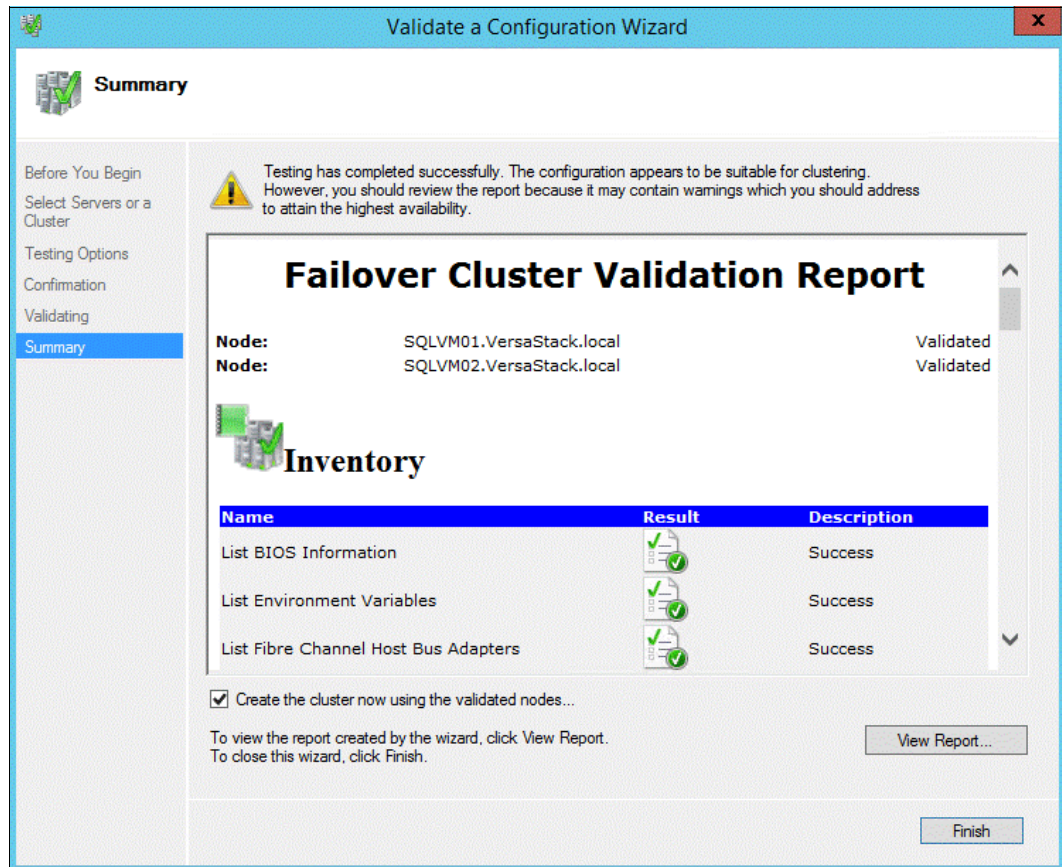


Figure 11-34 Validation report

7. If the validation is successful without any issues, select the **Create the cluster using the validated nodes** check box and click **Finish**.

- Enter a cluster name and IP address for the cluster and click **Next**, as shown in Figure 11-35.

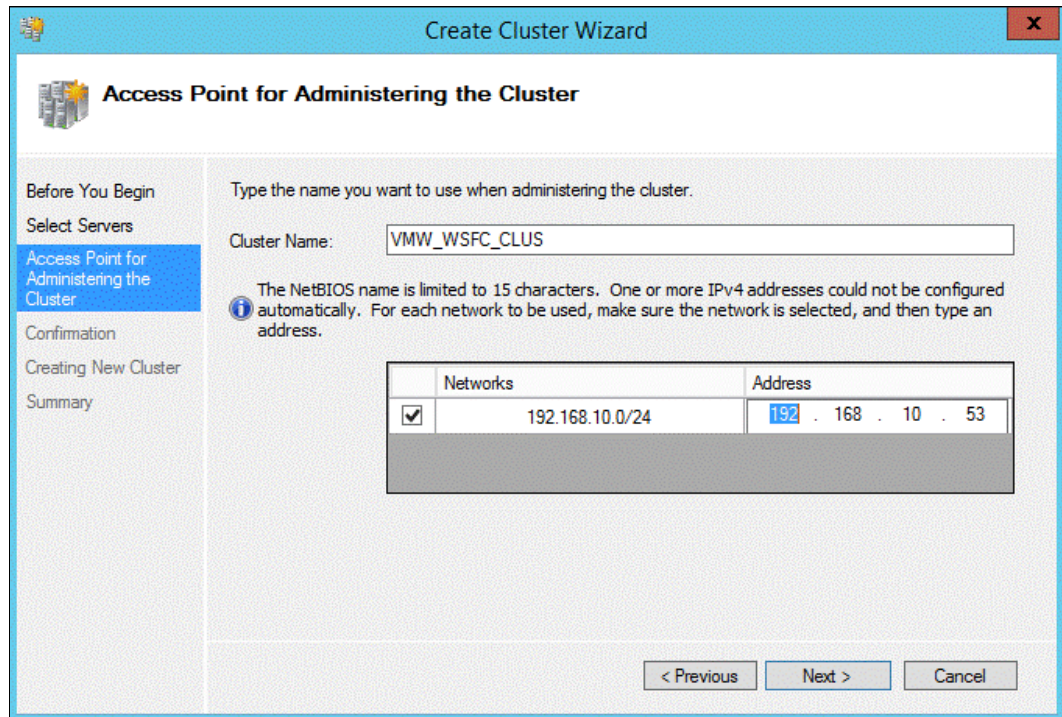


Figure 11-35 Create cluster

- Review the settings in the Confirmation window, select the **Add all eligible storage to the cluster** check box, and click **Next**, as shown in Figure 11-36.

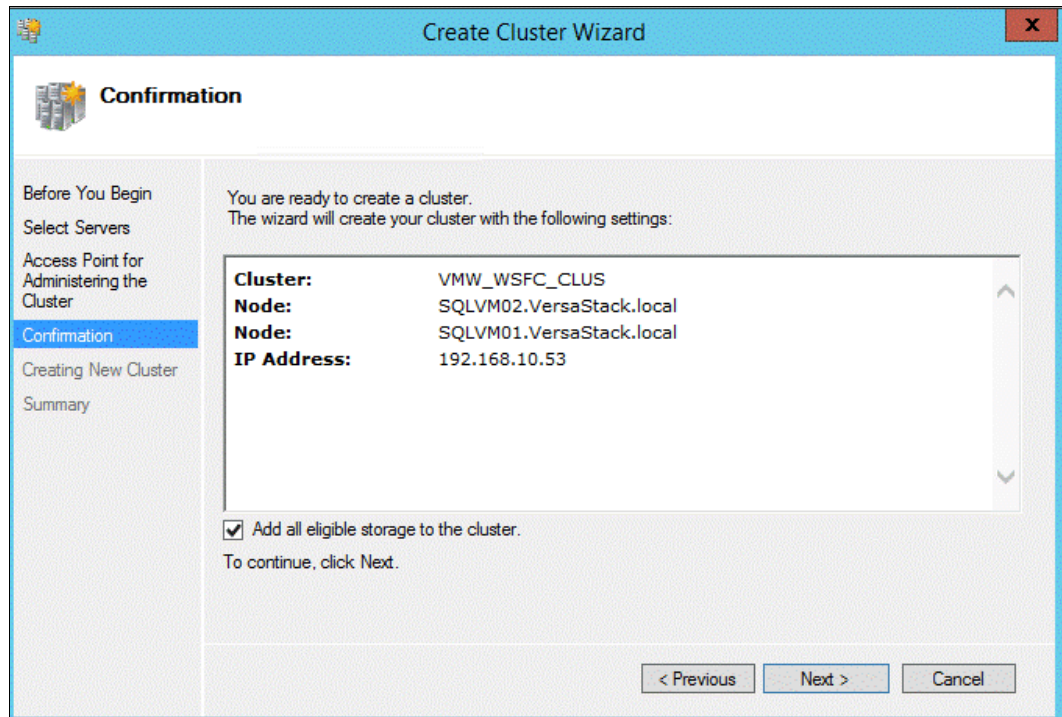


Figure 11-36 Confirmation

10. After the cluster is created, click **Finish** in the summary window, as shown in Figure 11-37.

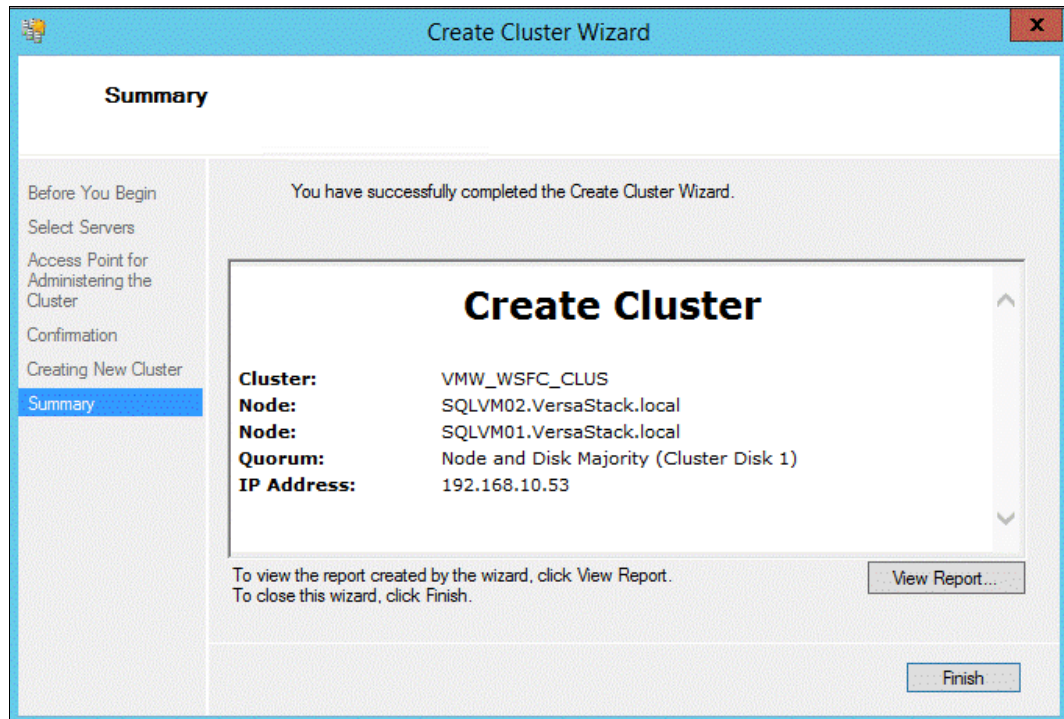


Figure 11-37 Cluster created

11. In the Failover Cluster Manager window, verify that the statuses of Cluster Core Resources, Network, and Storage are all online, as shown in Figure 11-38.

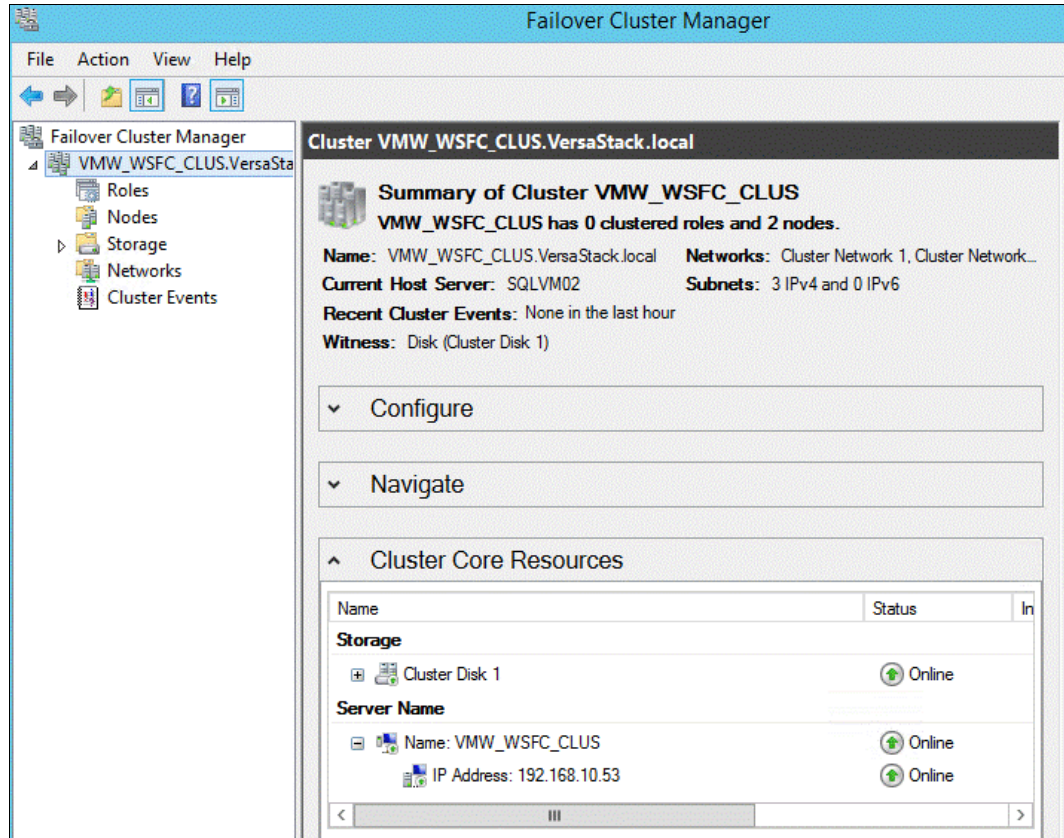


Figure 11-38 Verify status

12. Right-click only those Cluster Disks that will be used by SQL Server and select **Add to Cluster Shared Volumes**, as shown in Figure 11-39. Do not add the witness disk to the Cluster Shared Volumes.

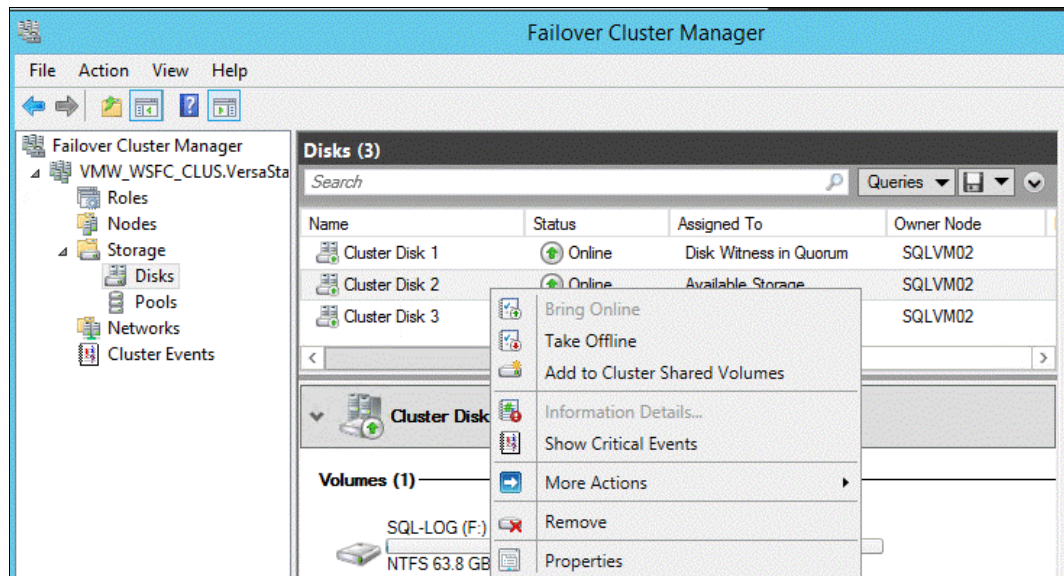


Figure 11-39 Add to cluster

13. Verify that the Cluster Shared Volumes status is online, as shown in Figure 11-40.

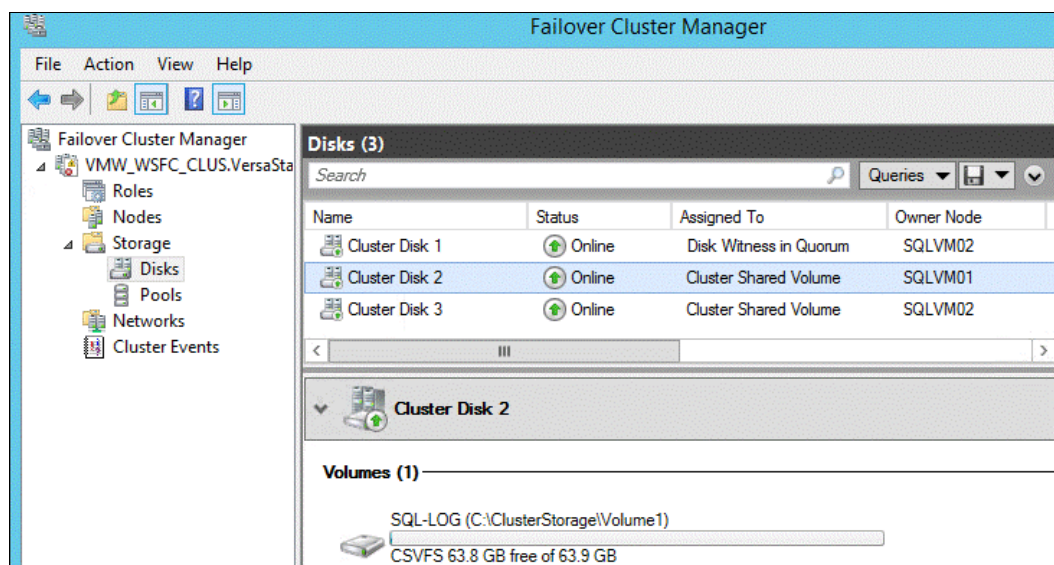


Figure 11-40 Verify the online status

11.4.1 Installing a DB2 Server Failover Cluster

The steps in 11.4, “Installing a Windows Server Failover Cluster” on page 196 describe how to deploy a Microsoft Windows Failover Cluster in virtual environment by using physical Raw Device Mapping (pRDM) volumes to store the database data.

With this setup configured, you can now deploy DB2 on this cluster. Instructions about how to perform this deployment are outlined in detail in Chapter 4, “DB2 with Microsoft Windows Failover Cluster”, of *High Availability and Disaster Recovery Options for DB2 for Linux, UNIX, and Windows*, SG24-7363.

Information specific to DB2 10.5, which is the version that is used in the environment in this book, can be found at:

https://www.ibm.com/support/knowledgecenter/?lang=en#!/SSEPGG_10.5.0/com.ibm.db2.1uw.admin.ha.doc/doc/c0007402.html

However, using pRDM to store the database data reduces the mobility of your VMs within the VMware virtual infrastructure and also imposes limitations with regards to VMware vStorage API for Data Protection (VADP) assisted backup solutions, such as the IBM Spectrum Protect for Virtual Environments that is used in the DB2 on VersaStack environment.

So, in a case where your expected database payload does not require the use of pRDMs, you can opt to deploy a DB2 with regular virtual machine disks (VMDKs) and use the built-in High Availability (HA) feature of DB2 itself. This setup allows for transparent vMotion across your infrastructure and can be used with both IBM Spectrum Protect for Virtual Environment backups at the VM level, and with IBM Spectrum Protect Snapshot FlashCopy based backups to have restorability at both the VM and entire datastore levels.

Section 11.6, “Installing DB2” on page 213 describes how to deploy DB2 in such a VM, and 11.9, “Configuring DB2 High Availability” on page 248 covers the HA configuration by creating a near-synchronous active-passive configuration.

11.5 Modifying the vSphere High Availability and Distributed Resource Scheduler settings for the Windows Server Failover Cluster virtual machines

When using WSFC in a vSphere HA or Distributed Resource Scheduler (DRS) environment, you must configure the hosts and VMs to use certain settings. The following sections describe how to configure the use of WSFC in a vSphere HA and DRS environment:

- ▶ Creating anti-affinity rules
- ▶ Enabling strict enforcement of anti-affinity rules
- ▶ Setting the Distributed Resource Scheduler automation level for clustered virtual machines
- ▶ Using vSphere Distributed Resource Scheduler groups and VM-Host affinity rules with clustered virtual machines

11.5.1 Creating anti-affinity rules

When you cluster VMs across physical hosts in a vSphere environment, you should keep the VMs on different physical hosts. When you enable vSphere HA and DRS in an environment where VMs are clustered across physical hosts, there are situations where the VMs might run on the same host because of their capabilities. Therefore, to avoid situations where clustered VMs run on the same host, you must configure VM-VM anti-affinity rules by completing the following steps:

1. In the vSphere Web Client, go to the cluster.
2. Click the **Manage** tab.

3. Click **Settings** → **DRS Rules**, and click **Add**, as shown in Figure 11-41.

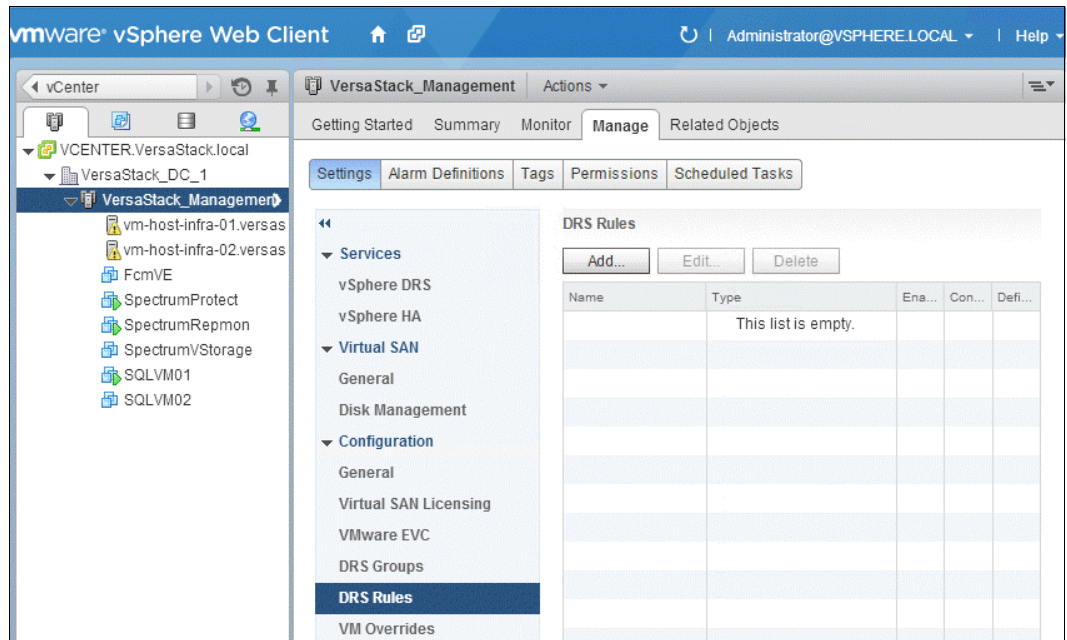


Figure 11-41 DRS Rules

4. Enter a name for the rule in the Rule dialog box, as shown in Figure 11-42.

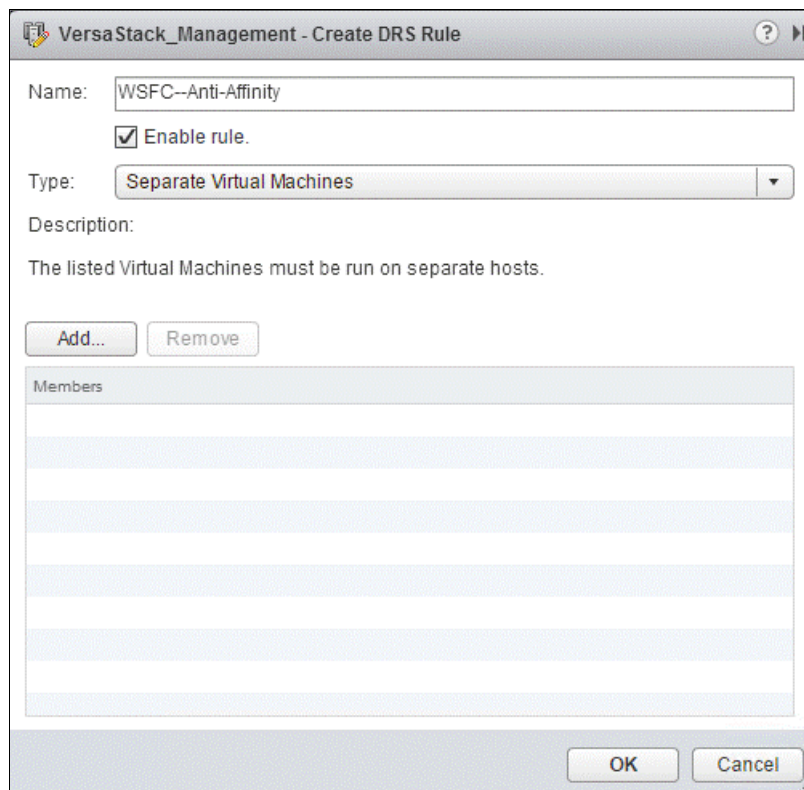


Figure 11-42 Create DRS Rule

5. From the Type drop-down menu, select the **Separate Virtual Machines** rule and click **Add**.
6. Select the two VMs to which the rule applies and click **OK** twice, as shown in Figure 11-43.

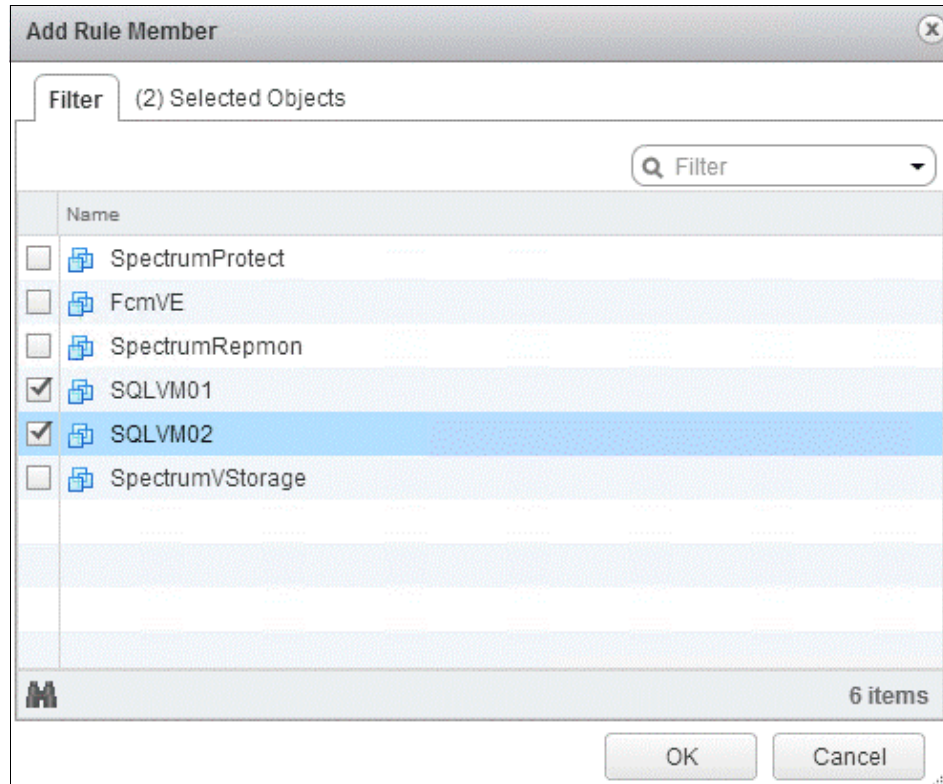


Figure 11-43 Add Rule Member

11.5.2 Enabling strict enforcement of anti-affinity rules

Setting the vSphere DRS advanced option “ForceAffinePoweron” to “1” enables strict enforcement of the anti-affinity rules. Complete the following steps:

1. In the vSphere Web Client, go to the cluster.
2. Click the **Manage** tab.
3. Click **Settings**, and under vSphere DRS, click **Edit**. The window that is shown in Figure 11-44 on page 207 opens.

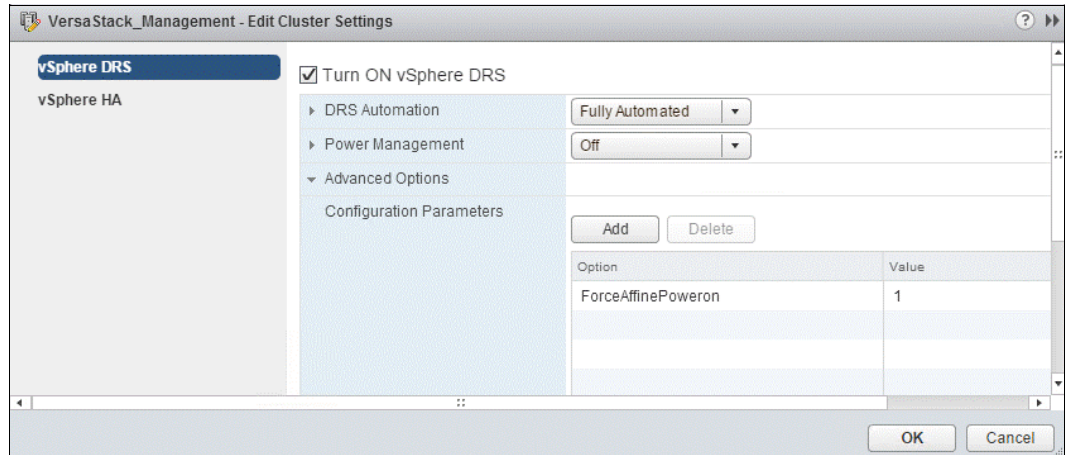


Figure 11-44 Set the Distributed Resource Scheduler options

4. Expand **Advanced Options** and click **Add**.
5. Enter “ForceAffinePoweron” into the Option column,
6. Enter “1” into the Value column and click **OK**.

11.5.3 Setting the Distributed Resource Scheduler automation level for clustered virtual machines

You must set the automation level of all VMs in a WSFC cluster to Partially Automated. Migration of WSFC clustered VMs is not recommended. Complete the following steps:

1. Browse to the cluster in the vSphere Web Client object navigator.
2. Click the **Manage** tab and click **Settings**.
3. Under Services, click **Edit**. The window that is shown in Figure 11-45 opens.

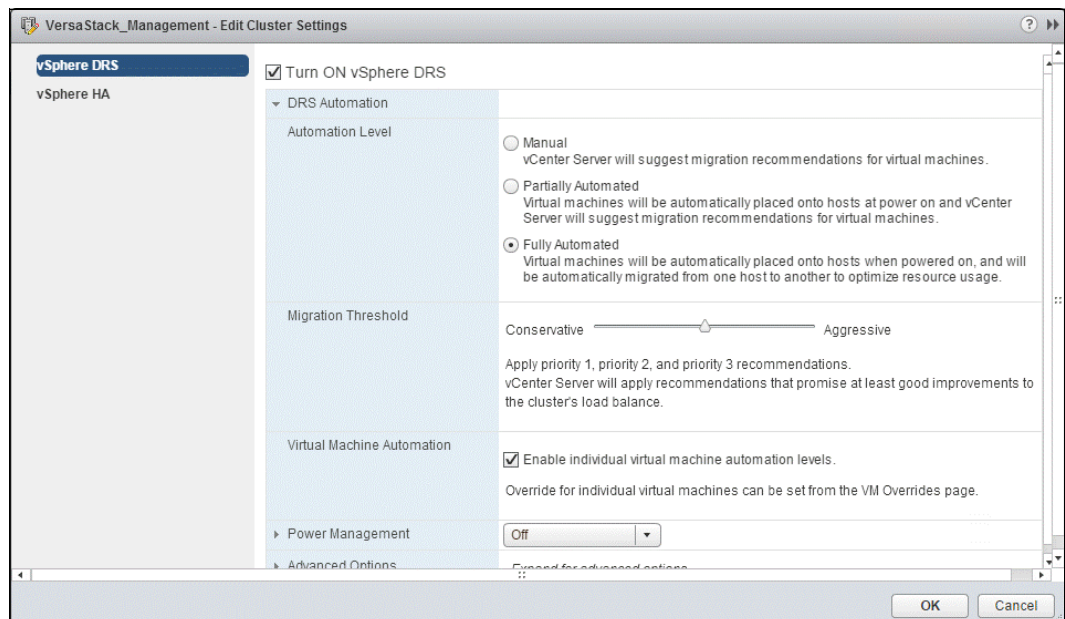


Figure 11-45 Automation level

- Expand **DRS Automation**, and under Virtual Machine Automation, select the **Enable individual virtual machine automation levels** check box and click **OK**.
- Under Configuration, select **VM Overrides** and click **Add**, as shown in Figure 11-46.

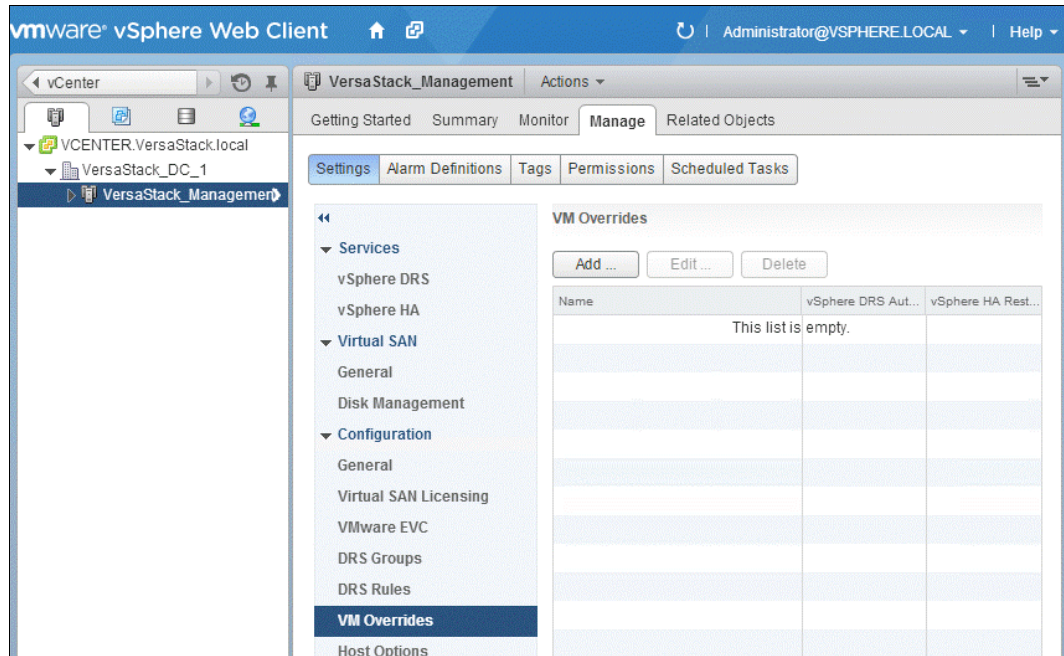


Figure 11-46 VM overrides

- Click the plus button, select the WSFC VMs in the cluster, and click **OK**.
- Click the **Automation level** drop-down menu, select **Partially Automated**, and click **OK**, as shown in Figure 11-47.

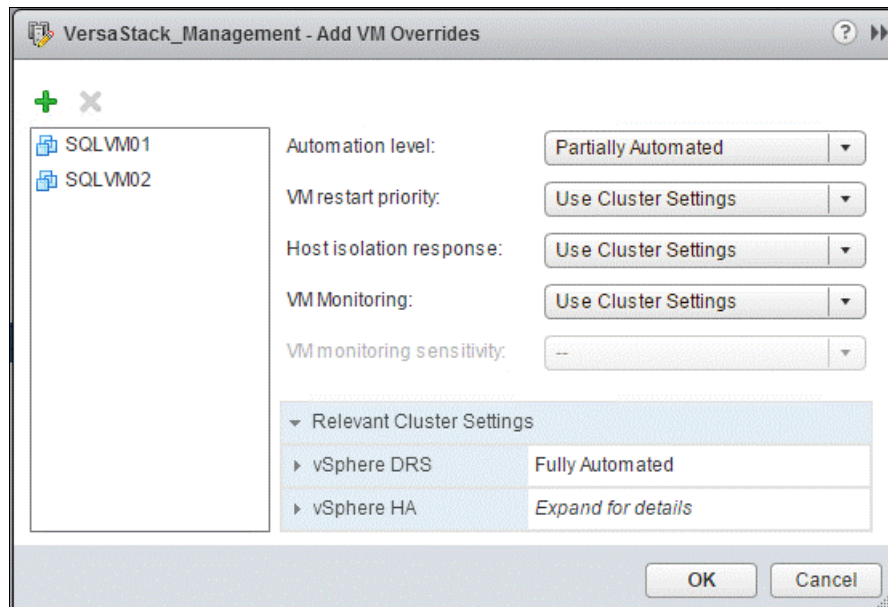


Figure 11-47 Choose settings

11.5.4 Using vSphere Distributed Resource Scheduler groups and VM-Host affinity rules with clustered virtual machines

In this section, you must complete the following tasks:

- ▶ Create two types of DRS groups by using the vSphere Web Client:
 - VM DRS groups containing at least one VM
 - Host DRS groups containing at least one host
- ▶ Set up VM-Host affinity rules for DRS groups (WSFC). A VM-Host anti-affinity rule establishes an anti-affinity relationship between a VM DRS group and a host DRS group.

Because vSphere HA does not obey VM-VM anti-affinity rules, it might put clustered VMs that are meant to stay apart on the same host. So, you also must create a VM-Host anti-affinity rule by setting up DRS groups and by using VM-Host anti-affinity rules, which are obeyed by vSphere HA.

For a cluster of VMs across physical hosts, each WSFC VM must be in a different VM DRS group, and linked to a different host DRS group with the affinity rule “Must run on hosts in group”.

Table 11-2 shows a configuration example where we created two VM DRS groups and two host DRS groups that are mapped as shown in the table.

Table 11-2 Configuration example

VM DRS group name	Member VM name	Mapped host DRS group name	Member host name
VMGroup_01	SQLVM01	HostGroup_01	vm-host-infra-01
VMGroup_02	SQLVM02	HostGroup_02	vm-host-infra-02

Creating a virtual machine DRS group (WSFC)

Before you can create a VM-Host affinity rule, you must create the VM DRS group and host DRS group to which the rule applies. This section describes how to create the VM DRS group. For information about creating the host DRS group, see “Creating a host DRS group (WSFC)” on page 211.

For a cluster of VMs across physical hosts, create one VM DRS group for each MSCS VM. For example, VMGroup_01 contains SQLVM01 and VMGroup_02 contains SQLVM02. Complete the following steps:

1. Browse to the cluster in the vSphere Web Client navigator.
2. Click the **Manage** tab.

3. Click **Settings**, click **DRS Groups**, and click **Add**, as shown in Figure 11-48.

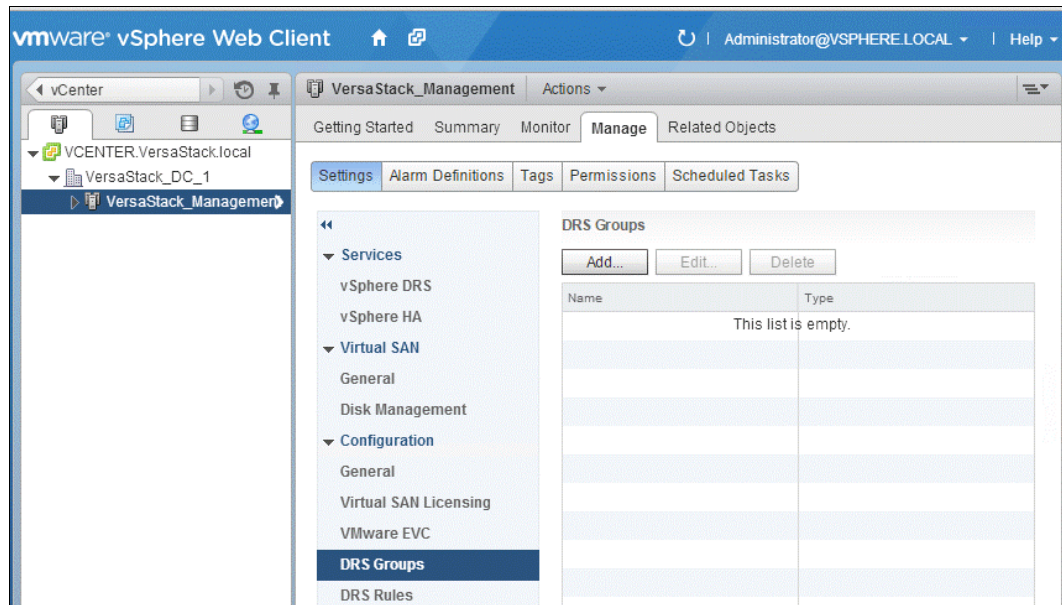


Figure 11-48 Add DRS groups

The window that is shown in Figure 11-49 opens.

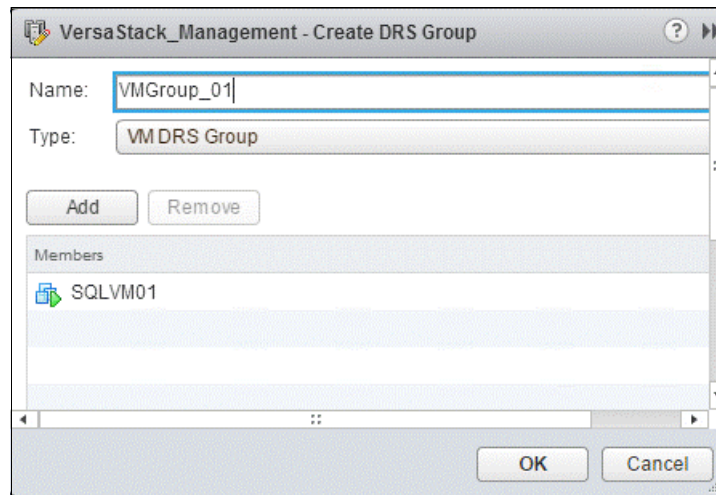


Figure 11-49 Create DRS Group

4. In the DRS Group dialog box, enter a name for the group.
5. Select **VM DRS Group** from the **Type** drop-down menu and click **Add**.
6. Select the check box next to a VM to add it. Continue this process until all the wanted VMs are added.
7. For a cluster of VMs across physical hosts, add one MSCS VM per group.
8. Click **OK**. Figure 11-50 on page 211 shows the completed process.

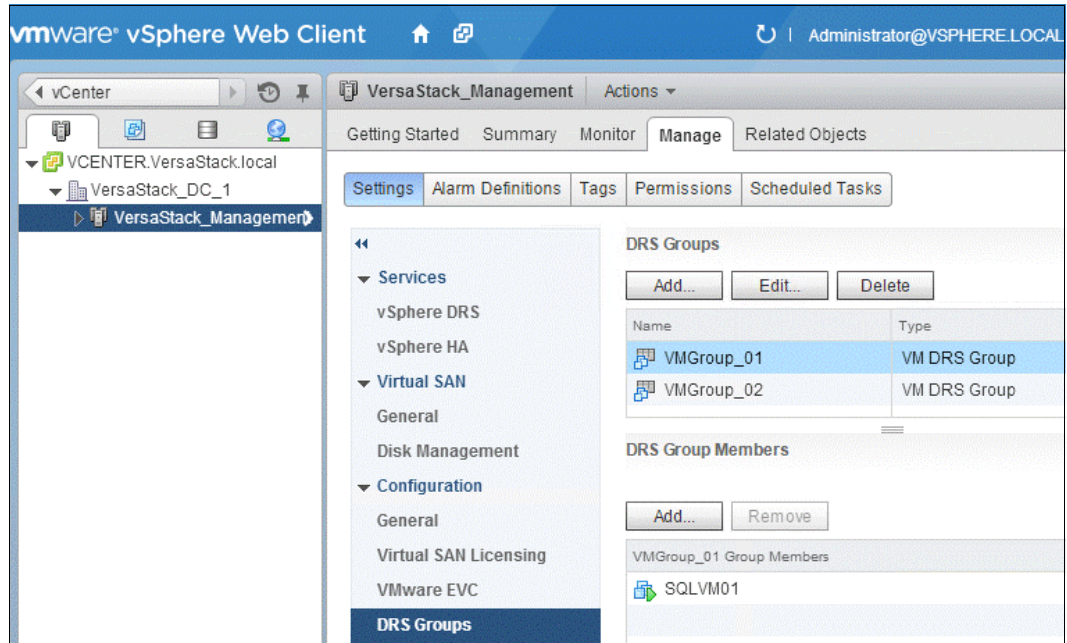


Figure 11-50 Complete

Creating a host DRS group (WSFC)

Before you can create a VM-Host affinity rule, you must create the host DRS group and the VM DRS group to which the rule applies. This section describes how to create the host DRS group. For information about creating the VM DRS group, see “Creating a virtual machine DRS group (WSFC)” on page 209

For a cluster of VMs across physical hosts, create groups with sets of hosts that do not overlap to ensure that the VMs that are placed in different host groups do not ever run on the same host simultaneously. Complete the following steps:

1. Browse to the cluster in the vSphere Web Client navigator.
2. Click the **Manage** tab.
3. Click **Settings**, click **DRS Groups**, and click **Add**.

4. In the DRS Group dialog box, enter a name for the group, as shown in Figure 11-51.

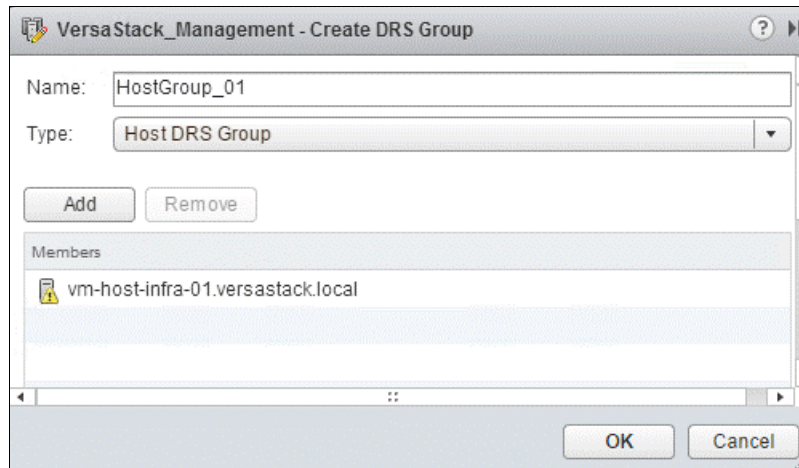


Figure 11-51 Create Host DRS group

5. Select **Host DRS Group** from the **Type** drop-down menu and click **Add**.
6. Click the check box next to a host to add it. Continue this process until all the wanted hosts are added.
7. Click **OK**, as shown in Figure 11-52.

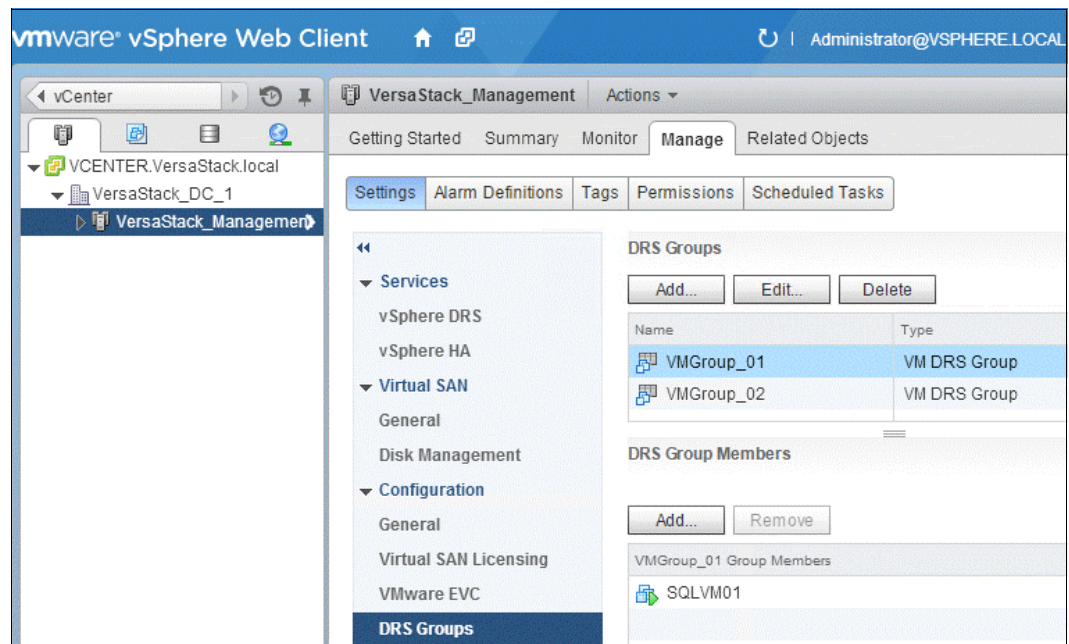


Figure 11-52 Add host group

Setting up the VM-Host affinity rules for DRS groups (WSFC)

You must create VM-Host affinity rules to specify whether the members of a selected VM DRS group can run on the members of a specific host DRS group. Complete the following steps:

1. Browse to the cluster in the vSphere Web Client navigator.
2. Click the **Manage** tab.

3. Click **Settings**, click **DRS Rules**, and click **Add**.
4. In the DRS Rule dialog box, enter a name for the rule.
5. From the Type menu, select **Virtual Machines to Hosts**.
6. Select the VM DRS group and the host DRS group to which the rule applies. For example, select **VMGroup_1** and **HostGroup_1**.
7. Select **Must run on hosts in group** and click **OK**, as shown in Figure 11-53.

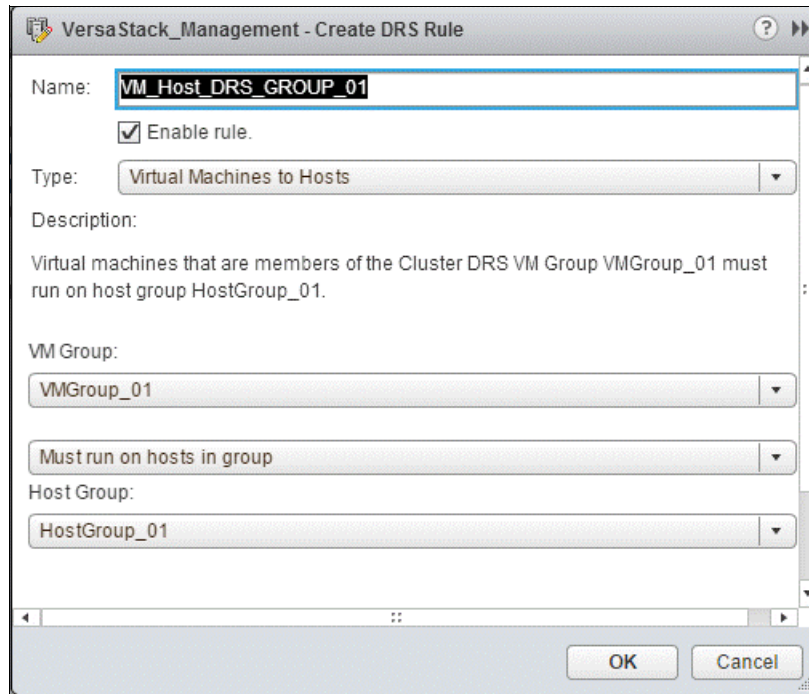


Figure 11-53 Select settings

The setup is now complete.

11.6 Installing DB2

Create a VM by following the instructions that are outlined in 11.1, “Creating virtual machines” on page 172 and deploy Windows Server 2012 R2 as described in 11.2, “Installing Windows Server 2012 R2” on page 179. Complete these steps twice to create the SQLVM01 and SQLVM02 VMs.

Because you want to deploy a small workload, deploy the base operating system, the DB2 database server, and a sample database on the same virtual disk. Naturally, separate virtual disks can be added to the VM to have separate locations for the database data and the database log files if needed.

In this example, we copy the following binary files to the source VMs:

- ▶ DB2 10.5 Fix Pack 6: v10.5fp6_ntx64_universal_fixpack.exe
- ▶ Data Studio Client V4.1.2: DATA_STUDIO_CLIENT_V4.1.2_FOR_WIN.zip
- ▶ GSDB sample database: GSDB_DB2_LUW_ZOS_v2r3.zip

You can obtain the DB2 and Data Studio Client through the IBM Passport Advantage® portal. The GSDB sample database can be downloaded from the following website:

https://www.ibm.com/support/knowledgecenter/#!/SS62YD_4.1.1/com.ibm.sampledata.go.doc/topics/download.html

To start the DB2 10.5 installation, complete the following steps:

1. Log on to the SQLVM01 VM, go to the folder where v10.5fp6_ntx64_universal_fixpack.exe is deployed, and extract the binary file as an administrator, as shown in Figure 11-54.

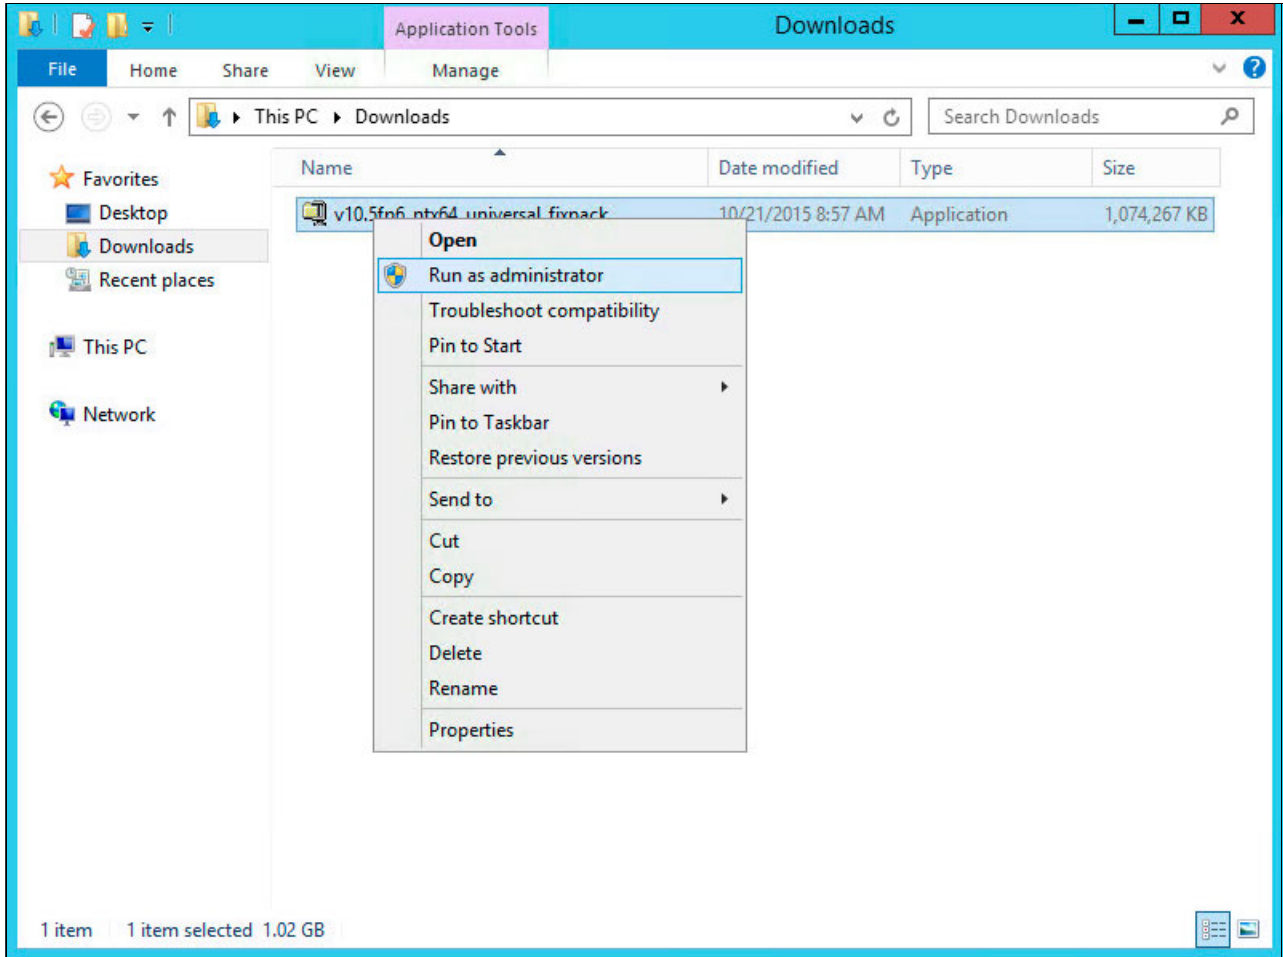


Figure 11-54 DB2 - extract the installation binary files

2. Specify a directory to which to extract the installation binary files, for example, c:\db2inst, and click **Unzip** to start the extraction, as shown in Figure 11-55 on page 215.

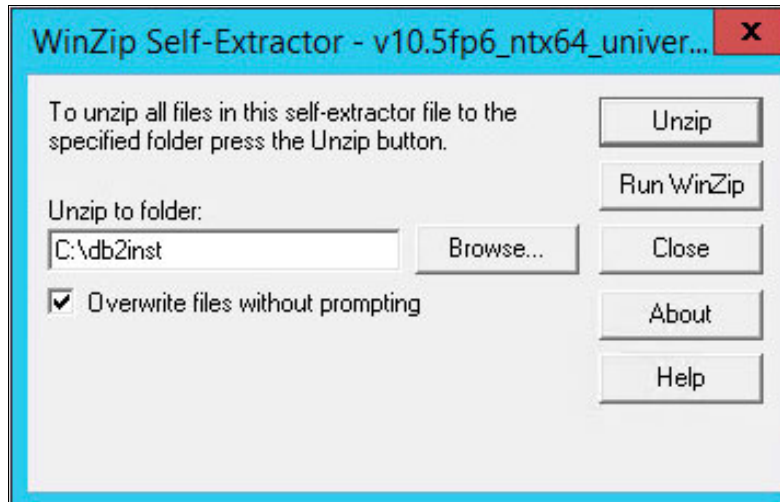


Figure 11-55 DB2 - define the location for the installation binary files

- Go to `c:\db2inst` in Windows Explorer, right-click `setup.exe`, and start the installation as an administrator, as shown in Figure 11-56.

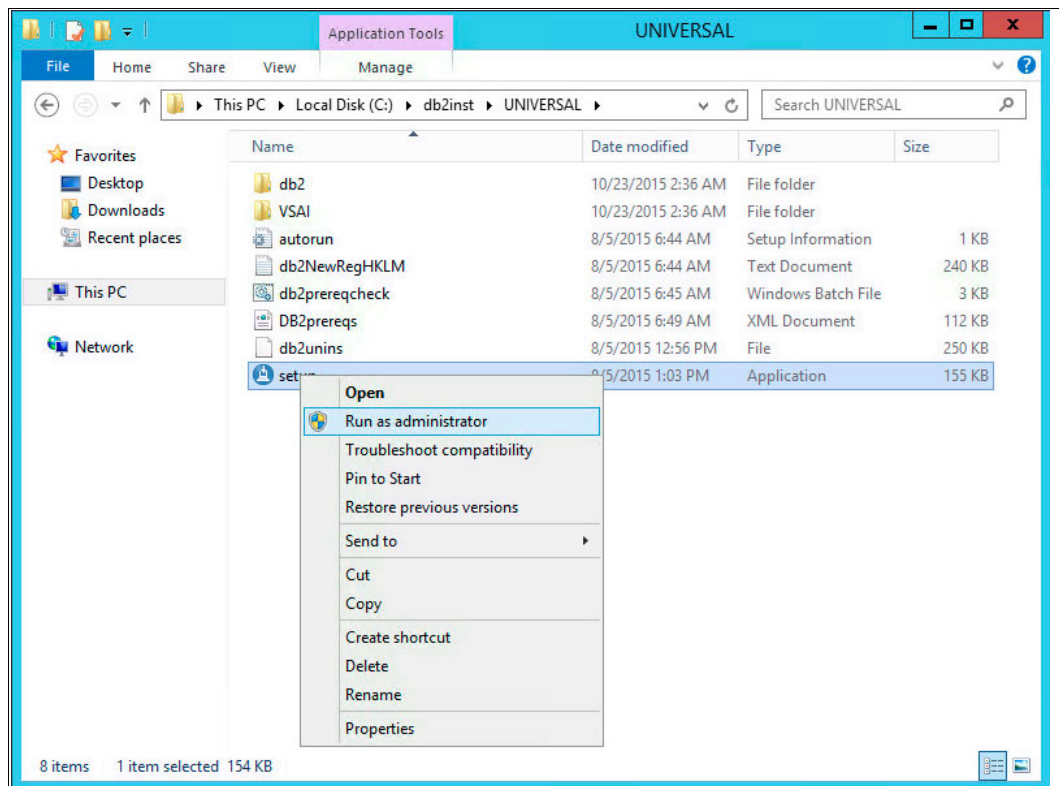


Figure 11-56 DB2 - run the setup program

The Welcome opens, as shown in Figure 11-57.

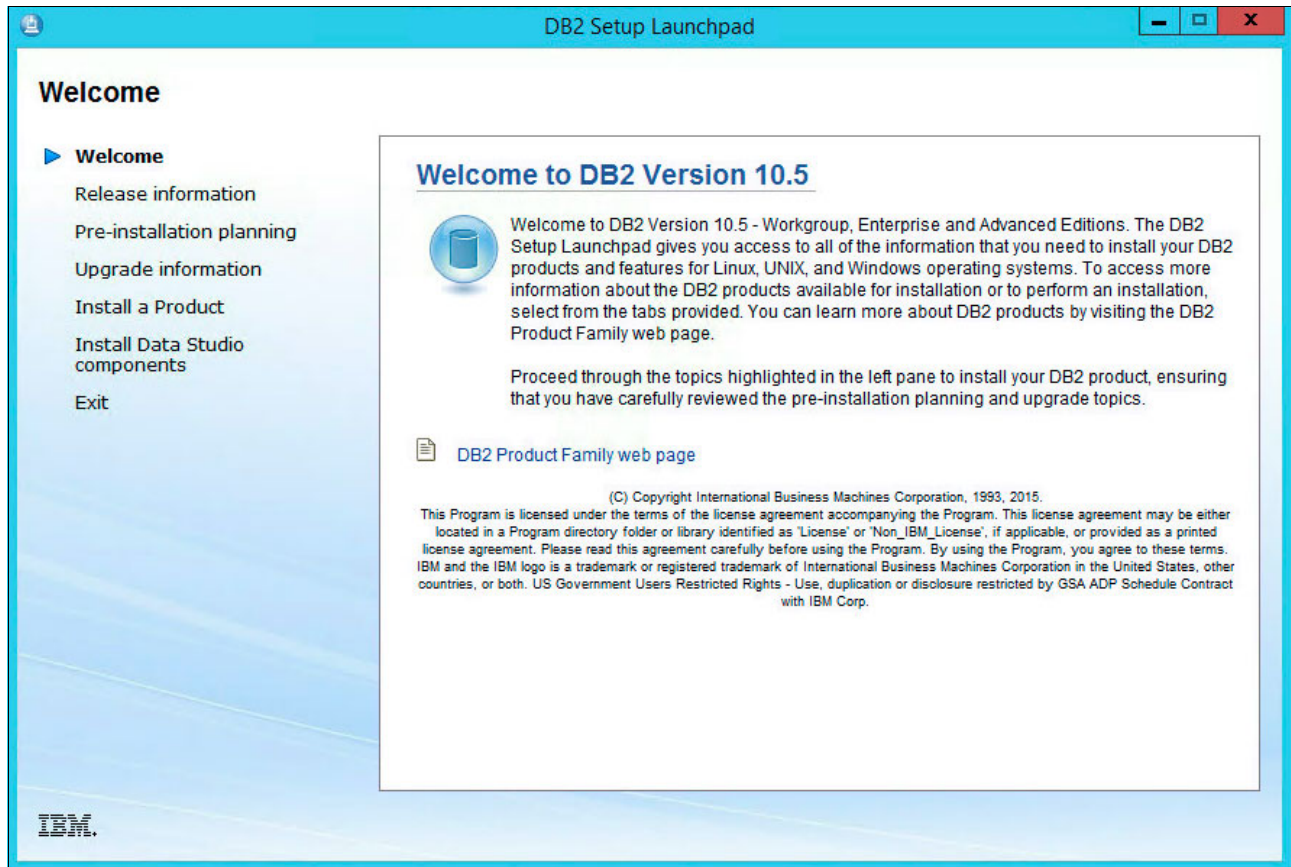


Figure 11-57 DB2 - Welcome to DB2 Version 10.5

4. In the left pane, click **Install a Product**, scroll down until you reach DB2 Enterprise Server Edition, and click **Install New**, as shown in Figure 11-58.

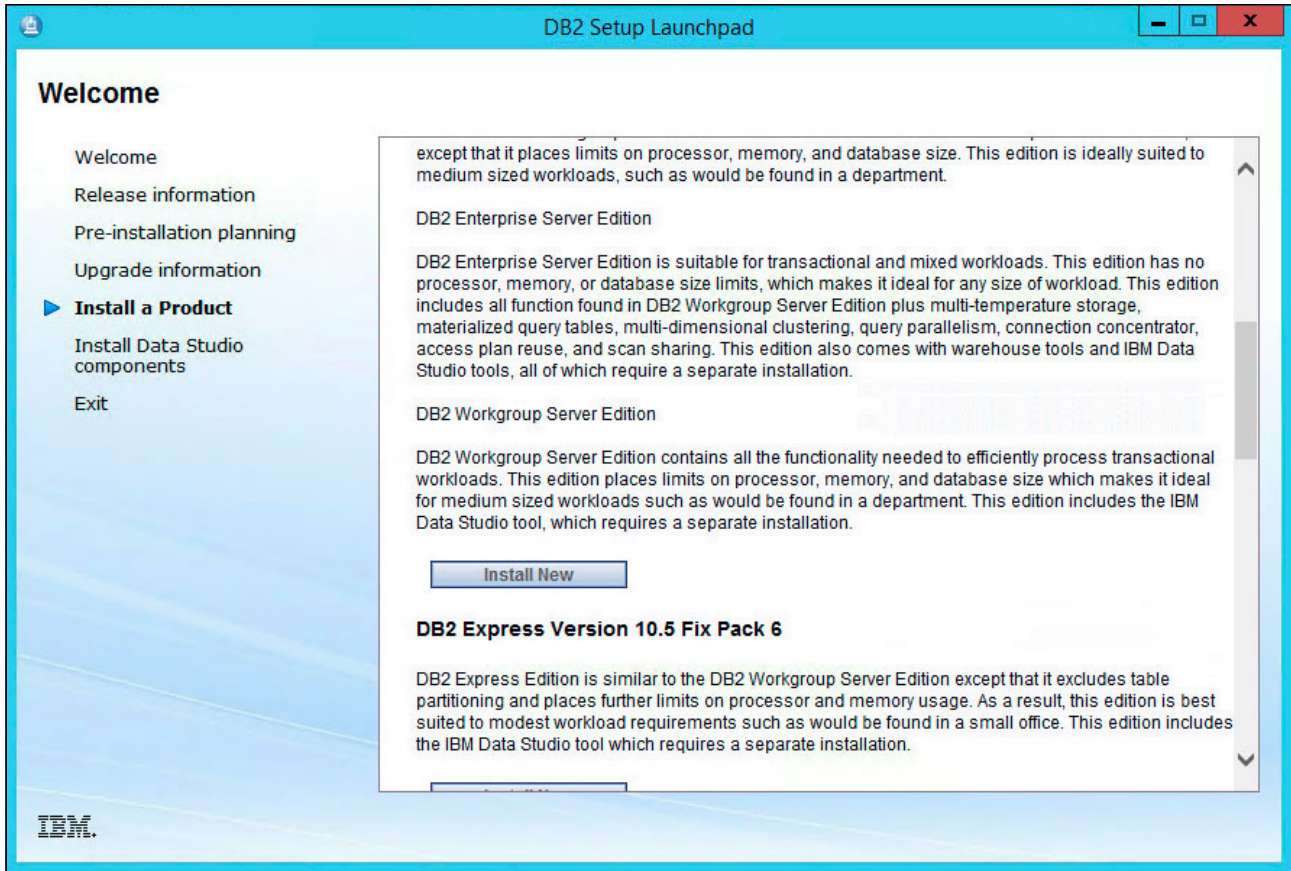


Figure 11-58 DB2 - install new DB2 Enterprise Server Edition

5. DB2 Setup - DB2 Server Edition starts. Click **Next** to continue, as shown in Figure 11-59.

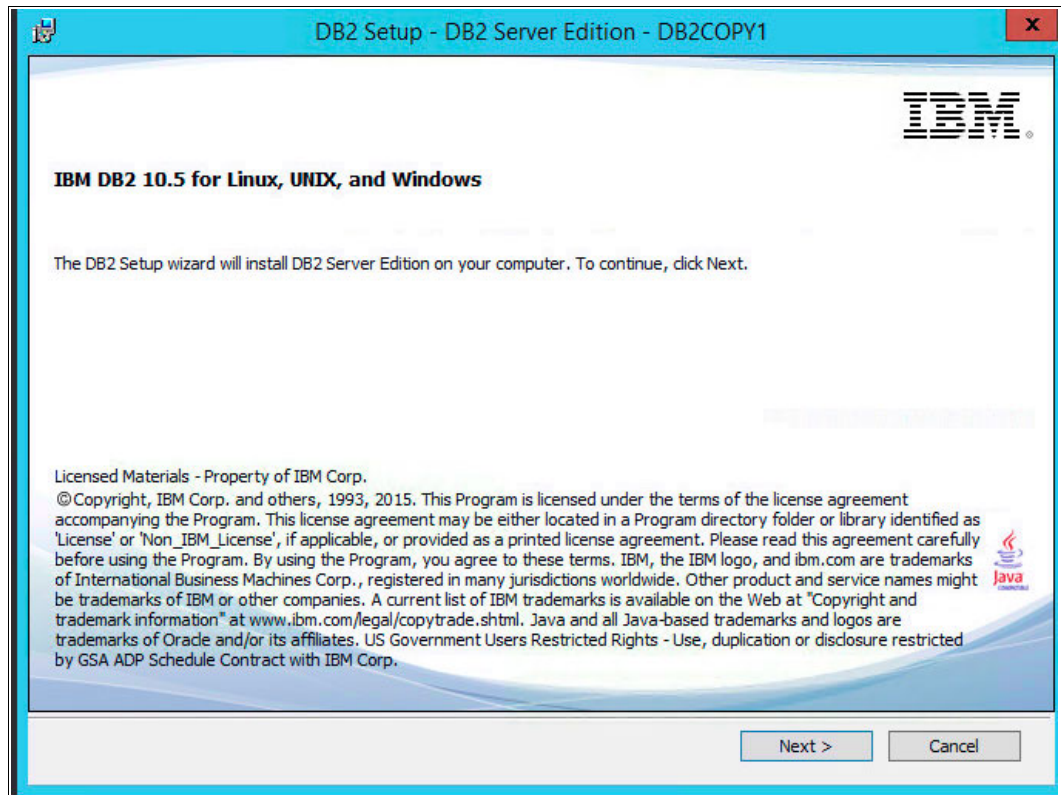


Figure 11-59 DB2 - DB2 Server Edition Setup

6. Select **I accept both the IBM and the non-IBM terms** and click **Next**, as shown in Figure 11-60.

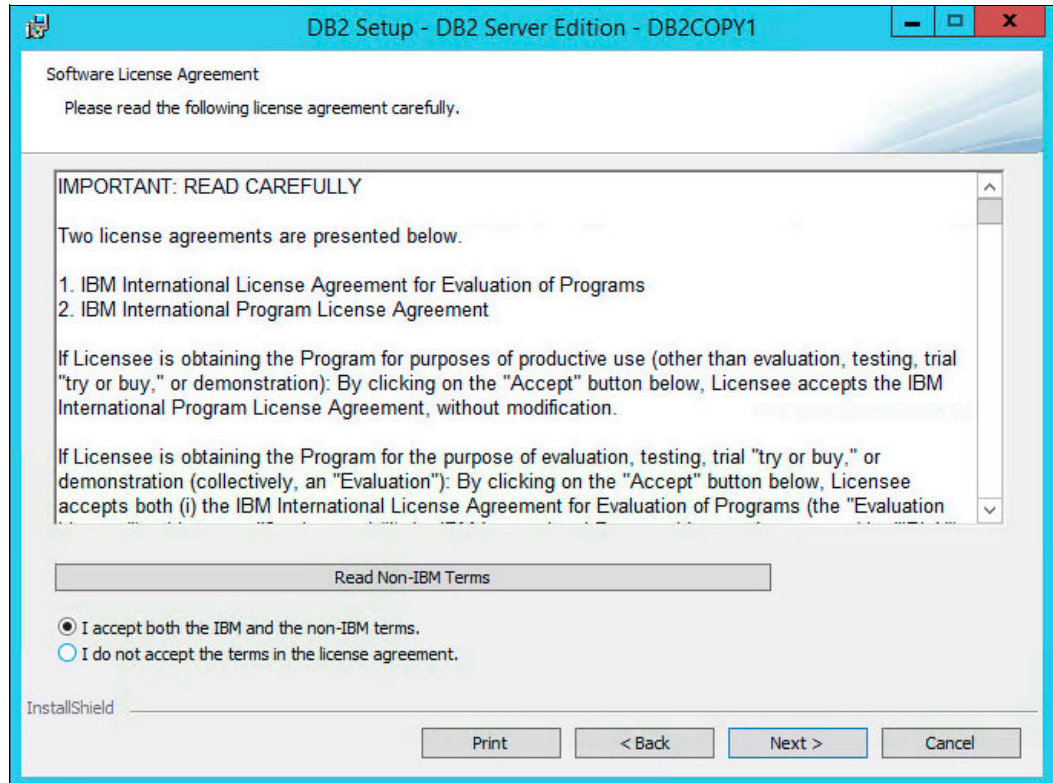


Figure 11-60 DB2 - accept IBM and non-IBM terms

7. Select **Typical** as the installation type and click **Next** to continue, as shown in Figure 11-61.

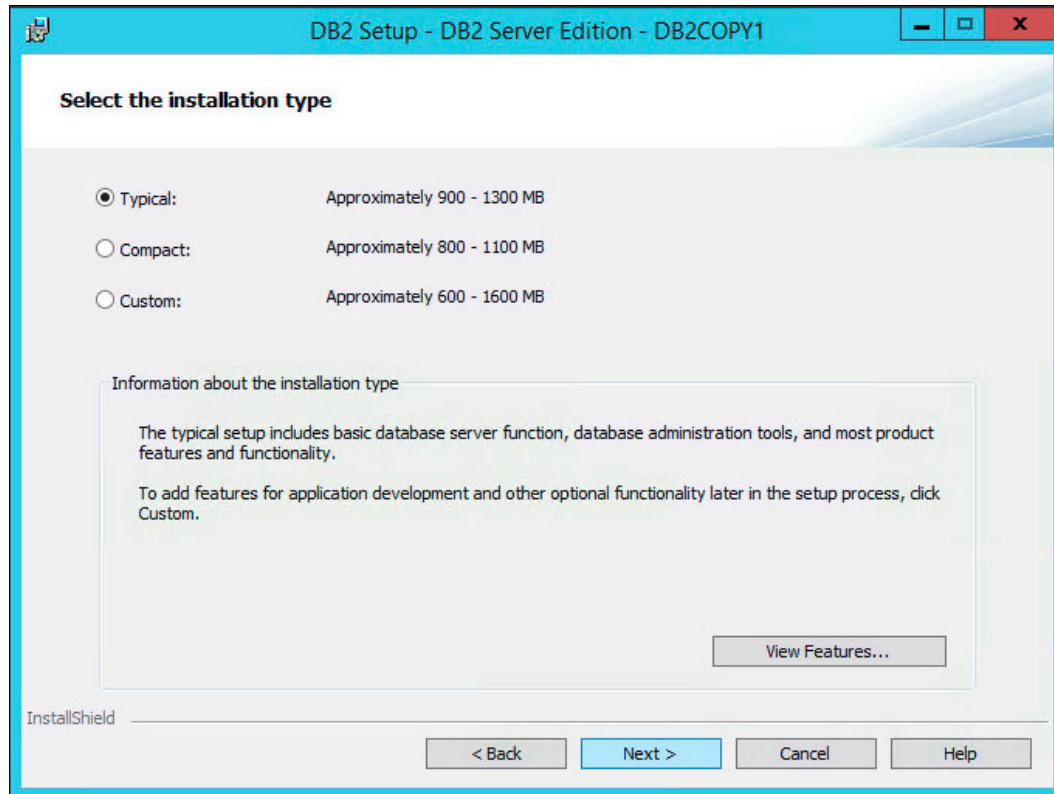


Figure 11-61 DB2 - select a Typical installation

8. Keep the default option **Install DB2 server Edition on the computer and save my settings in a response file** and click **Next** to continue, as shown in Figure 11-62.

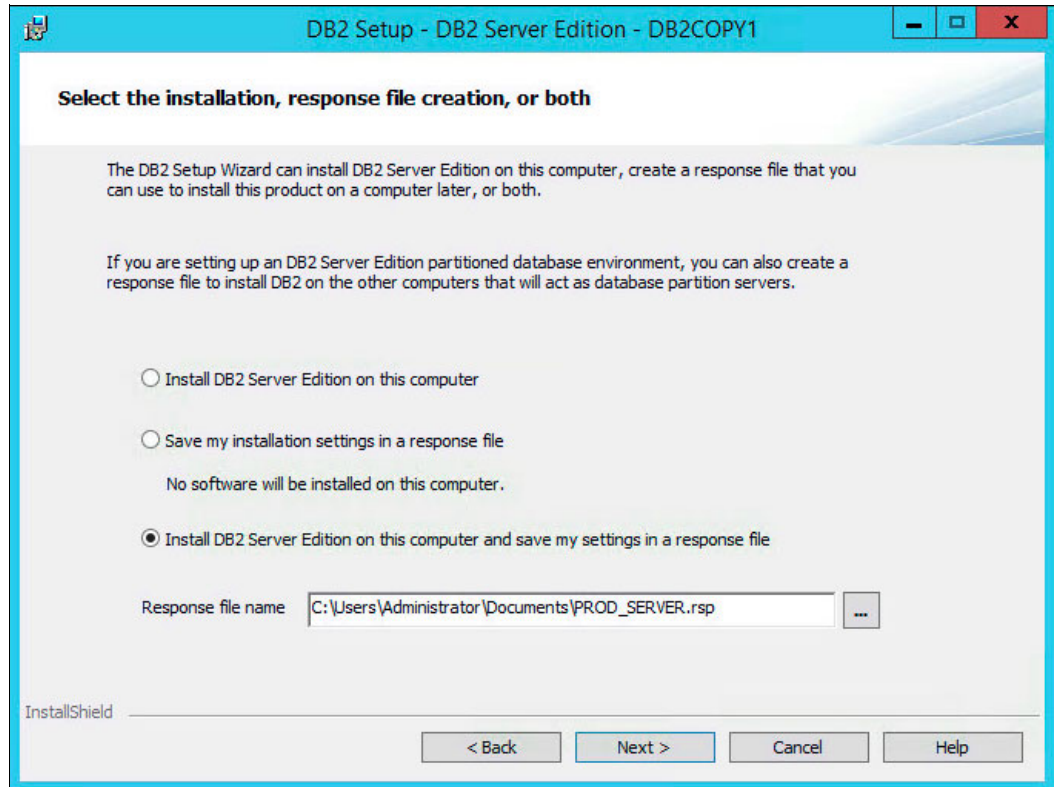


Figure 11-62 DB2 - install DB2 server and save the settings to a response file

9. Keep the default path to deploy the DB2 binary files and click **Next** to continue, as shown in Figure 11-63.

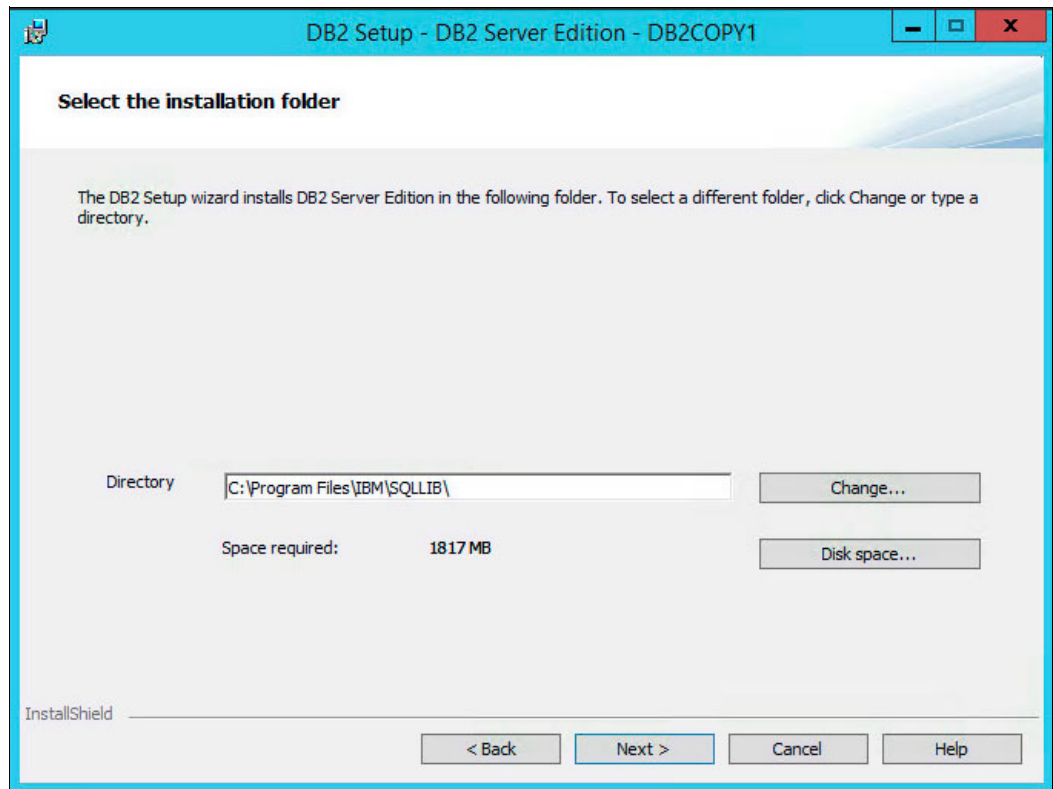


Figure 11-63 DB2 - installation folder

10. Keep the default location and startup settings for the IBM SSH server and click **Next** to continue, as shown in Figure 11-64.

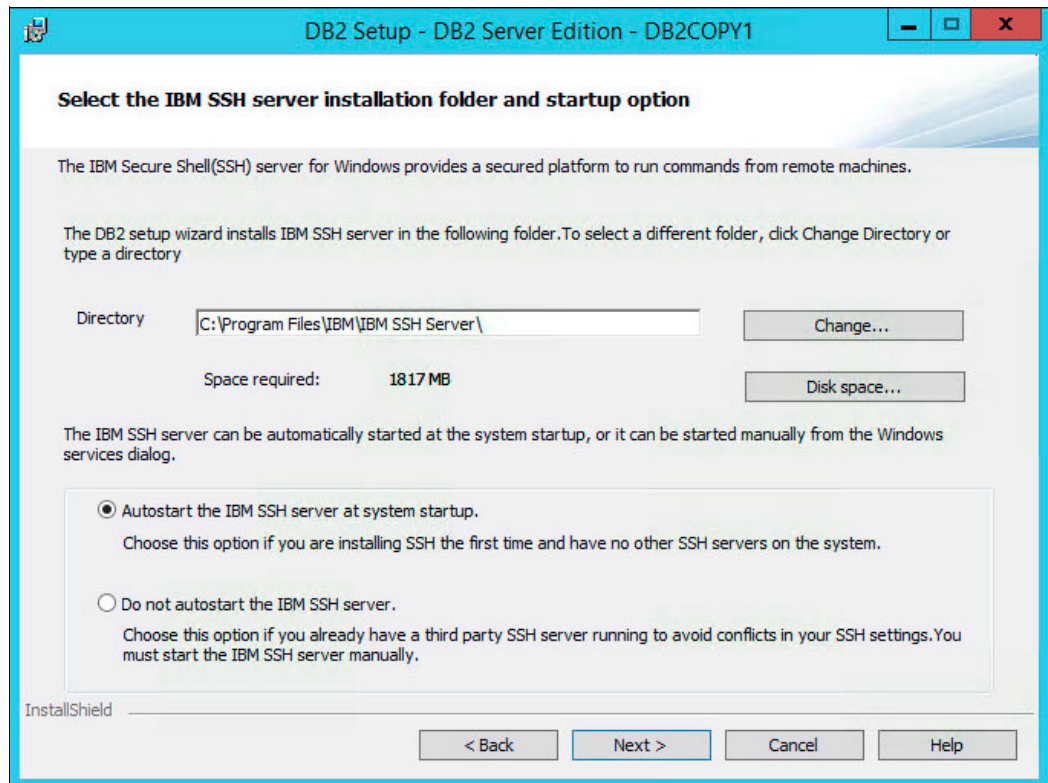


Figure 11-64 DB2 - IBM SSH server installation location

11. Specify db2admin as the user for the DB2 Administration Server (DAS) with Object00 as the password and click **Next** to continue, as shown in Figure 11-65.

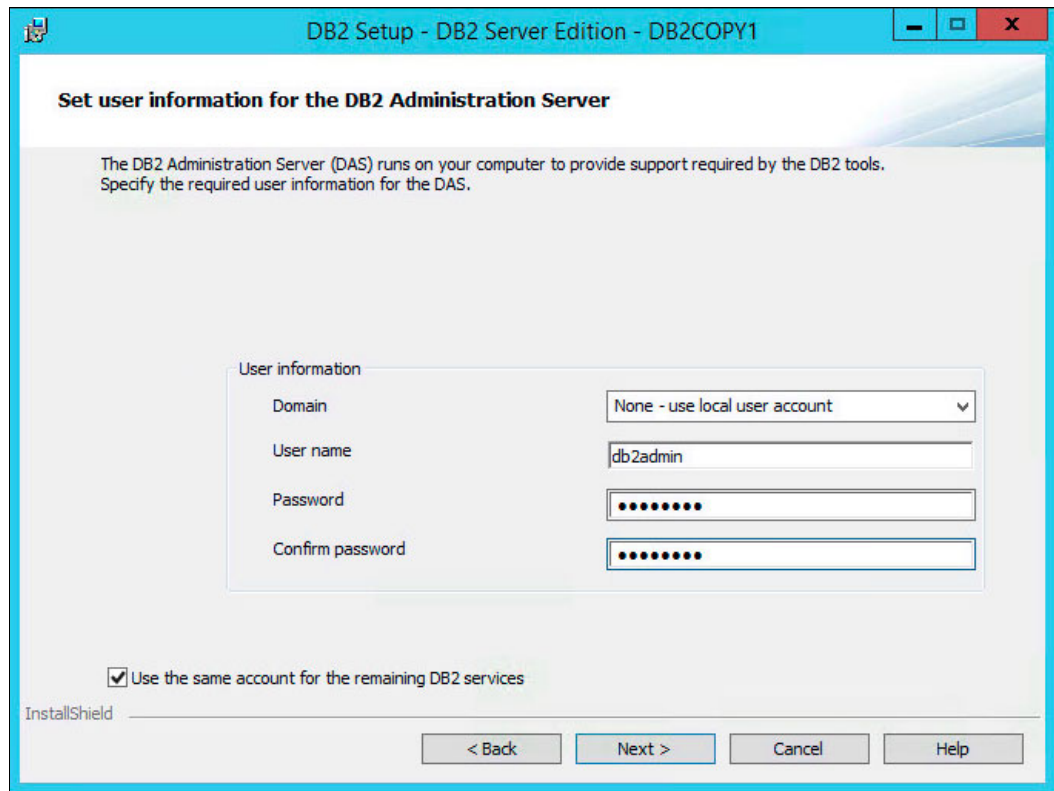


Figure 11-65 DB2 - define a DAS user account

12. Keep the default DB2 instance and click **Next** to continue, as shown in Figure 11-66.

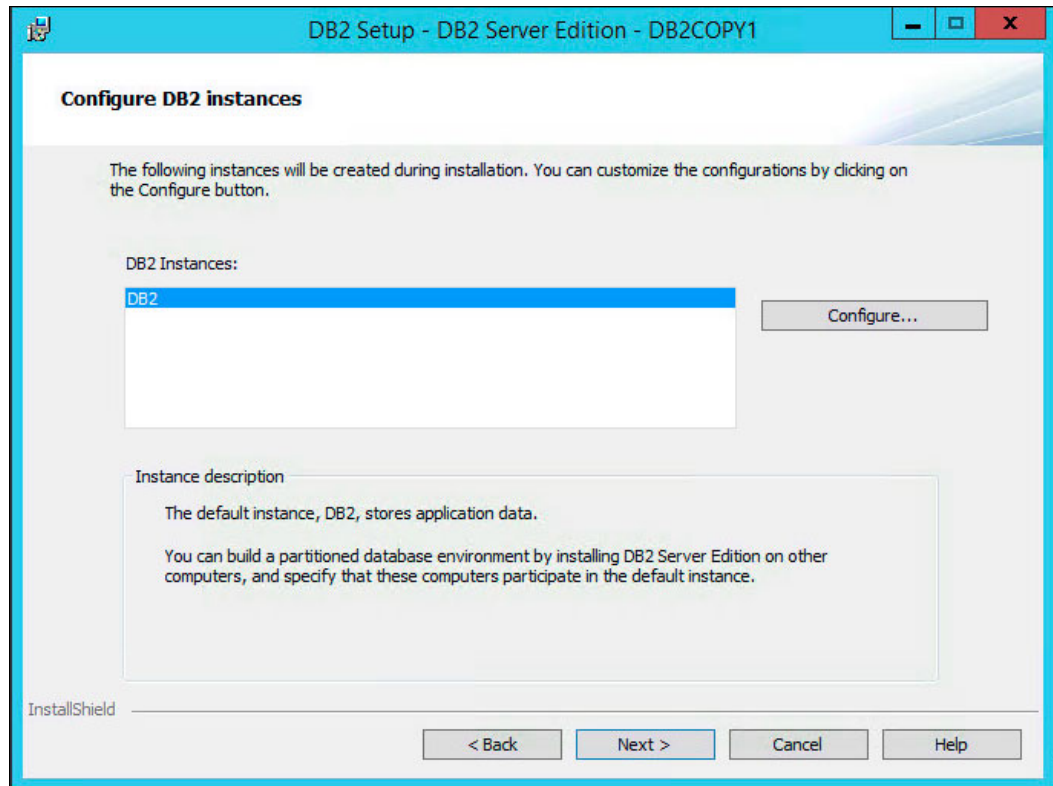


Figure 11-66 DB2 - define the default DB2 instance

13. Specify the host name or IP address of your SMTP server for the DB2 server to send notifications to and click **Next** to continue, as shown in Figure 11-67.

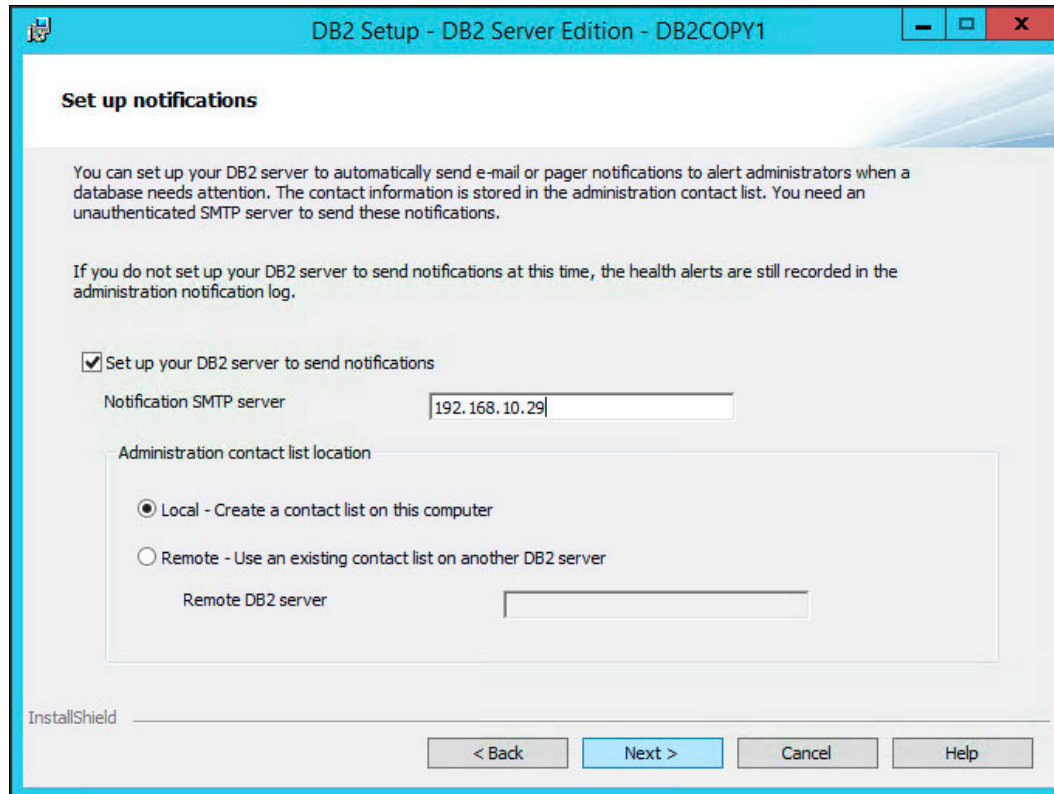


Figure 11-67 DB2 - configure the SMTP notification server

14. Specify a contact to receive the email notifications and click **Next** to continue, as shown in Figure 11-68.

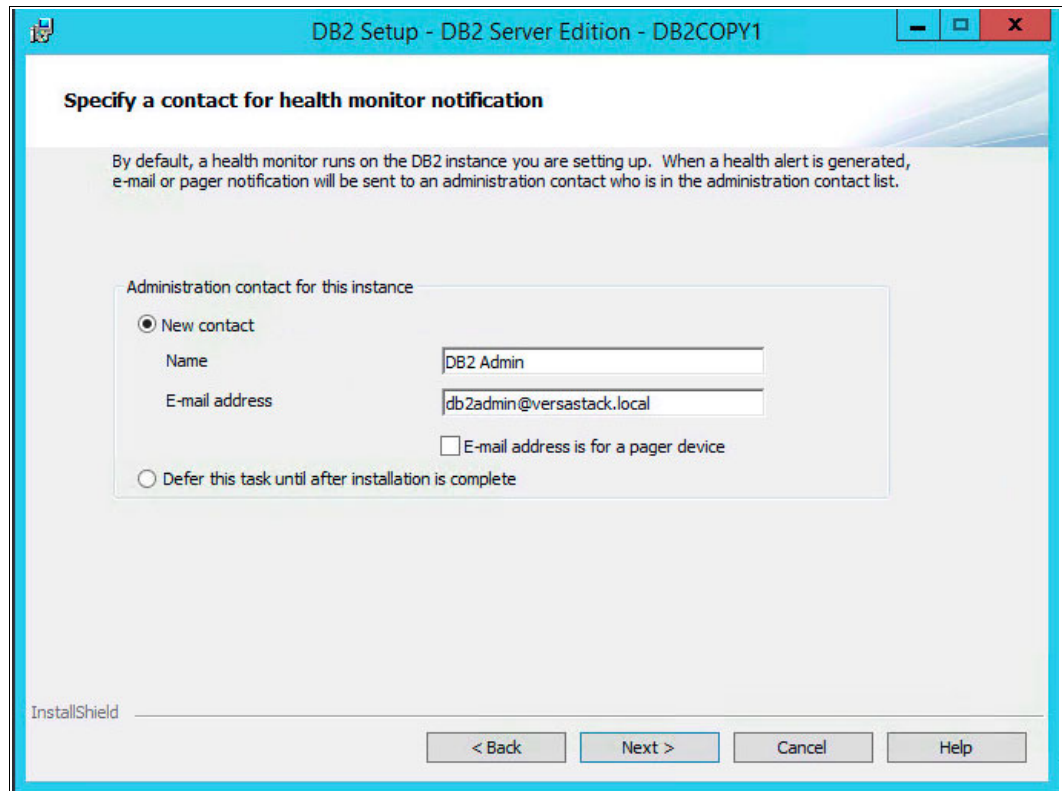


Figure 11-68 DB2 - instance administration contact

15. Optionally, enable operating systems security for DB2 objects and click **Next** to continue, as shown in Figure 11-69.

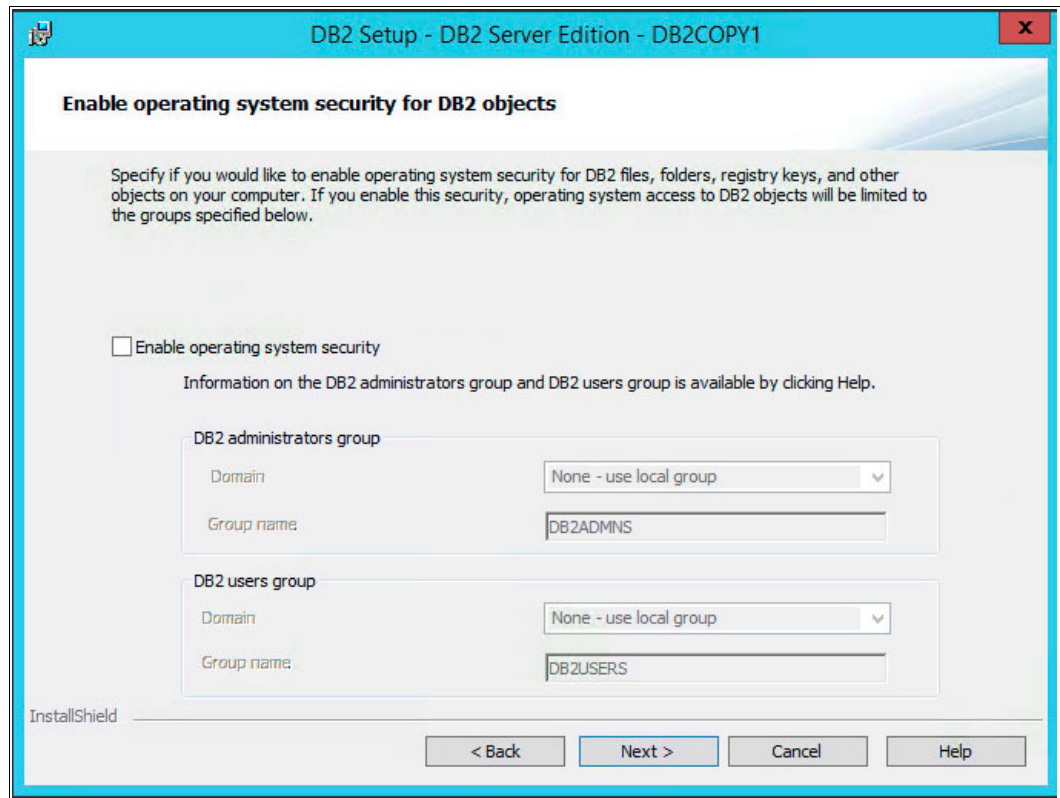


Figure 11-69 DB2 - enable operating system security for DB2 objects

16. Review the settings and click **Finish** to start the installation, as shown in Figure 11-70.

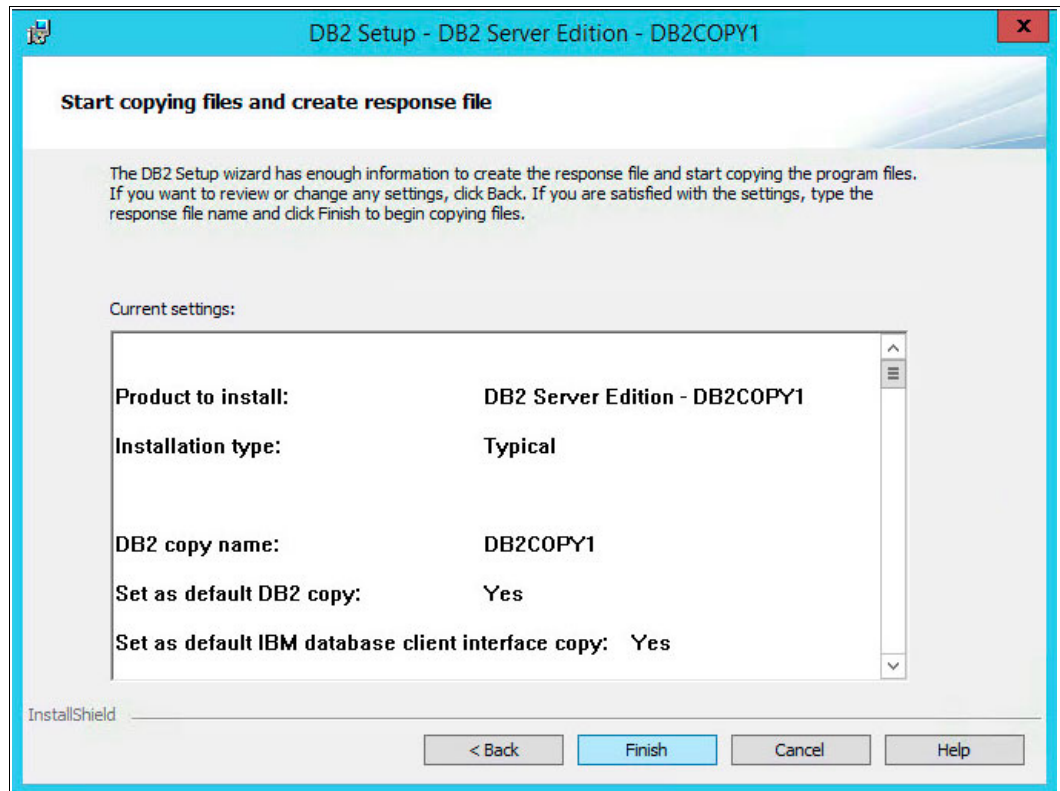


Figure 11-70 DB2 - finish the installation wizard

Do not install additional components after the installation finishes. Close the setup windows instead.

Figure 11-71 shows the Setup is complete window.

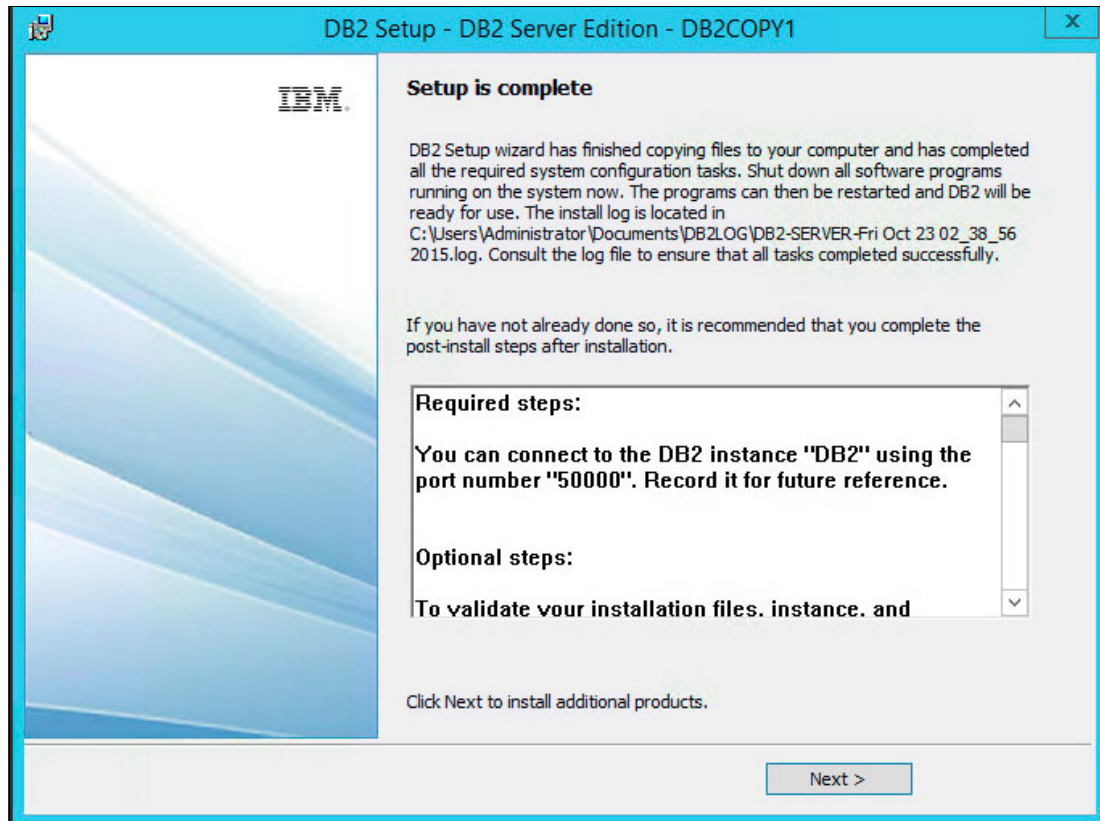


Figure 11-71 DB2 - instance deployment is complete

11.7 Installing DB2 Data Studio Client

IBM Data Studio consists of the Data Studio client and the Data Studio web console, which provide the tools that you need for developing database applications, administering databases and database instances, and tuning queries.

The Data Studio client provides an integrated development environment (IDE) for routine and Java application development, database and database instance administration, and query tuning.

The Data Studio web console provides health and availability monitoring features and job creation and management tools for your databases.

To install the DB2 Data Studio Client, complete the following steps:

1. Extract the DATA_STUDIO_CLIENT_V4.1.2_FOR_WIN.zip file, go to the extracted folder location, and double-click launchpad.exe to start the Data Studio Client installation, as shown in Figure 11-72.

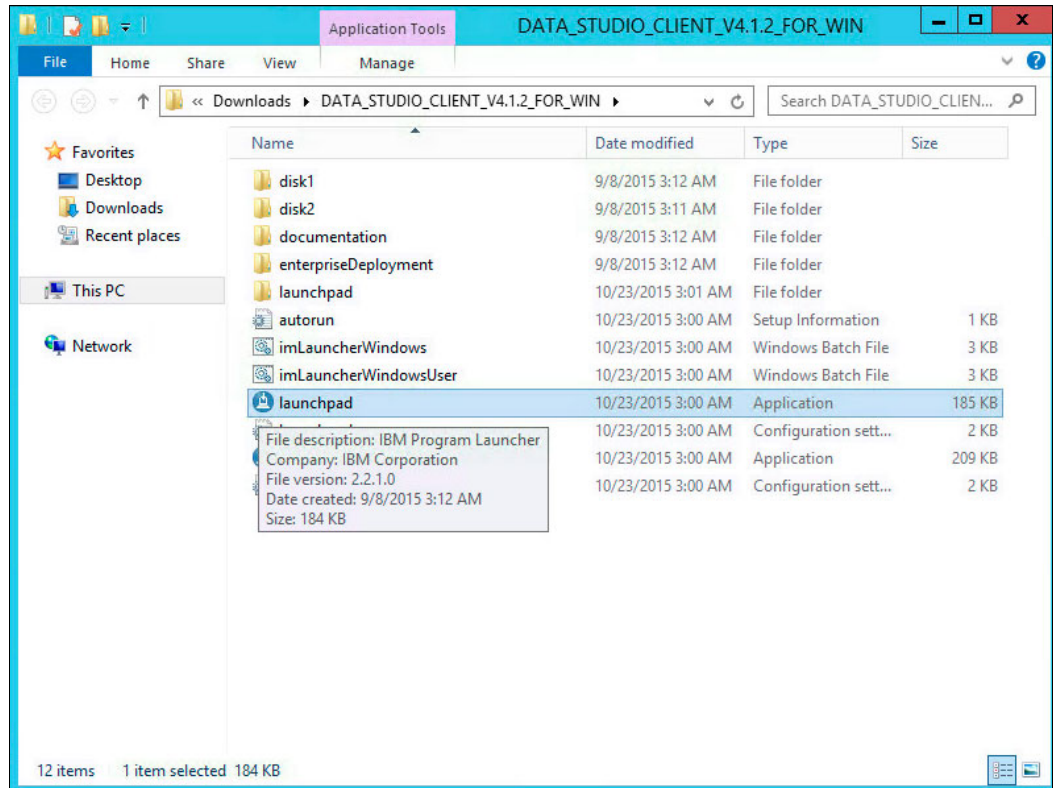


Figure 11-72 Data Studio Client - installation launchpad

2. In the Data Studio Client Welcome window, click **Install or Update Product**, as shown in Figure 11-73.

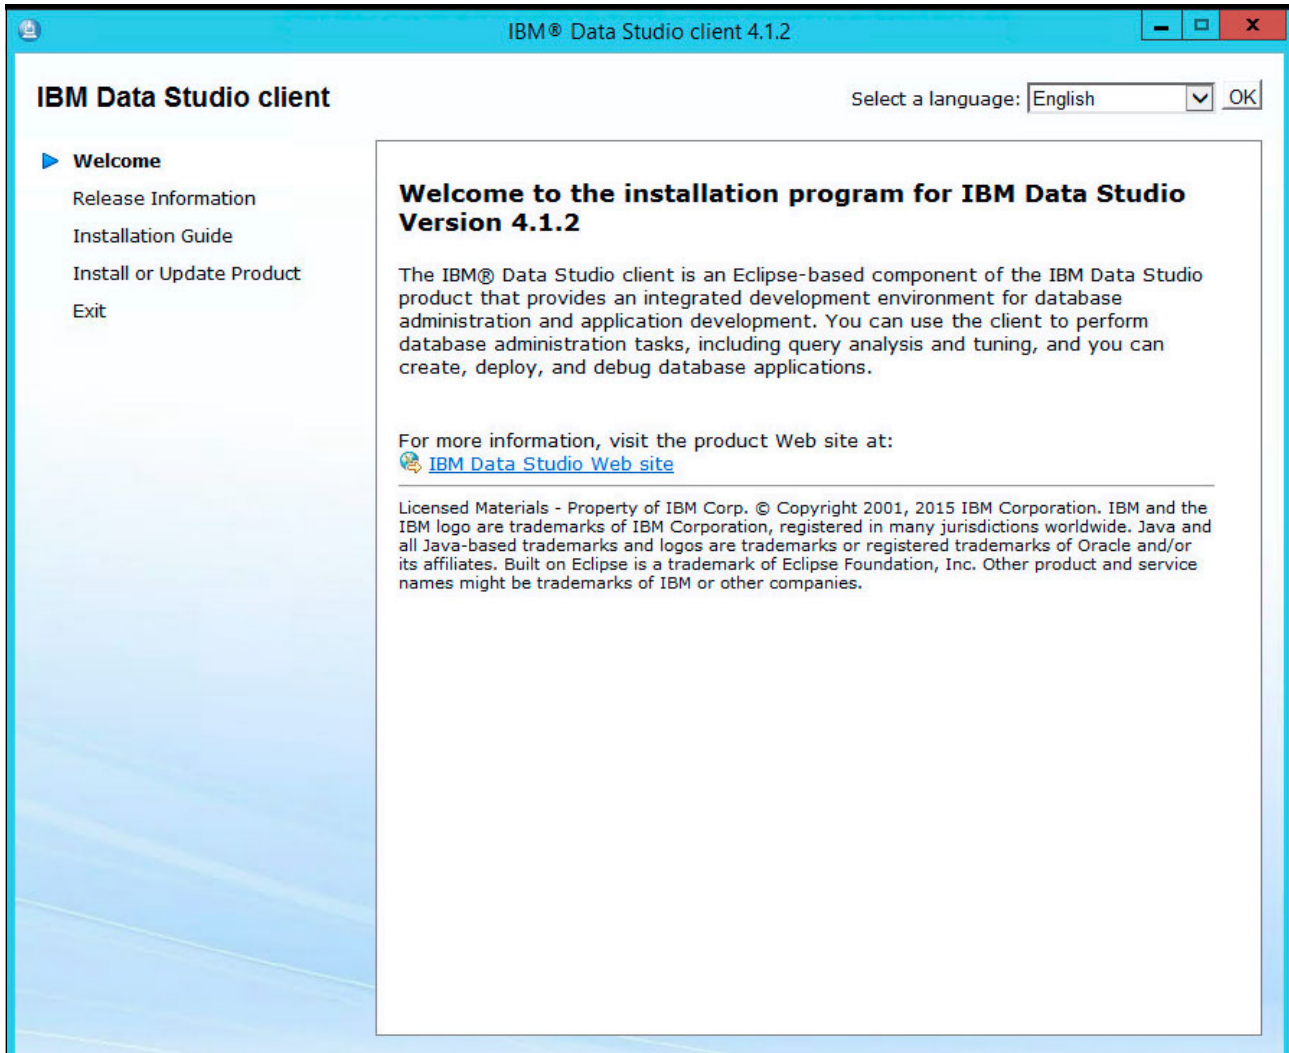


Figure 11-73 Data Studio Client - installation program

- The deployment of the IBM Installation Manager, a common installer component that is used by IBM software such as IBM DB2 Data Studio and IBM Spectrum Protect and IBM Spectrum Control servers, starts. Click **Next** to continue, as shown in Figure 11-74.

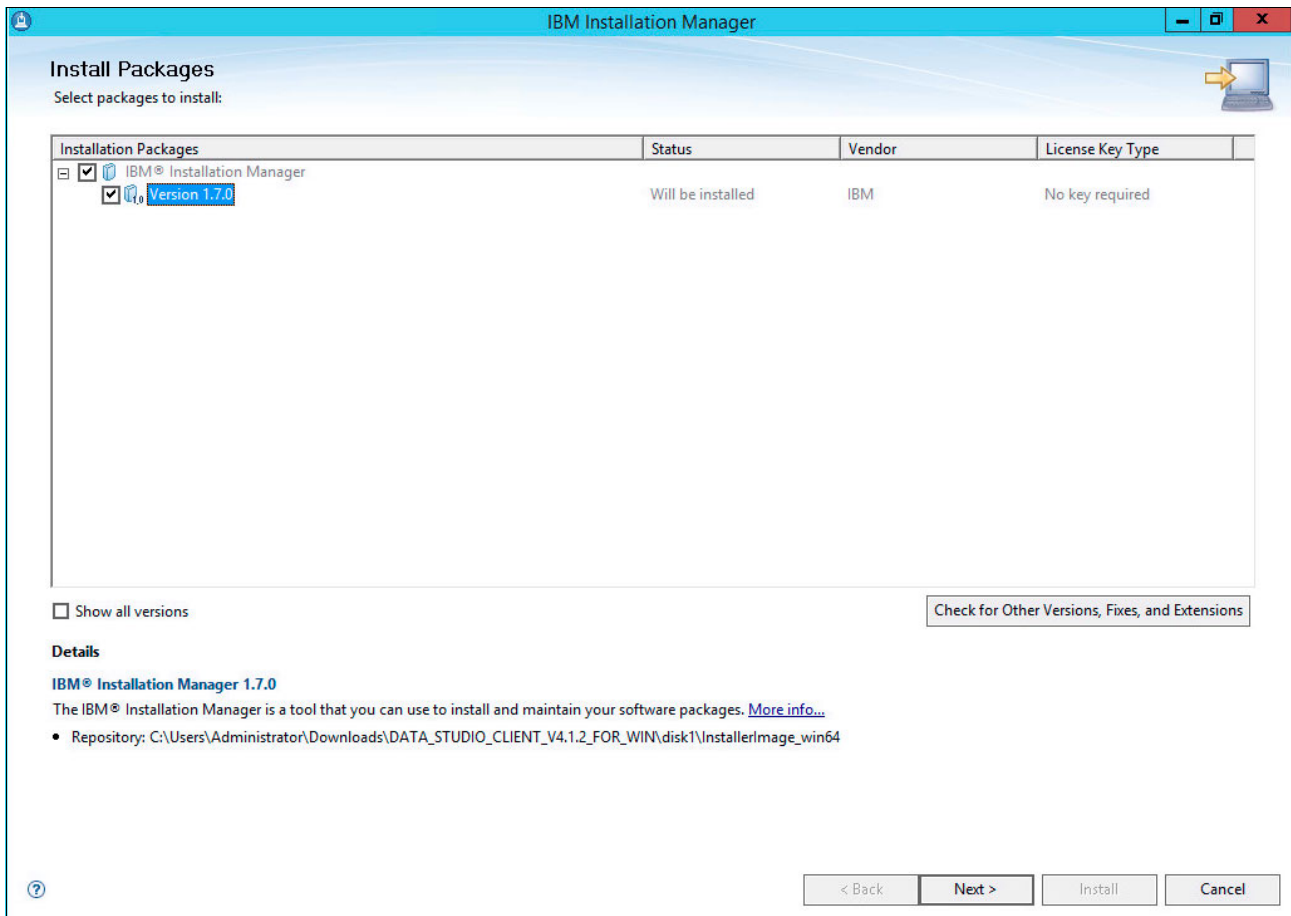


Figure 11-74 Data Studio Client - IBM Installation Manager deployment

4. Accept the license agreement for the IBM Installation Manager and click **Next** to continue, as shown in Figure 11-75.

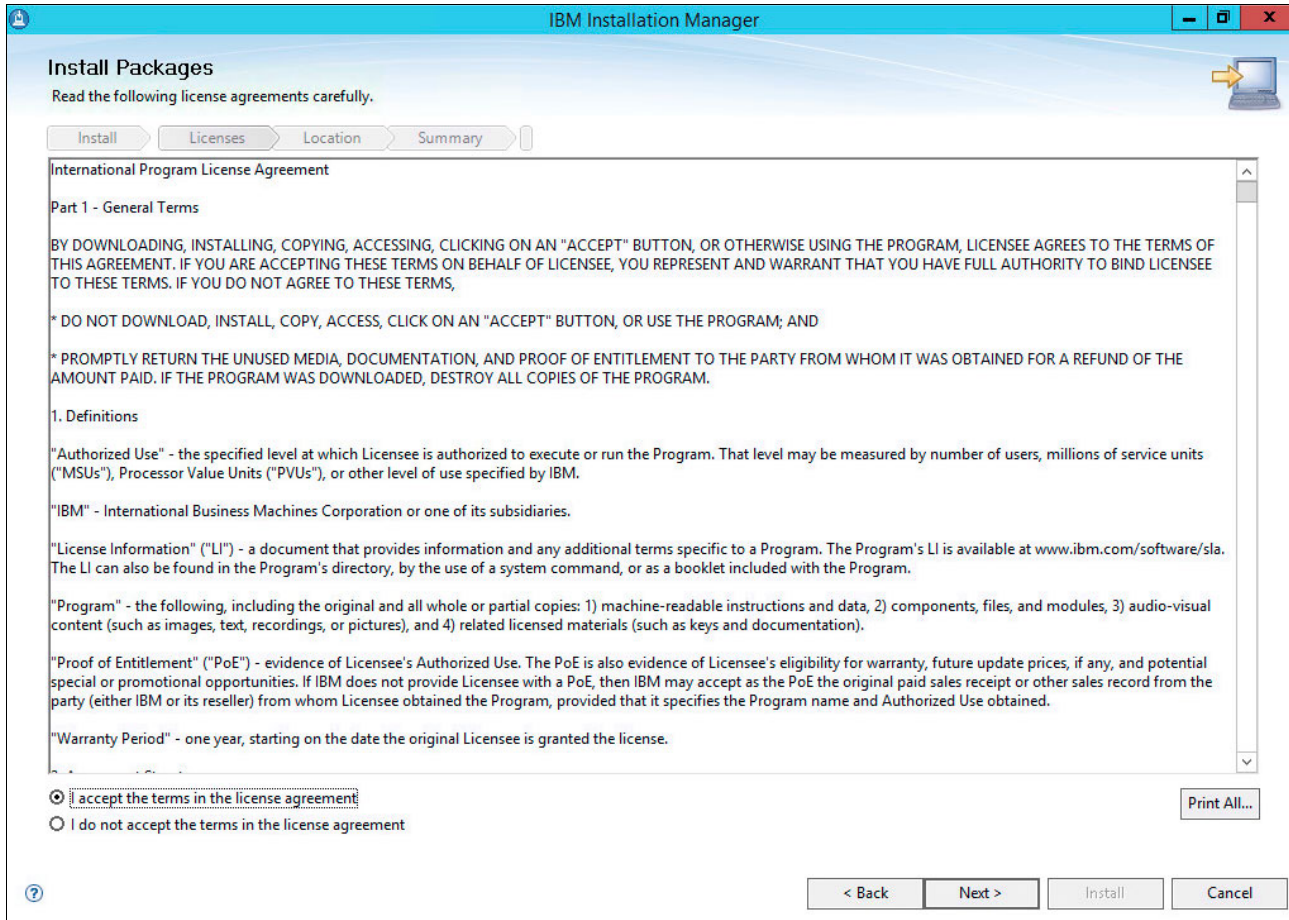


Figure 11-75 IBM Installation Manager - accept the license terms

5. Keep the default installation directory and click **Next** to continue, as shown in Figure 11-76.




Figure 11-76 IBM Installation Manager - installation location

6. Click **Install** to start the IBM Installation Manager deployment, as shown in Figure 11-77.

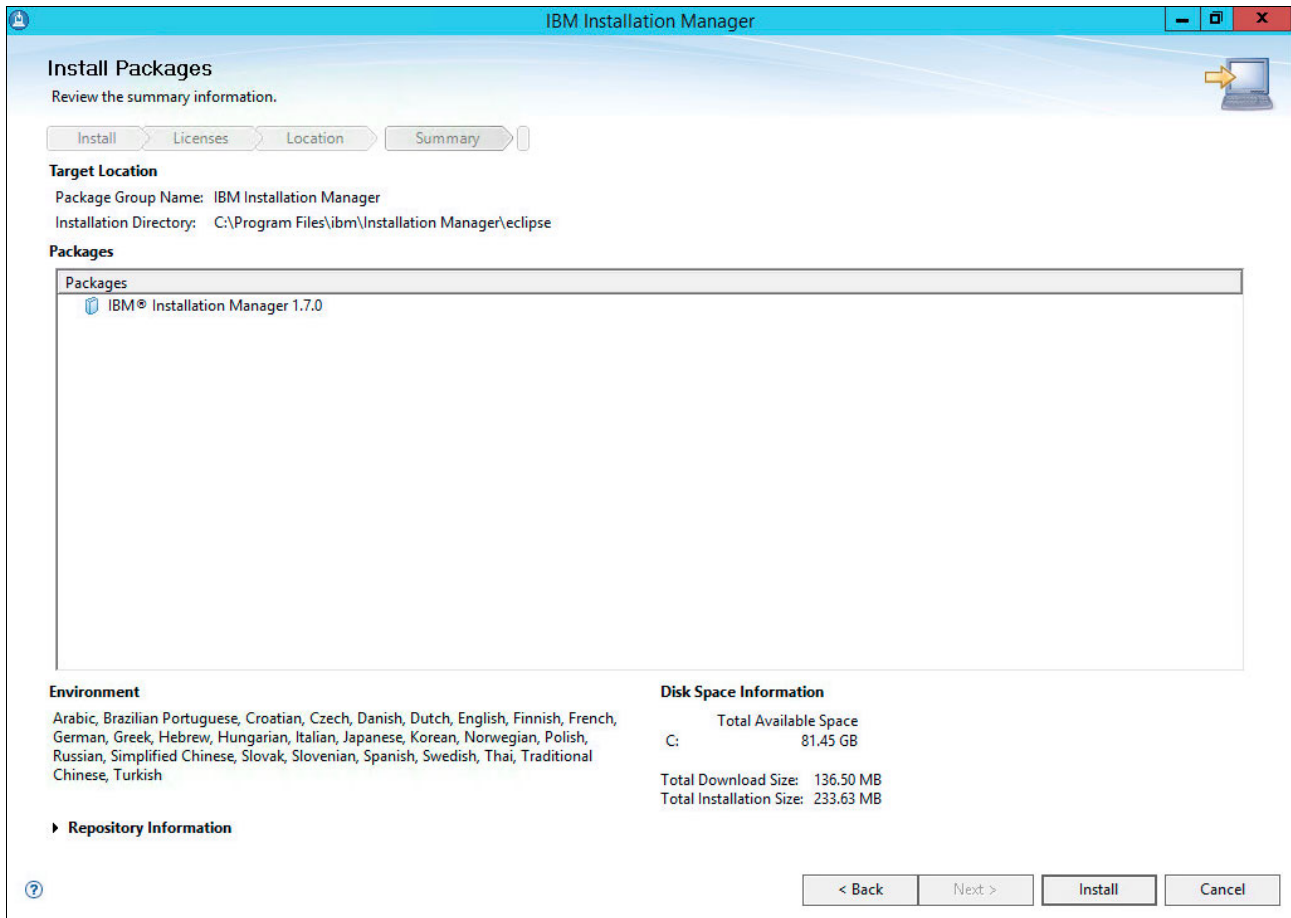


Figure 11-77 IBM Installation Manager - installation

7. After the installation, the IBM Installation Manager restarts. Click **Install** to select the software packages to be installed, as shown in Figure 11-78.

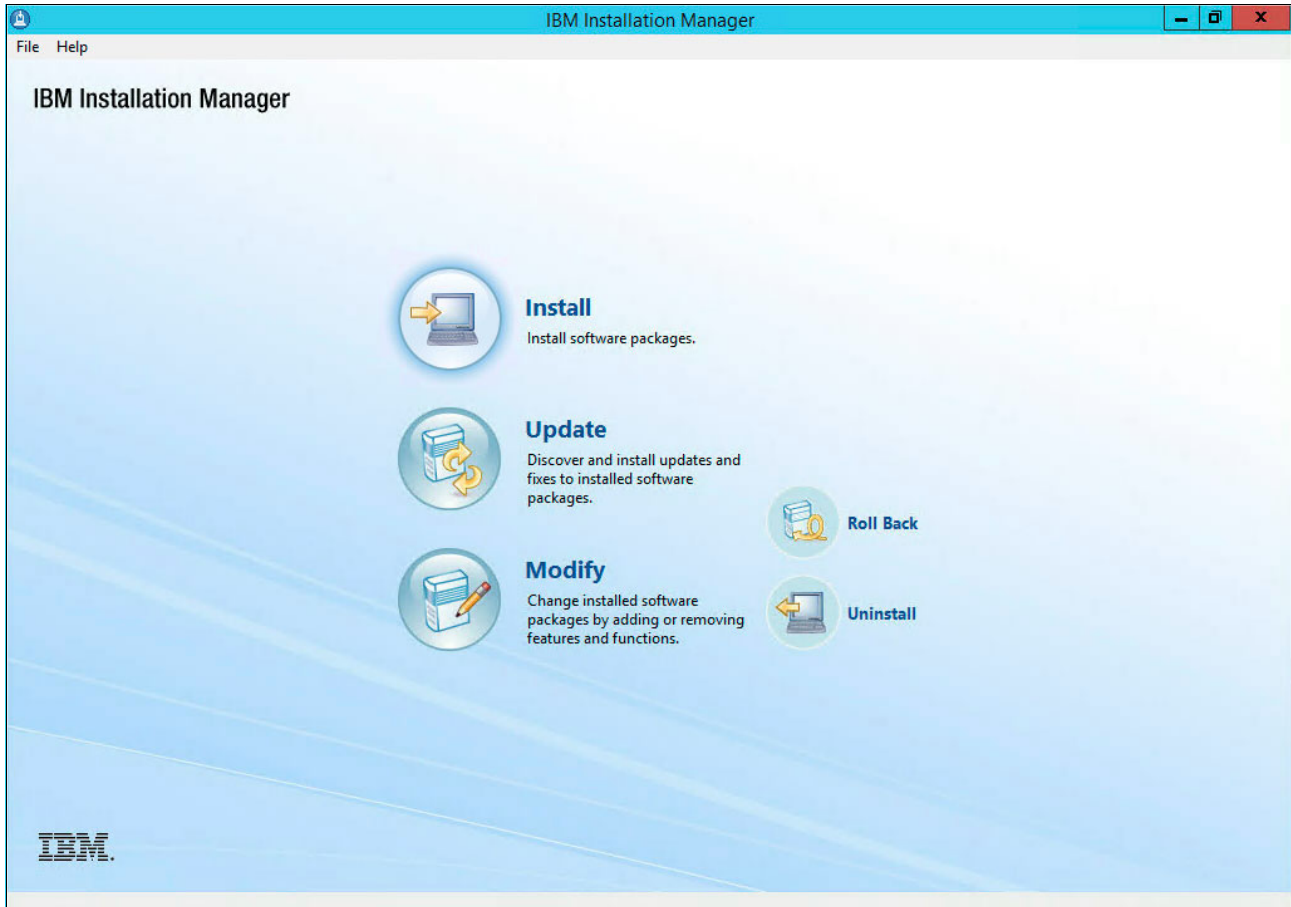


Figure 11-78 IBM Installation Manager - landing page

8. IBM Data Studio Client V4.1.2 is selected by default. Click **Next** to continue, as shown in Figure 11-79.

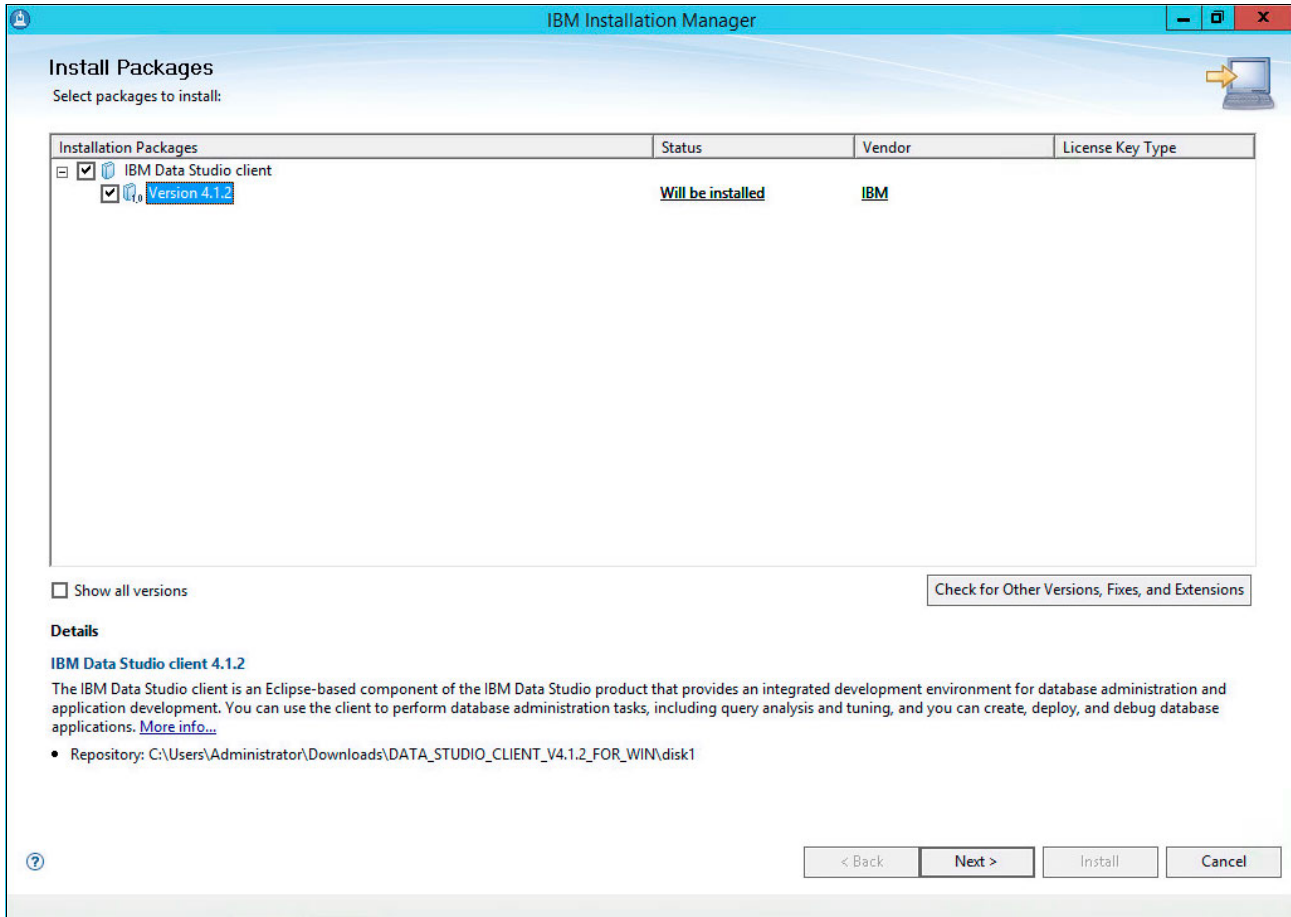


Figure 11-79 Data Studio Client - Version 4.1.2 installation packages selection

9. Accept the license agreement and click **Next** to continue, as shown in Figure 11-80.

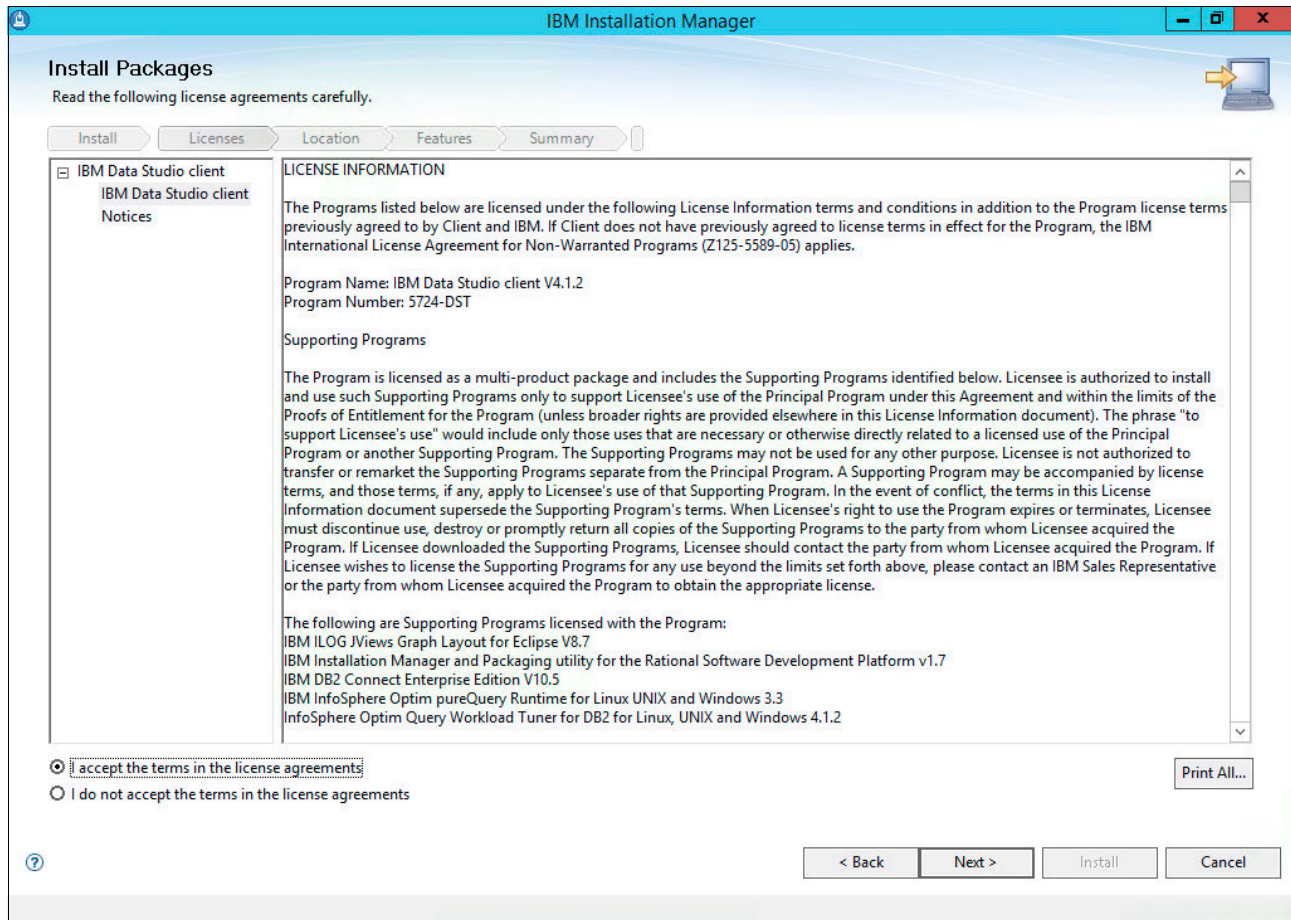


Figure 11-80 Data Studio Client - license agreement

10. Keep the default shared resource directory and click **Next** to continue, as shown in Figure 11-81.

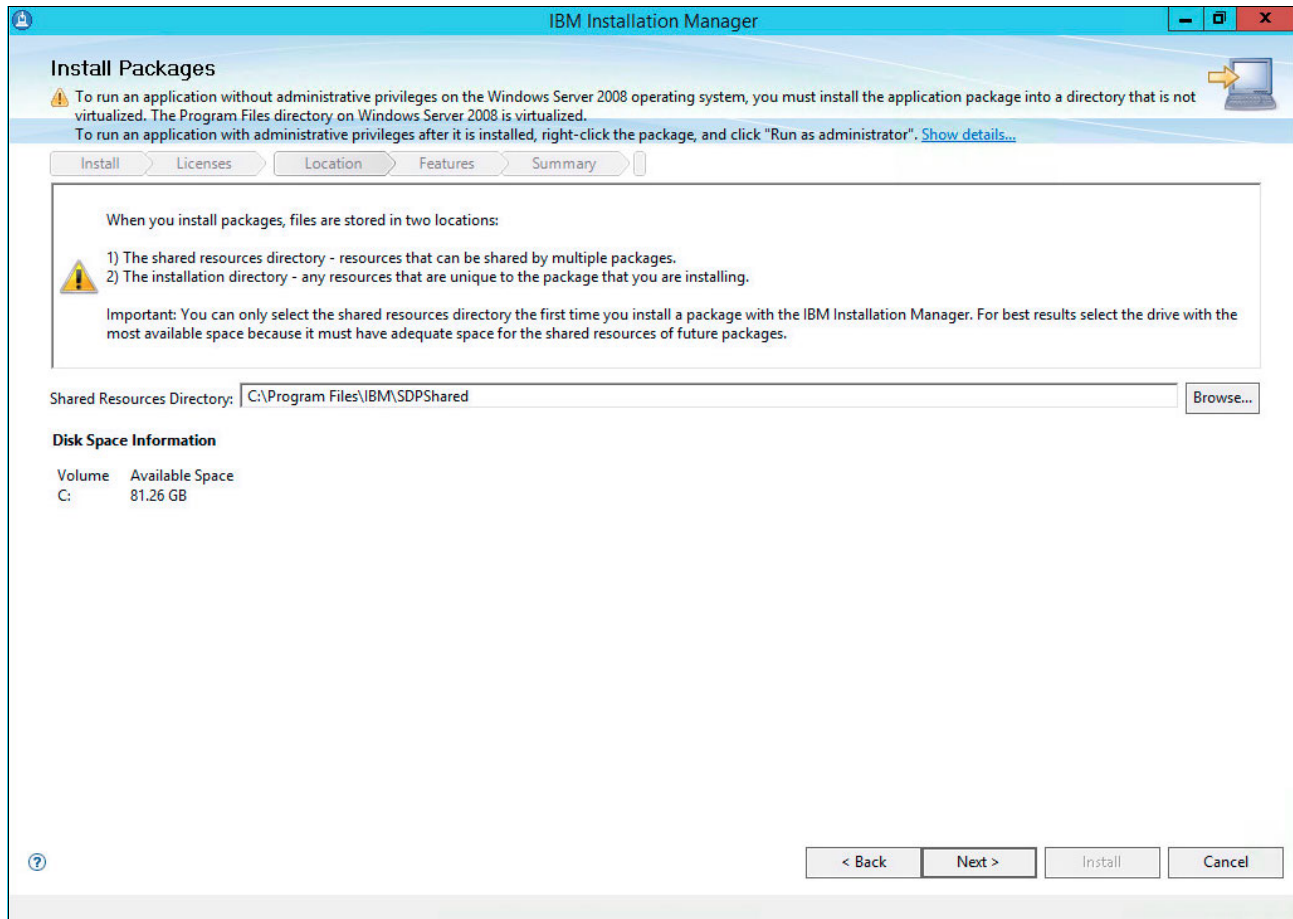


Figure 11-81 Data Studio Client - specify the Shared Resource Directory

11. Keep the default location for the IBM Data Studio Package Group and click **Next** to continue, as shown in Figure 11-82.




Figure 11-82 Data Studio Client package group installation Location

12. Optionally, specify additional language support for the Data Studio Client and click **Next** to continue, as shown in Figure 11-83.

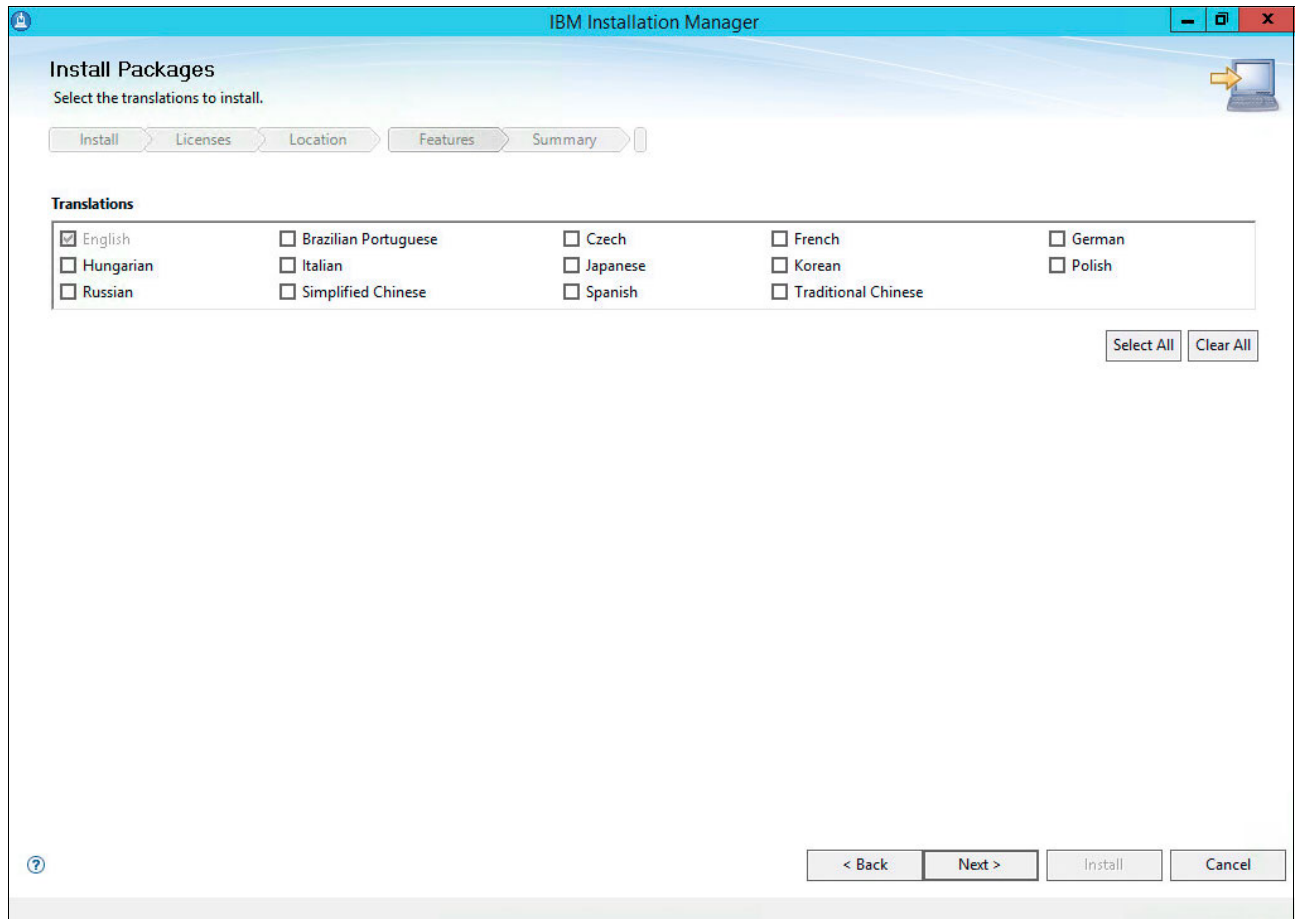


Figure 11-83 Data Studio Client - language translations

13. Keep the default Full Production Option and click **Next** to continue, as shown in Figure 11-84.

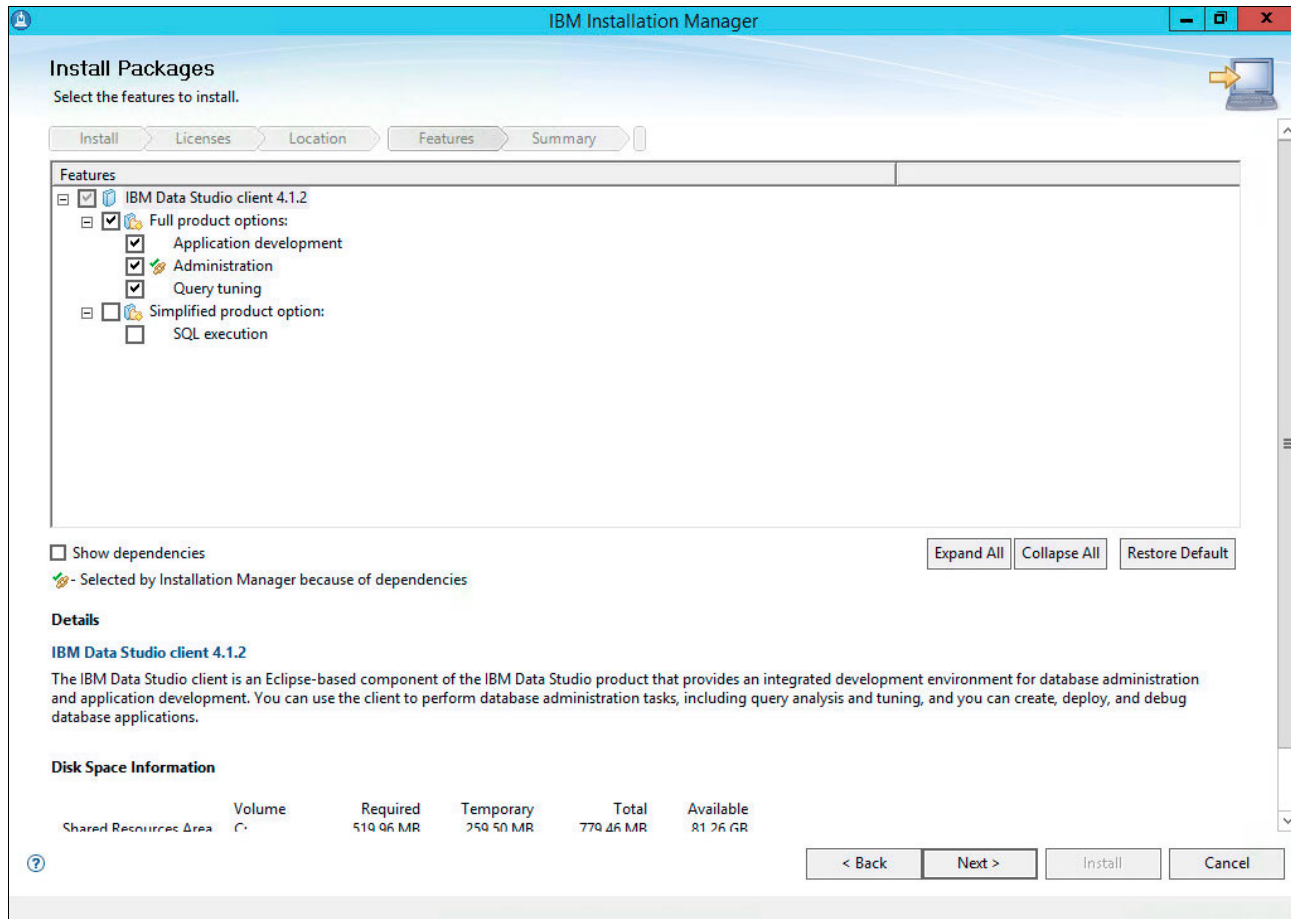


Figure 11-84 Data Studio Client - select product options

14. Click **Install** to start the installation process, as shown in Figure 11-85.

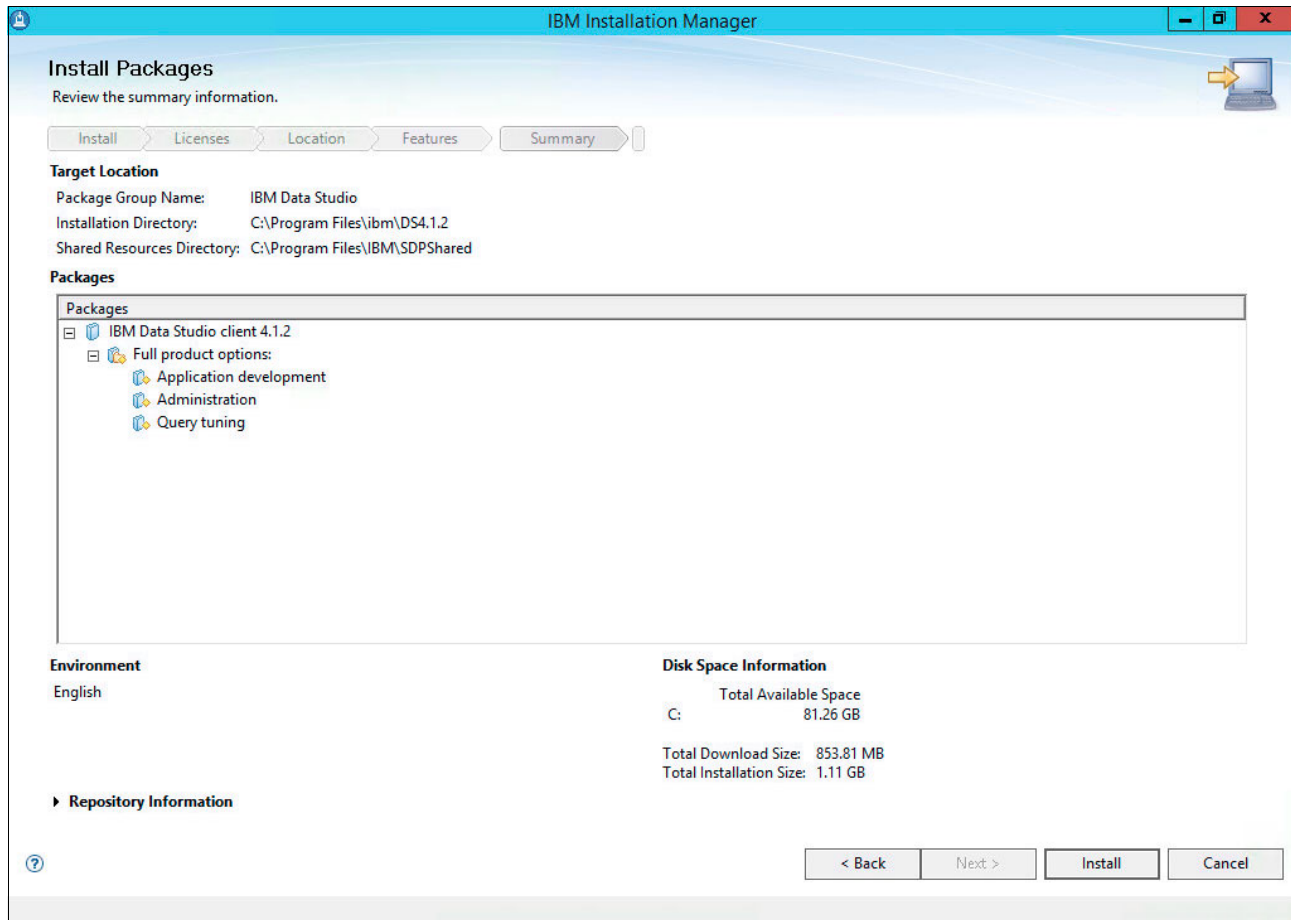


Figure 11-85 Data Studio Client Installation

15. After the installation completes, you return to the IBM Installation Manager window. Click **File/Exit** to finish the installation.

You completed the deployment of the DB2 management software on the SQLVM01 VM. You can use the Data Studio client to accomplish the following tasks:

- ▶ Connect to both the primary DB2 instance on SQLVM01 and the standby DB2 instance on SQLVM02.
- ▶ Configure database logging on the sample databases.
- ▶ Set up high availability and disaster recovery.

For an overview of the Data Studio functions and the key tasks that it can perform, go to the following website:

https://www.ibm.com/support/knowledgecenter/#/SS62YD_4.1.1/com.ibm.datatools.ds.release.doc/topics/dsoverview.html

In the next section, you configure a sample database on both DB2 servers.

11.8 Deploying the DB2 GSDB sample database

The GSDB database is a rich and realistic database that contains sample data for the Sample Outdoor Company, which is a fictional outdoor equipment retailer. The database includes data for operations, human resources, sales, marketing, and finance.

The database must be deployed on both the SQLVM01 and SQLVM02 VMs. To do so, complete the following steps:

1. Go to the folder where you extracted GSDB_DB2_LUW_ZOS_v2r3.zip. Go to the DB2\win subdirectory, right-click the setupGSDB.bat file, and select **Run as Administrator**, as shown in Figure 11-86.

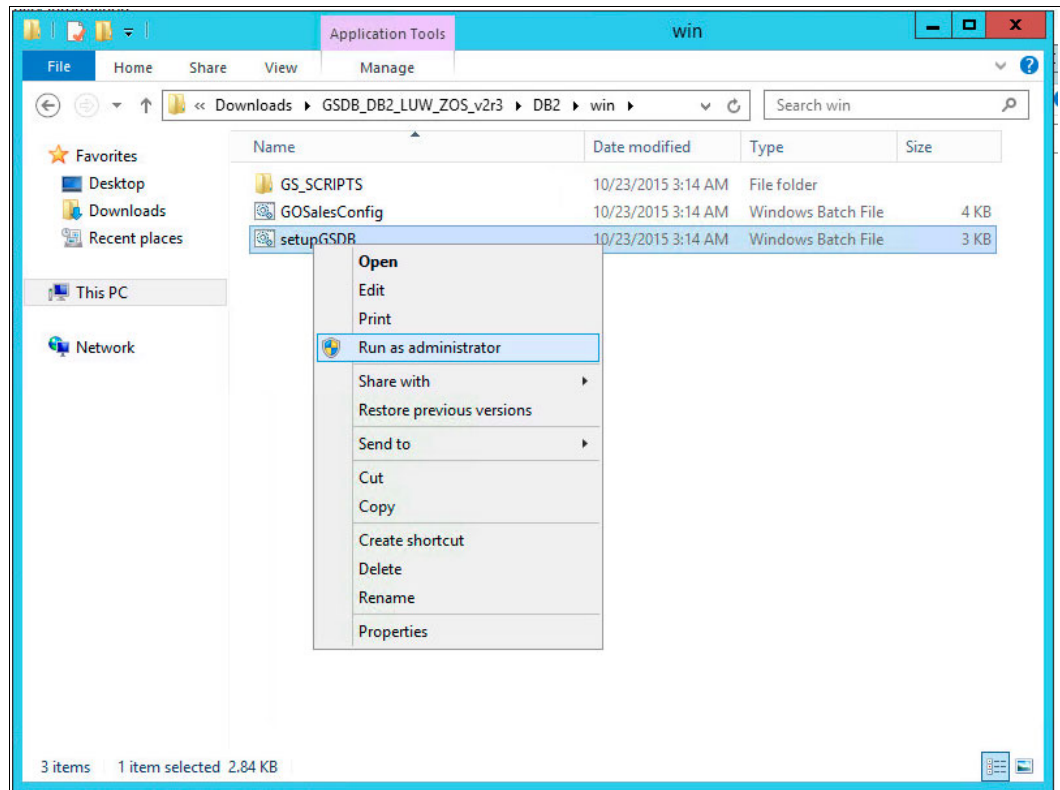


Figure 11-86 GSDB - deployment

2. Provide the database name that you want to use for the GOSales sample data, for example, VersaStackDB2, and press Return, as shown in Figure 11-87.

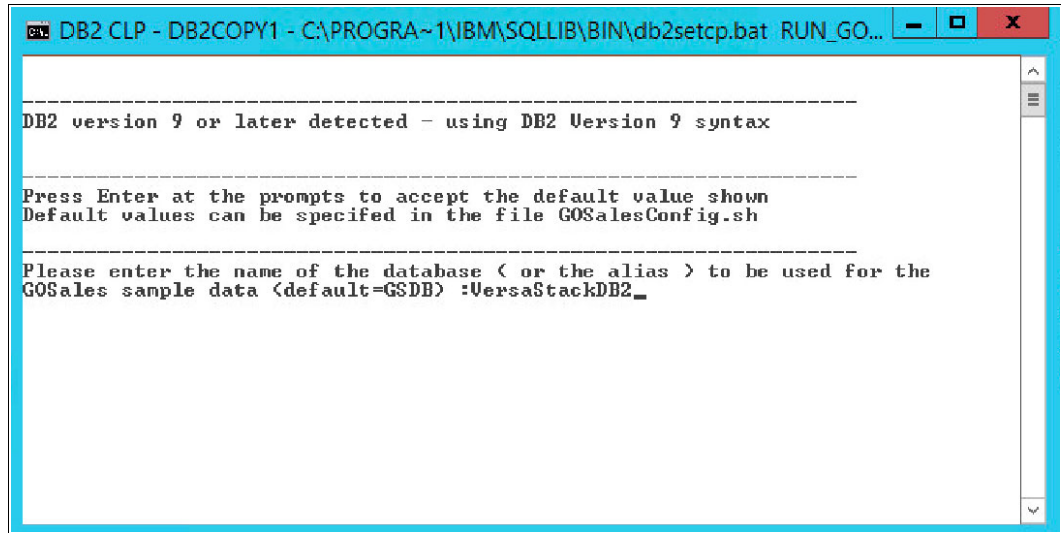


Figure 11-87 GSDB - define the database name

3. Enter Y to create the VersaStackDB2 GSDB-based sample database, as shown in Figure 11-88.

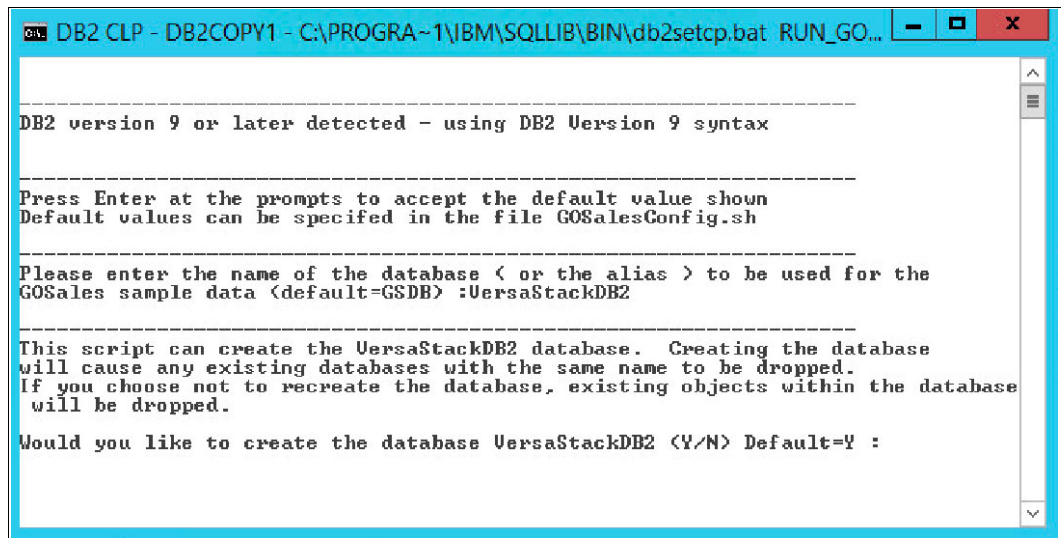


Figure 11-88 GSDB - confirm DB creation

4. In the example environment, we created a database with the same name. Confirm that this database will be dropped as part of the deployment by entering Y, as shown in Figure 11-89 on page 247.


```
DB2 CLP - DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat RUN_GO...
-----
DB2 version 9 or later detected - using DB2 Version 9 syntax
-----
Press Enter at the prompts to accept the default value shown
Default values can be specified in the file GOSalesConfig.sh
-----
Please enter the name of the database ( or the alias ) to be used for the
GOSales sample data (default=GSDB) :VersaStackDB2
-----
This script can create the VersaStackDB2 database.  Creating the database
will cause any existing databases with the same name to be dropped.
If you choose not to recreate the database, existing objects within the database
will be dropped.
Would you like to create the database VersaStackDB2 (Y/N) Default=Y :
-----
The database VersaStackDB2 already exists and will be dropped.
Are you sure you want to drop the existing database? (Y/N) Default=Y :
```

Figure 11-89 GSDB - drop existing database confirmation

5. Use db2admin for the DB2 admin user and press Enter, as shown in Figure 11-90.

```
DB2 CLP - DB2COPY1 - C:\PROGRA~1\IBM\SQLLIB\BIN\db2setcp.bat RUN_GO...
Default values can be specified in the file GOSalesConfig.sh
-----
Please enter the name of the database ( or the alias ) to be used for the
GOSales sample data (default=GSDB) :VersaStackDB2
-----
This script can create the VersaStackDB2 database.  Creating the database
will cause any existing databases with the same name to be dropped.
If you choose not to recreate the database, existing objects within the database
will be dropped.
Would you like to create the database VersaStackDB2 (Y/N) Default=Y :
-----
The database VersaStackDB2 already exists and will be dropped.
Are you sure you want to drop the existing database? (Y/N) Default=Y :
-----
Enter the DB2 administration username for
creating the database and importing the data.
If no value is provided the local user will attempt
to connect without a password.
Please enter the DB2 admin user name :db2admin
```

Figure 11-90 GSDB - specify DB2 admin user

6. Enter Y to start the database creation, as shown in Figure 11-91.

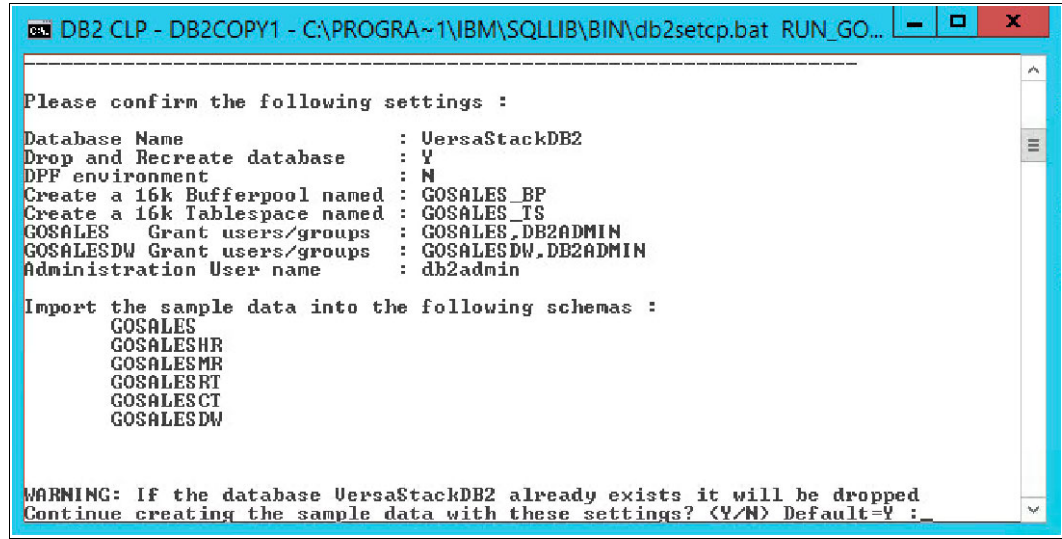


Figure 11-91 GSDB - confirm database creation

The batch command window closes after the creation completes. Repeat the same procedure for SQLVM02.

11.9 Configuring DB2 High Availability

In 11.8, “Deploying the DB2 GSDB sample database” on page 245, you deployed the GOSales sample database on the DB2 instances on SQLVM01 and SQLVM02.

In this section, you use the Data Studio client to set up a connection to this database on both DB2 instances.

Then, you configure the database log rotation by defining a local archive and backup repository for both databases.

Lastly, you use the Setup HADR wizard to define a near-synchronous relationship between the primary database on SQLVM01 and the standby secondary database on SQLVM02.

This wizard restores the backup from the primary database into the secondary database, starts the secondary in HA standby mode, and completes the configuration by starting the database on SQLVM01 as the HA primary DB.

Here is an overview of the available HADR synchronization modes:

► SYNC (synchronous)

This mode provides the greatest protection against transaction loss. However, using it might result in the longest transaction response time among the four modes. In this mode, log writes are considered successful only when logs are written to log files on the primary database and when the primary database receives acknowledgment from the standby database that the logs also are written to log files on the standby database. The log data is guaranteed to be stored at both sites.

► NEARSYNC (near synchronous)

Although this mode has a shorter transaction response time than synchronous mode, it also provides slightly less protection against transaction loss. In this mode, log writes are considered successful only when the log records are written to the log files on the primary database and when the primary database receives acknowledgment from the standby system that the logs also are written to main memory on the standby system. Loss of data occurs only if both sites fail simultaneously and if the target site has not transferred to nonvolatile storage all of the log data that it has received.

► ASYNC (asynchronous)

Compared with the SYNC and NEARSYNC modes, the ASYNC mode results in shorter transaction response times, but might cause greater transaction losses if the primary database fails. In ASYNC mode, log writes are considered successful only when the log records are written to the log files on the primary database and are delivered to the TCP layer of the primary system's host machine. Because the primary system does not wait for an acknowledgment from the standby system, transactions may be considered committed when they are still on their way to the standby database.

► SUPERASYNC (super asynchronous)

This mode has the shortest transaction response time, but also has the highest probability of transaction losses if the primary system fails. This mode is useful when you do not want transactions to be blocked or experience elongated response times because of network interruptions or congestion. In this mode, the HADR pair is never in peer state or disconnected peer state. The log writes are considered successful when the log records are written to the log files on the primary database. Because the primary database does not wait for acknowledgment from the standby database, transactions are considered committed regardless of the state of the replication of that transaction.

Because the transaction commit operations on the primary database are not affected by the relative slowness of the HADR network or the standby HADR server, the log gap between the primary database and the standby database might continue to increase. It is important to monitor the log gap because it is an indirect measure of the potential number of transactions that might be lost if a true disaster occurs on the primary system. In disaster recovery scenarios, any transactions that are committed during the log gap are not available to the standby database. Therefore, monitor the log gap by using the `hadr_log_gap` monitor element. If the log gap is not acceptable, investigate the network interruptions or the relative speed of the standby database node and take corrective measures to reduce the log gap.

11.9.1 Connecting to the DB2 instance with the Data Studio Client

Complete the following steps:

1. Log on to the SQLVM01 system and start the Data Studio V4.1.2 Client, as shown in Figure 11-92.

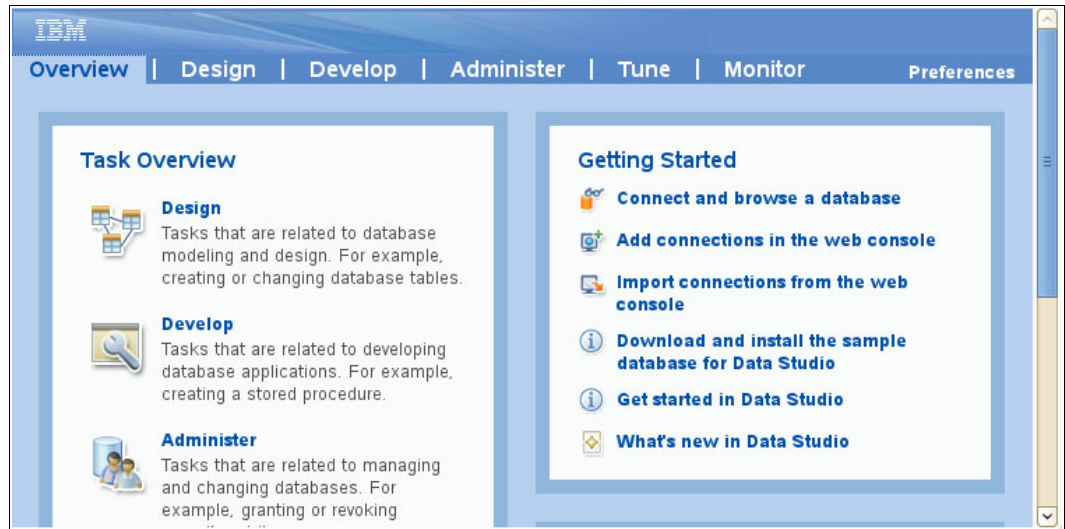


Figure 11-92 Data Studio - client overview window

2. In the Overview window, click **Connect and browse a database** in the Getting Started pane. In the Data Sources subpane, right-click **Database Connections** and select New, as shown in Figure 11-93 on page 251.

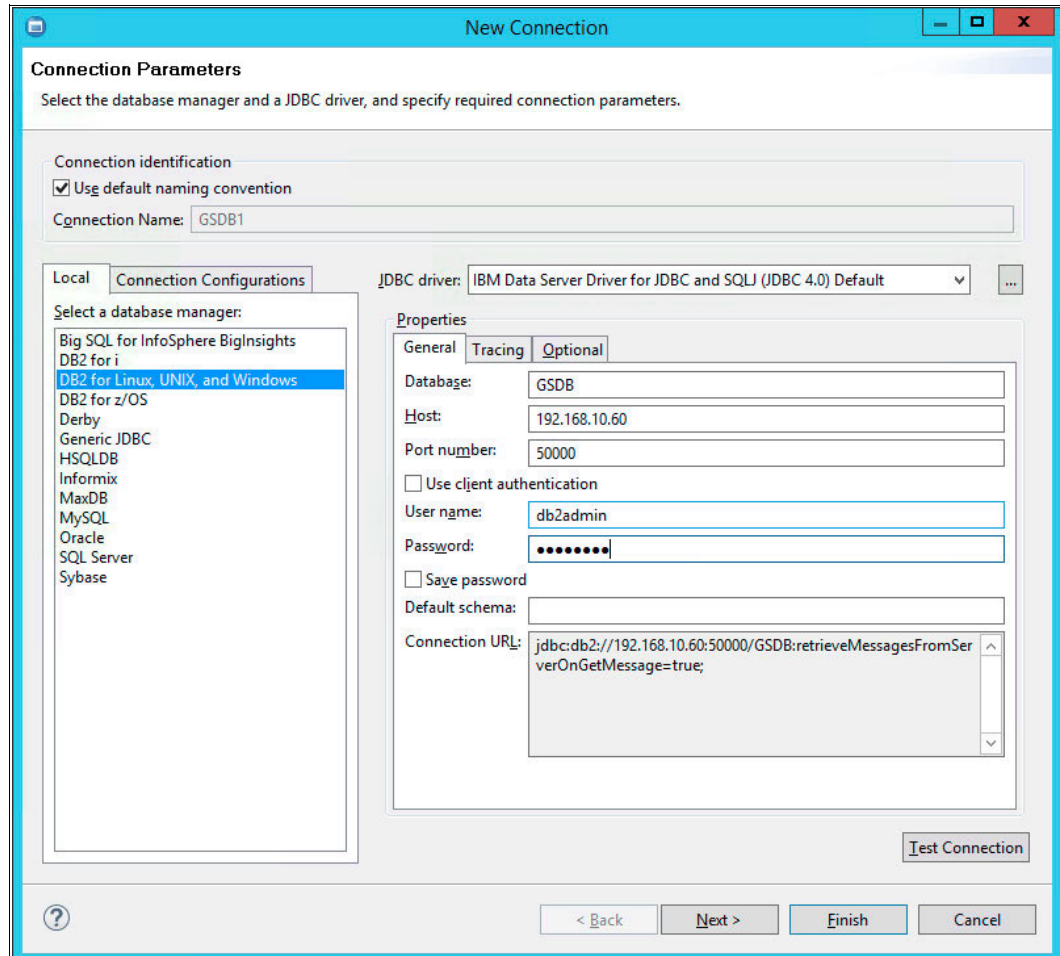


Figure 11-93 Data Studio - new database connection configuration

3. Specify GSDB for the database, the IP address that is used for SQLVM01 (192.168.10.61), and db2admin with Object00 as the password. Click **Test Connection** to verify the connection and click **Finish**.

A new connection appears under the All Databases section in the Data Studio Administration Explorer left pane.

4. Repeat these steps to set up a connection to the GSDB database that is deployed on SQLVM02 (192.168.10.62).

11.9.2 Configuring the database transaction logging

Complete the following steps:

1. Create two directories to hold the archive logs (c:\archlog) and the database backups (c:\dbbackup).
2. In the Data Studio client, expand the 192.168.10.61 entry, right-click the GSDB database, and select **Set Up and Configure/Configure Database Logging**, as shown in Figure 11-94.




Figure 11-94 GSDB - Configure Database Logging

3. In the Configure Database Logging GSDB section, specify the following items:
 - Logging Type:
 - Database logging type: Archive.
 - Archive log handling: Automatic DB2 archive with File System for primary archive log Media Type and c:\archlog for the archive log location.
 - Backup Image: Select **File System** for the media path, click **Add**, and specify c:\dbbackup for the backup location
 - Click **Run** to update the DB configuration and run a database backup.
4. Repeat this process for the GSDB that is hosted on SQLVM02 under 192.168.10.62.

Note: In this example, we use the local file system for the archive log and database backup repository. With the IBM Spectrum Protect server deployed in the DB2 on VersaStack environment, you can specify **TSM** (IBM Spectrum Protect) as the destination for both the archive logs and the database backups. This approach is covered in detail in Chapter 13, “IBM Spectrum Protect integration” on page 355.

11.9.3 Configuring the HADR setup

As part of the HADR configuration, a backup of the primary GSDB database is restored on the secondary GSDB database instances. Because we used the File System for the backup target when configuring the database logging, the content of the `c:\dbbackup` directory on SQLVM02 must be replaced by the content of the corresponding directory on SQLVM01.

Complete the following steps:

1. Within the Data Studio client, expand the 192.168.10.61 entry, right-click the GSDB database, and select **Setup HADR**, as shown in Figure 11-95.

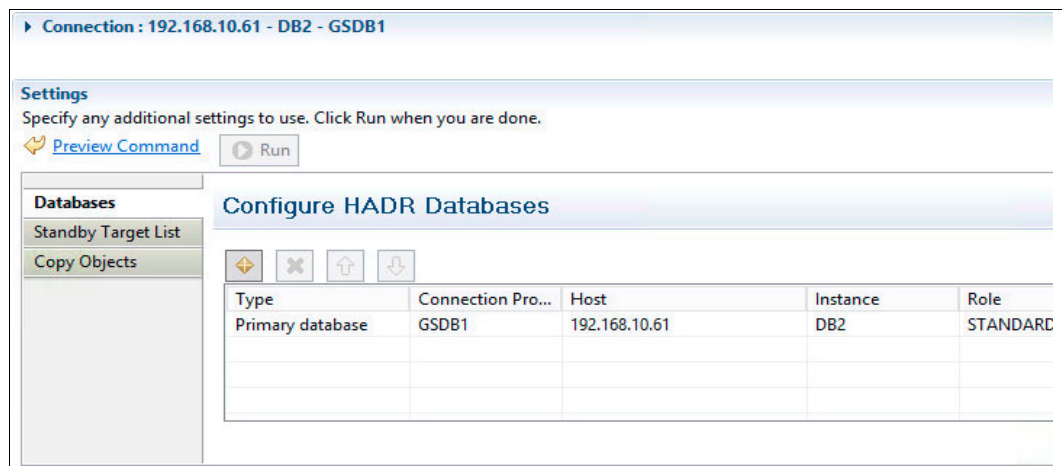


Figure 11-95 Data Studio - set up HADR GSDB

2. Click the plus icon to add the Principal Standby system and select that line item:
 - a. In the Connection profile, select the existing GSDB2 connection (that corresponds with the connection that we created for the SQLVM02 DB2). The Instance name is automatically populated with the main DB2 instance that we configured during the deployment.
 - b. Specify a port number for the HADR service name or port number, for example, 60000. Make sure that this port is not in use by other applications or blocked by the operating systems firewall.
 - c. Leave the other default options.
3. Click the Primary database line item:
 - a. Specify a port number for the HADR service name or port number, for example, 60000. Make sure that this port is not in use by other applications or blocked by the operating systems firewall.
 - b. In the Client Reroute Options, select **192.168.10.62** as the alternative host name and **50000** for the alternative port number.
4. Click **Run** to run the HADR setup and let it run to successful completion.

Example 11-1 shows the sample output when manually configuring the DB2 instance for HADR setup when using the DB2 command windows.

Example 11-1 DB2 command window HADR setup

```
--Primary Database Configuration Commands
-- Configure database for client reroute - 192.168.10.62(DB2) - GSDB
UPDATE ALTERNATE SERVER FOR DATABASE GSDB USING HOSTNAME 192.168.10.62 PORT 50000
-- Update HADR configuration parameters on database - 192.168.10.61(DB2) - GSDB
UPDATE DB CFG FOR GSDB USING HADR_LOCAL_HOST 192.168.10.61
UPDATE DB CFG FOR GSDB USING HADR_LOCAL_SVC 60000
UPDATE DB CFG FOR GSDB USING HADR_REMOTE_HOST 192.168.10.62
UPDATE DB CFG FOR GSDB USING HADR_REMOTE_SVC 60000
UPDATE DB CFG FOR GSDB USING HADR_REMOTE_INST DB2
UPDATE DB CFG FOR GSDB USING HADR_TIMEOUT 120
UPDATE DB CFG FOR GSDB USING HADR_TARGET_LIST 192.168.10.62:60000
UPDATE DB CFG FOR GSDB USING HADR_SYNCMODE NEARSYNC
UPDATE DB CFG FOR GSDB USING HADR_PEER_WINDOW 0
UPDATE DB CFG FOR GSDB USING HADR_SPOOL_LIMIT 0
UPDATE DB CFG FOR GSDB USING HADR_REPLAY_DELAY 0
UPDATE DB CFG FOR GSDB USING BLOCKNONLOGGED NO
--Primary Database Quiesce Commands
CONNECT TO GSDB
QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS
UNQUIESCE DATABASE
CONNECT RESET
--Standby Database Configuration Commands for 192.168.10.62 -DB2
-- Update HADR configuration parameters on database - 192.168.10.62(DB2) - GSDB
UPDATE DB CFG FOR GSDB USING HADR_LOCAL_HOST 192.168.10.62
UPDATE DB CFG FOR GSDB USING HADR_LOCAL_SVC 60000
UPDATE DB CFG FOR GSDB USING HADR_REMOTE_HOST 192.168.10.61
UPDATE DB CFG FOR GSDB USING HADR_REMOTE_SVC 60000
UPDATE DB CFG FOR GSDB USING HADR_REMOTE_INST DB2
UPDATE DB CFG FOR GSDB USING HADR_TIMEOUT 120
UPDATE DB CFG FOR GSDB USING HADR_TARGET_LIST 192.168.10.61:60000
UPDATE DB CFG FOR GSDB USING HADR_SYNCMODE NEARSYNC
UPDATE DB CFG FOR GSDB USING HADR_PEER_WINDOW 0
UPDATE DB CFG FOR GSDB USING HADR_SPOOL_LIMIT 0
UPDATE DB CFG FOR GSDB USING HADR_REPLAY_DELAY 0
UPDATE DB CFG FOR GSDB USING BLOCKNONLOGGED NO
-- Start HADR on standby database - 192.168.10.62(DB2) - GSDB
db2 start hadr on database GSDB as standby
-- Start HADR on primary database - 192.168.10.61(DB2) - GSDB
db2 start hadr on database GSDB as primary
```



IBM Spectrum Control integration

This chapter describes how the Spectrum Control software suite complements the built-in functions of the VersaStack hardware components. It covers how the IBM Tivoli Productivity Center SmartCloud Virtual Storage Edition and IBM FlashCopy Manager offer the following functions:

- ▶ Cloud-enabled, pro-active, and event driven storage management
- ▶ Real-time performance monitoring, historical data analysis, and automated reporting
- ▶ Advanced data protection technologies that use hardware-assisted snapshots
- ▶ VMware vCenter and VMware Web Client integration

12.1 Spectrum Control overview

This section describes the components of IBM SmartCloud Virtual Storage Center (VSC) that are applicable to the VersaStack setup. This section also describes the VSC offerings and licensing model overview.

VSC V5.2 provides efficient virtualization, management, and data protection for heterogeneous storage environments. VSC helps IT storage managers migrate to an agile cloud-based storage environment and manage it effectively without having to replace existing storage systems. This powerful offering removes the physicality of storage, and also the complexity that is associated with managing multivendor infrastructures.

VSC V5.2 offers a storage virtualization platform, capabilities for storage virtualization management, and instant copy management. VSC V5.2 delivers to customers, under one licensed software product, the complete set of functions that are available in the IBM Tivoli Storage Productivity Center, the functions and capabilities that are associated with the IBM System Storage SAN Volume Controller (including copy services), and the capabilities of the IBM Tivoli Storage FlashCopy Manager. With VSC, you can now get all of the advanced capabilities of what was previously Tivoli Storage Productivity Center Standard Edition, and with the VSC V5.2 license, you get all of the advanced analytics functions. This powerful solution enables organizations to optimize provisioning, capacity, availability, reporting, and management for virtual storage.

12.2 Storage hypervisor

This section introduces the concepts of *server hypervisor* and *storage hypervisor*. It also has an overview of the IBM Storage Hypervisor, which is integrated with the VSC V5.2.

Server hypervisor

In cloud computing, a *server hypervisor* has the following key attributes, which provide effective resource utilization, cost savings, and flexibility to the business:

- ▶ Pooled physical resources are consumed by virtual machines (VMs), resulting in high asset utilization.
- ▶ VMs are mobile, giving administrators their choice of physical server and location.
- ▶ A common set of value capabilities and centralized management are provided for VMs, regardless of what physical server on which they are running.

Storage hypervisor

A *storage hypervisor* is a rapidly emerging way of describing the same value aspects, but in a storage context:

- ▶ Consolidation and cost: Storage pooling increases utilization and decreases costs.
- ▶ Business availability: Data mobility of virtual volumes can improve availability.
- ▶ Application support: Tiered storage optimization aligns storage costs with required application service levels.

IBM Storage Hypervisor

The IBM Storage Hypervisor offers the following features (shown in Figure 12-1):

- ▶ Virtualizes storage resources from multiple arrays, vendors, and datacenters, which are pooled together and accessed from anywhere.
- ▶ Standardized storage services are selected from a service catalog.
- ▶ Storage volumes move dynamically based on workload balancing policies.
- ▶ Self-service provisioning uses automation to allocate capacity.
- ▶ Pay-per-use storage resources, so users are aware of the impact of their consumption and service-level choices.




Figure 12-1 IBM Storage Hypervisor

IBM Storage Hypervisor is part of VSC V5.2, which includes storage virtualization, storage virtualization management, and storage snapshot management that are tightly integrated with advanced analytics to deliver a robust storage cloud solution. This solution ultimately helps businesses to optimize provisioning, capacity, availability, data protection, reporting, and management for virtualized storage.

12.3 IBM SmartCloud Virtual Storage Center component model

As shown in Figure 12-2, VSC V5.2 includes core functions from three IBM offerings: Storage management through IBM Tivoli Storage Productivity Center, storage virtualization through IBM System Storage SAN Volume Controller, and application-aware data protection through IBM Tivoli Storage FlashCopy Manager.

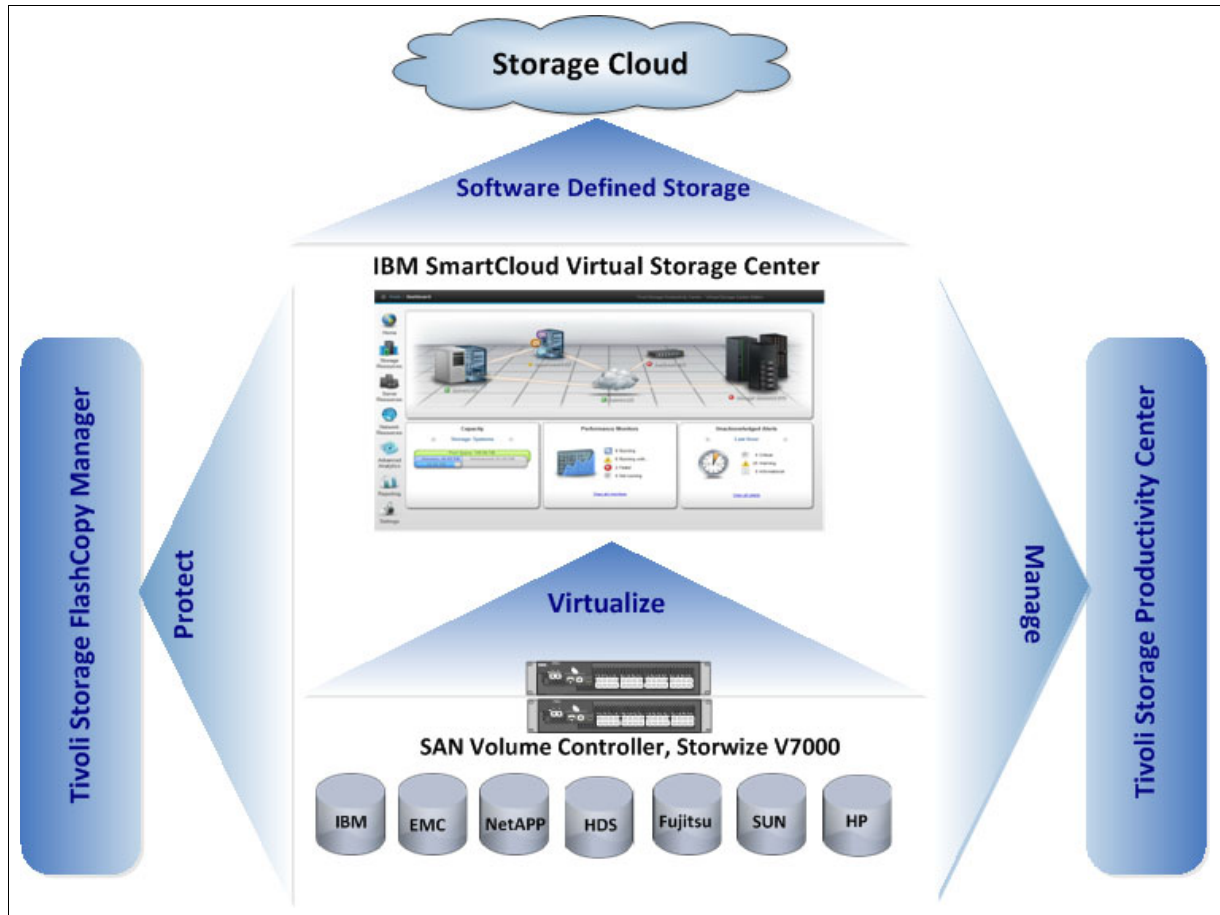


Figure 12-2 Overview diagram of IBM SmartCloud Virtual Storage Center

12.3.1 Storage management

The storage management component in VSC V5.2 provides advanced storage infrastructure and data management capabilities. The Tivoli Storage Productivity Center component that is available in VSC includes all the capabilities of Tivoli Storage Productivity Center V5.2. It uniquely provides all the advanced functions that were available in the past as part of Tivoli Storage Productivity Center Standard Edition and Tivoli Storage Productivity Center for Replication. Unique to the VSC V5.2 Storage Analytics Engine is data management with file system and database scanning and analysis, data placement, user quotas, and an advanced management GUI to help simplify virtual storage administration.

The storage management component of the VSC solution improves visibility, control, and automation for data and storage infrastructures, including storage systems, devices, and SAN fabrics, and is integrated with SAN Volume Controller functions for auto-tiering and workload-aware placement across the datacenter.

Tivoli Storage Productivity Center, the storage management component of VSC, helps simplify provisioning, performance management, and data replication processes (Figure 12-3).

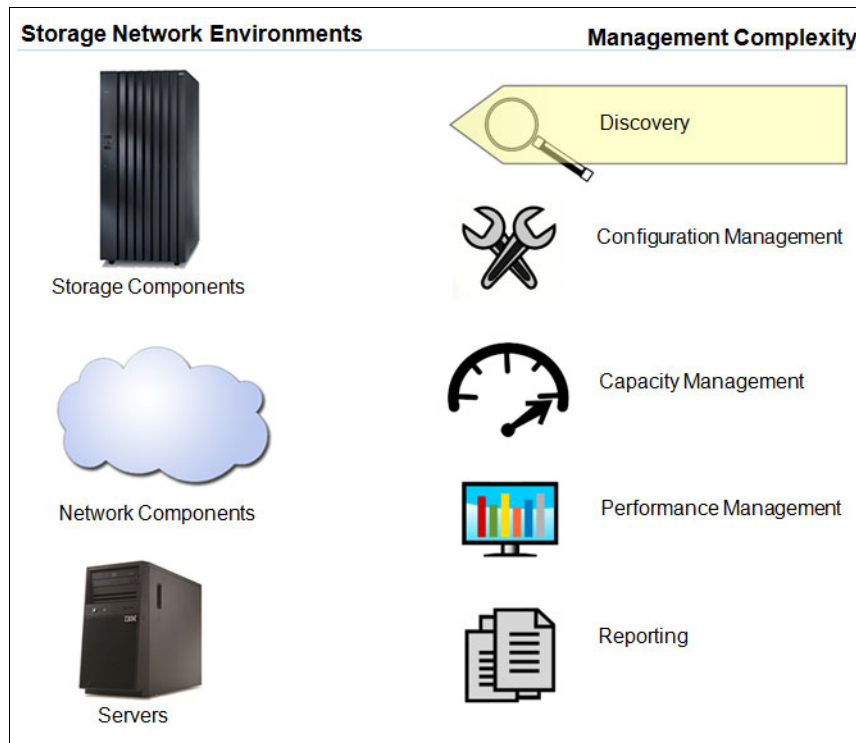


Figure 12-3 IBM SmartCloud Virtual Storage Center storage management and Tivoli Storage Productivity Center

IBM Tivoli Storage Productivity Center provides these capabilities, all from a single GUI:

- ▶ Database, host, file-system, and file-level capacity analytics
- ▶ Storage performance management
- ▶ Tiered storage analysis
- ▶ Trend analysis
- ▶ SAN planning and provisioning
- ▶ Performance optimization
- ▶ SAN fabric performance management

Note: For more information about VSC offerings and licensing, see 12.6, “IBM SmartCloud Virtual Storage Center offerings” on page 274.

Tivoli Storage Productivity Center can generate threshold alerts and forward them to SNMP receivers. Tivoli Storage Productivity Center provides many ready-to-use reports, as shown in Figure 12-4.

- Public Folders
 - IBM Tivoli Storage Productivity Center Predefined Reports
 - Fabrics and Switches
 - Switches
 - Performance of One Switch
 - Compare Performance of Multiple Switches
 - Compare Performance of One Switch over Time Ranges
 - Switch Ports
 - Most Active Switch Ports
 - Performance of One Switch Port
 - Compare Performance of Multiple Switch Ports
 - Compare Performance of One Switch Port over Time Ranges
 - Hypervisors
 - Hypervisors Capacity
 - Most Active Hypervisors
 - Summarized Performance of Volumes by Hypervisor
 - Performance of Volumes by Hypervisor
 - Hypervisor Data Stores Capacity
 - Hypervisor Disks Capacity
 - Servers
 - Filesystems
 - File Systems Capacity
 - Servers Capacity
 - Most Active Servers
 - Summarized Performance of Volumes by Server
 - Performance of Volumes by Server
 - Server Disks Capacity
 - Storage Systems
 - Controllers, Modules, or Nodes
 - Most Active Controllers or Modules
 - Performance of One Controller or Module
 - Most Active Nodes
 - Performance of One Node
 - Compare Performance of Multiple Controllers and Modules
 - Compare Performance of Multiple Nodes
 - Disks
 - Disks Capacity
 - Most Active Disks
 - Performance of One Disk
 - Compare Performance of Multiple Disks
 - Host Connections
 - Most Active Host Connections
- Summarized Performance of Volumes by Host Connection
 - Performance of Volumes by Host Connection
 - IO Groups
 - Most Active IO Groups
 - Performance of One IO Group
 - Compare Performance of Multiple IO Groups
 - Managed Disks
 - Managed Disks Capacity
 - Most Active Managed Disks
 - Performance of One Managed Disk
 - Compare Performance of Multiple Managed Disks
 - Pools
 - Pools Capacity
 - Pools Historical Capacity
 - Most Active Pools
 - Performance of One Pool
 - Compare Performance of Multiple Pools
 - Compare Performance of One Pool over Time Ranges
 - Ports
 - Most Active Ports
 - Performance of One Port
 - Compare Performance of Multiple Ports
 - RAID Arrays
 - Most Active RAID Arrays
 - Performance of One RAID Array
 - Compare Performance of Multiple RAID Arrays
 - Volumes
 - Volumes Capacity
 - Volumes Historical Capacity
 - Most Active Volumes
 - Performance of One Volume
 - Compare Performance of Multiple Volumes
 - Storage Systems Capacity
 - Storage Systems Historical Capacity
 - Most Active Storage Systems
 - Performance of One Storage System
 - Compare Performance of Multiple Storage Systems
 - Compare Performance of One Storage System over Time Ranges
 - Performance Data Export
 - Storage Resource Relationships Summary
 - Storage Resource Relationships Summary (Configurable)
 - File System to Volume Relationships

Figure 12-4 Tivoli Storage Productivity Center reports that are ready to use

These reports can be scheduled to run periodically. Additional custom reports can be created with IBM Cognos®. For more information about IBM Cognos reports, see *IBM Tivoli Storage Productivity Center V5.1 Technical Guide*, SG24-8053.

12.3.2 Storage virtualization

The IBM SAN Volume Controller virtualization engine moves the storage control function into the storage network, allowing disk storage to be managed as a single virtual pool, which supports many disk vendors (Figure 12-5).

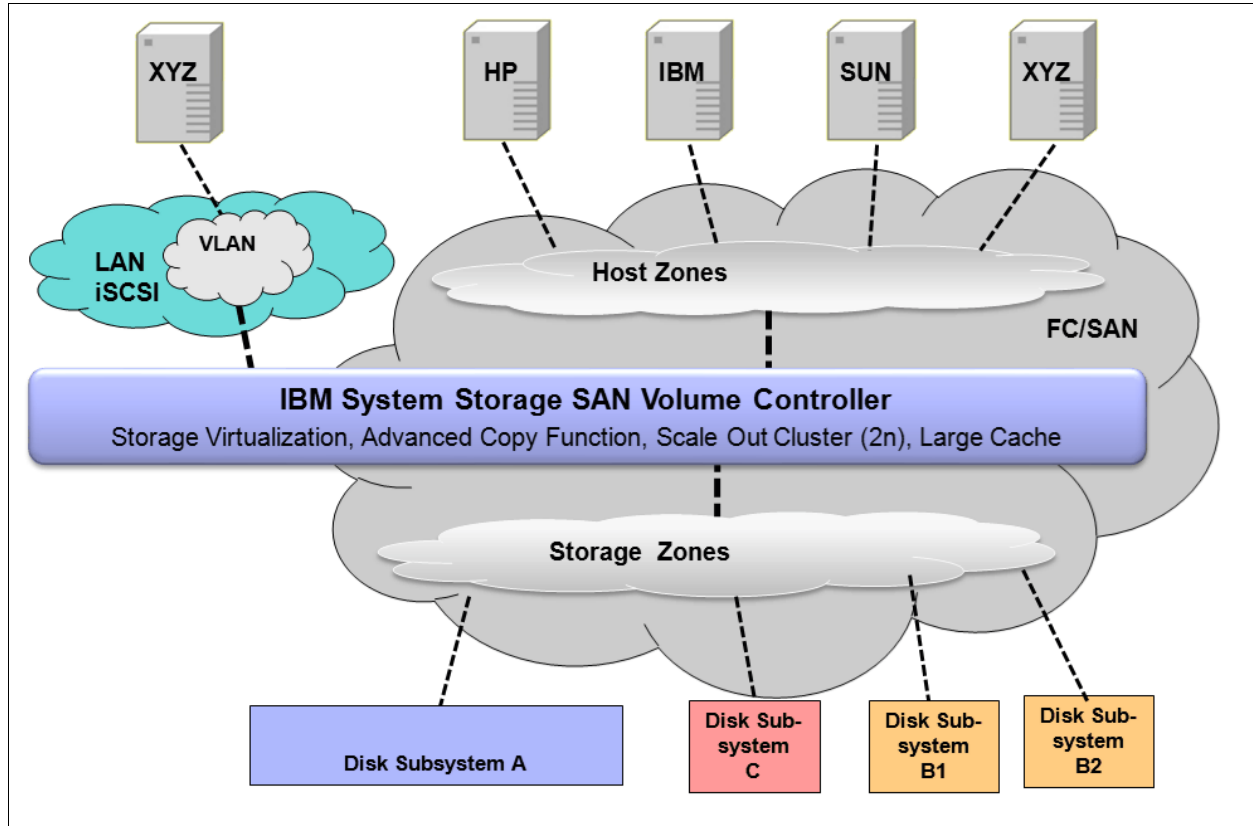


Figure 12-5 SAN Volume Controller conceptual and topology overview

SAN Volume Controller can pool storage volumes into a reservoir of capacity for centralized management. Virtualization with SAN Volume Controller eliminates the boundaries among disk and flash systems, which simplifies management and enables IT operations to focus on managing storage as a resource to meet business requirements rather than as a set of boxes. The RAID array from an external storage system or from internal disks (a Storwize V7000 storage system, as shown in our example in Figure 12-5) is presented to a SAN Volume Controller or Storwize V7000 storage system as *Managed Disks (MDisks)*. A set of MDisks forms a storage pool from which extents are taken to create the volumes, which can be identified by logical unit numbers (LUNs). The volumes, now in virtualized mode, are presented to the hosts. In this sense, the hosts no longer see the back-end disks directly, and the SAN Volume Controller or Storwize V7000 storage system behaves like a controller that is provisioning LUNs to the hosts.

To achieve multi-tenancy over the same physical SAN infrastructure, storage pools can be created that are specific to each tenant from a specific set of managed disks and assign them to the specific tenant hosts, as shown in Figure 12-6.

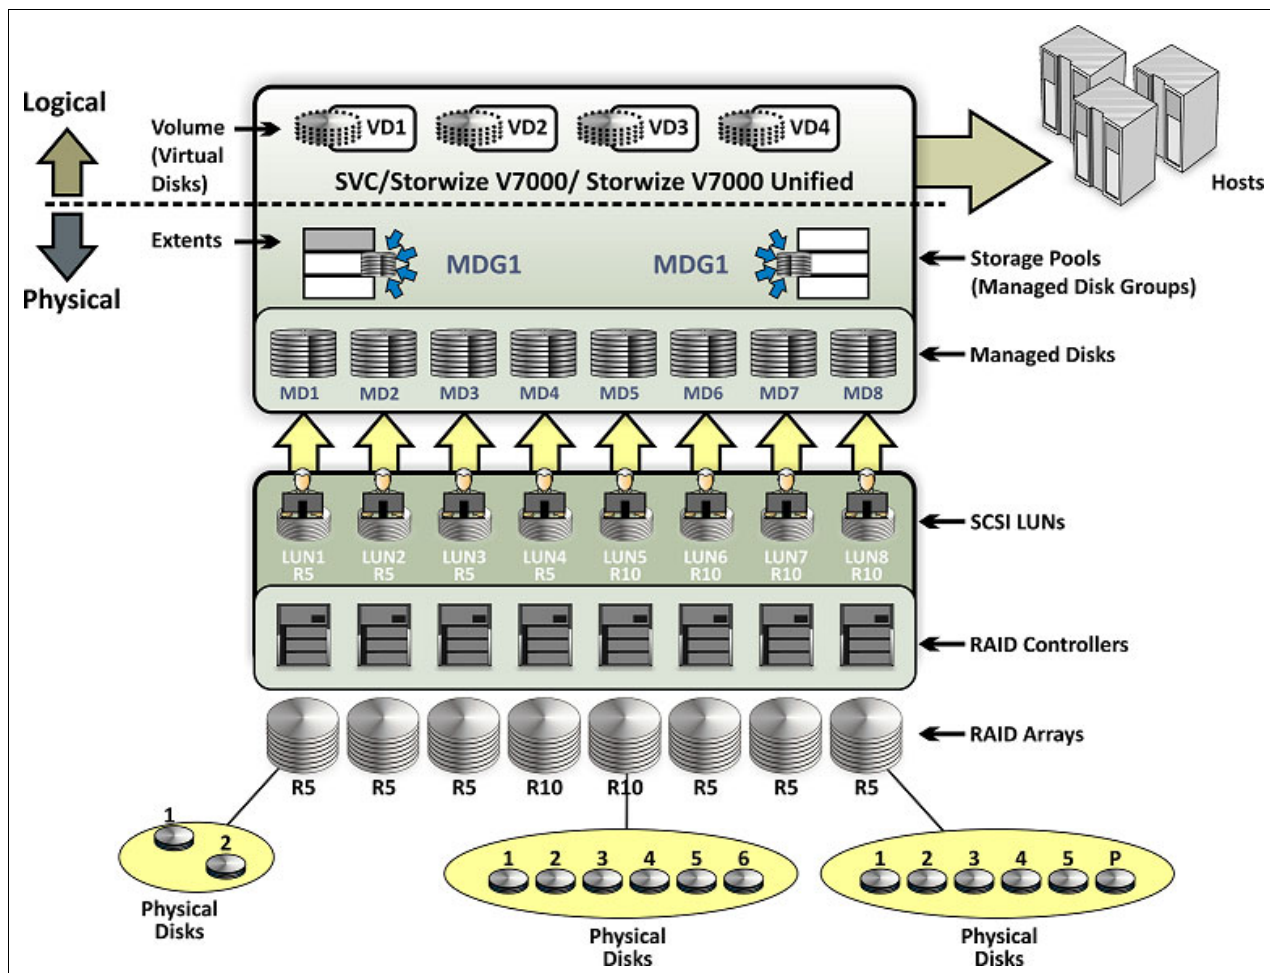


Figure 12-6 SAN Volume Controller storage virtualization concepts summary

The SAN Volume Controller component of VSC reduces labor, reduces and removes planned migration outages, and improves utilization. Storage virtualization with IBMSAN Volume Controller supports a heterogeneous, multivendor environment, with common management and services. SAN Volume Controller allows for nondisruptive changes to the storage environment without impacting host applications. SAN Volume Controller with Infrastructure Lifecycle Management (ILM) intelligent storage analytics provides policy-based automated data placement and tier movement.

Here are the key characteristics of SAN Volume Controller:

- ▶ Highly scalable: A SAN Volume Controller *cluster* scales horizontally through the addition of node pairs to a maximum of four node pairs (or eight nodes) per cluster.
- ▶ Host-independent: Supports multiple operating systems, including Windows, Linux, IBM AIX®, HP-UX, and so on.
- ▶ Storage controller-independent: Supports storage from multiple vendors, including IBM, EMC, HDS, Oracle, Hewlett-Packard, and others.

SAN Volume Controller offers the following services:

- ▶ Creation and management of storage pools that are attached to the SAN.
- ▶ Block-level virtualization.
- ▶ Provisioning of advanced functions across the SAN, such as advanced copy services (point-in-time copy, instant copy, synchronous remote copy, Metro Mirror and asynchronous remote copy, and Global Mirror).
- ▶ Thin provisioning.
- ▶ Real-time compression: The IBM Real-time Compression option can be added as a separately priced license. For more information about this topic, see 12.6, “IBM SmartCloud Virtual Storage Center offerings” on page 274.
- ▶ Data migration: Move volumes within or between storage controllers (within the same physical virtualization boundary).
- ▶ Growing or shrinking volumes.
- ▶ IBM Easy Tier helps administrators control storage growth more effectively by balancing MDisks within a pool, and by moving low-activity or inactive data into a hierarchy of lower-cost storage. Administrators can free disk space on higher-value storage for more important, active data.

The SAN Volume Controller is incorporated into the IBM Spectrum™ family as Spectrum Virtualize and is incorporated in the IBM Storwize V7000 storage system that is part of the VersaStack offering.

12.3.3 Application-aware data protection

With the Tivoli Storage FlashCopy Manager component of VSC, the data backup and restore component in IBM SmartCloud VSC V5.2 provides fast application-aware backups and restores by using advanced snapshot technologies that are available with IBM storage systems (Figure 12-7).

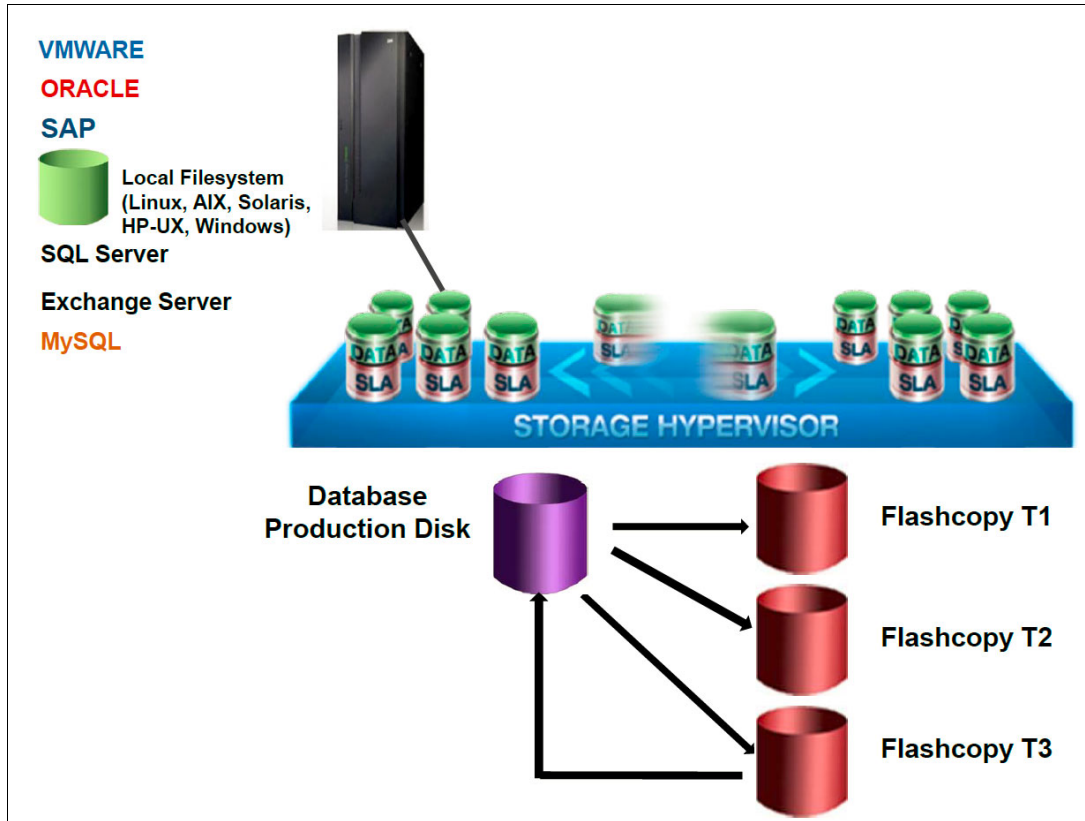


Figure 12-7 High-level overview of Tivoli Storage FlashCopy Manager application-aware copy solution

FlashCopy Manager uses advanced IBM storage hardware snapshot technology to help create a high-performance, low-impact application data protection solution. The Storage FlashCopy function operates at the level of virtual volumes (VDisks), that is, it copies whole volumes. The FlashCopy function is intended to create copies of data that may be used for purposes such as disk-to-disk backups, parallel processing (multiple applications processing different copies of the same data), and testing by using a copy of production data. The copy that is created by the FlashCopy function can be used almost immediately. SAN Volume Controller within the Storwize V7000 storage system can perform a background copy of all data from the source to the target or it can copy data only when an update occurs. It delivers high levels of data protection for business-critical applications through integrated application snapshot backup and restore capabilities.

Storage administrators can control the speed of the background copy to limit the impact that the copy has on other SAN Volume Controller activities. The first time that FlashCopy is used, the copy takes place as *normal*, which means that a full copy of data occurs from the source to the target VDisk. When changes are made, only the changes are copied to the target. A total of 256 copies can be made from the same source VDisk, which can be incremental or non-incremental, or a mix of both.

FlashCopy integrates with IBM System Storage DS8000®, FlashSystem V840, IBM SAN Volume Controller, IBM Storwize V7000 and V5000 storage systems, and IBM XIV® Storage System products. For Microsoft Windows environments, Tivoli Storage FlashCopy Manager also supports other hardware that can perform Microsoft Volume Shadow Copy Services (VSS) functions, such as IBM N Series, and IBM System Storage DS3000, DS4000®, and DS5000™ storage systems.

Here is an explanation of the FlashCopy Manager solution (see Figure 12-7 on page 264):

1. Starting at the left, the Application system, which is also called the Production system, has the production database on it. More important, this is the data that you want to protect. The applications can be Oracle, SQL, IBM DB2, SAP, Exchange, files systems, or VMware. FlashCopy Manager also provides the interfaces for custom applications to take snapshots of the data.
2. Following the black arrow, you see that the application data is on the LUN that sits on the SAN Volume Controller and on its back-end storage. Using FlashCopy Manager, when you take the backup of the database, you have local snapshot versions that represent the application data at some point. When you want to restore the data, use FlashCopy Manager to restore from any one of these snapshot versions, including the latest of a point-in-time snapshot. If you have Tivoli Storage Manager, you can then offload your backups to Tivoli Storage Manager and manage your data through Tivoli Storage Manager and FlashCopy Manager.

12.4 IBM SmartCloud Virtual Storage Center features

VSC helps reduce storage administration complexity and costs in the following ways:

- ▶ Improving storage utilization
- ▶ Making better use of existing storage and controlling storage growth expenditures
- ▶ Improving application availability and simplified data migrations
- ▶ Making changes to storage and moving data without taking down applications
- ▶ Simplifying storage management
- ▶ Improving efficiency and productivity for storage management staff
- ▶ Providing advantages with a software-defined storage architecture model
- ▶ Enabling greater choice (lower cost) when buying storage and lowering software costs
- ▶ Improving application recovery time and recovery point objectives (RTO and RPO)
- ▶ Providing application-aware hardware-based snapshots
- ▶ Providing network-based replication

Here are the outstanding features of IBM SAN Volume Controller:

- ▶ Efficient by design
- ▶ Self-optimizing
- ▶ Cloud agility

12.4.1 Efficient by design

Organizations must spend less of their IT budgets on storage and storage administration so that they can spend more on new, revenue-generating initiatives. VSC has built-in efficiency features that help users avoid purchasing add-ons or additional licenses or deal with complicated integration issues.

VSC has these advanced efficiency features:

- ▶ **Storage virtualization**

This is a foundational technology for clouds and software-defined environments (SDEs). Without virtualization, storage capacity utilization averages about 50%, but virtualized storage enables up to 90% utilization by enabling online data migration for load balancing. With VSC, you can virtualize your storage resources from multiple storage systems and vendors. You can use pooling storage devices to access capacity from any storage system, which is a significant advantage over the limitations that are inherent in traditional storage.
- ▶ **Simplified user experience**

VSC provides an advanced GUI and a VMware vCenter plug-in to reduce administration complexity. Administrators can do common tasks consistently, over multiple storage systems, even those from different vendors. The IBM storage GUI enables simplified storage provisioning with intelligent presets and embedded preferred practices, and integrates context-sensitive performance management throughout.
- ▶ **Near-instant, application-aware backup and restore**

To reduce downtime in high-availability virtual environments, critical applications such as email and databases require near-instant backups that have little or no impact on application performance. Application-aware snapshot backups can be performed frequently throughout the day to reduce the risk of data loss. VSC simplifies administration and recovery from snapshot backups.

12.4.2 Self-optimizing

Self-optimizing storage adapts automatically to workload changes to optimize application performance, eliminating most manual tuning efforts. IBM SmartCloud VSC includes these self-optimizing features:

- ▶ **IBM Tiered Storage Optimizer**

VSC uses performance metrics, advanced analytics, and automation to enable storage optimization on a large scale. It can optimize storage volumes across different storage systems and VM vendors and brands. The Tiered Storage Optimizer feature can reduce the unit cost of storage by as much as 50%, based on deployment in a large IBM datacenter.
- ▶ **IBM Easy Tier**

VSC helps optimize flash storage with automated tiering for critical workloads. Easy Tier helps make the best use of available storage resources by automatically moving the most active data to the fastest storage tier, which helps applications and virtual desktop environments run up to three times faster.
- ▶ **Thin provisioning and efficient remote mirroring**

Thin provisioning helps automate provisioning and improve productivity by enabling administrators to focus on overall storage deployment and utilization, and also on longer-term strategic requirements without being distracted by routine storage-provisioning requests. IBM Metro Mirror and Global Mirror functions automatically copy data to remote sites as it changes, enabling fast failover and recovery. These capabilities are integrated into the advanced GUI so that they become easier to deploy.

12.4.3 Cloud agility

Cloud computing is all about agility. Storage for clouds must be as flexible and service-oriented as the applications it supports. VSC can convert existing storage into a private storage cloud with no “rip and replace” required. You can use the solution to adapt to the dynamic storage needs of cloud applications by providing storage virtualization, automation, and integration for cloud environments.

Here are the agile features of the solution:

- ▶ OpenStack cloud application provisioning

VSC includes an OpenStack Cinder volume driver that enables automated provisioning by using any of the storage systems that are controlled by VSC. OpenStack cloud applications can access multiple storage tiers and services, without added complexity.

- ▶ Self-service portal

VSC can provide provisioning automation for self-service storage portals (such as IBM SmartCloud Storage Access), which enable immediate responses to service requests while eliminating manual administration tasks.

- ▶ Pay-per-use invoicing

VSC integrates with IBM SmartCloud Cost Manager and other chargeback systems to enable flexible usage accounting for storage resources. VSC can become the single source for usage metrics across storage area networks (SANs), network-attached storage, and direct-attached storage.

12.5 IBM SmartCloud Virtual Storage Center interfaces

IBM focuses on supporting four SDEs, as shown in Figure 12-8.

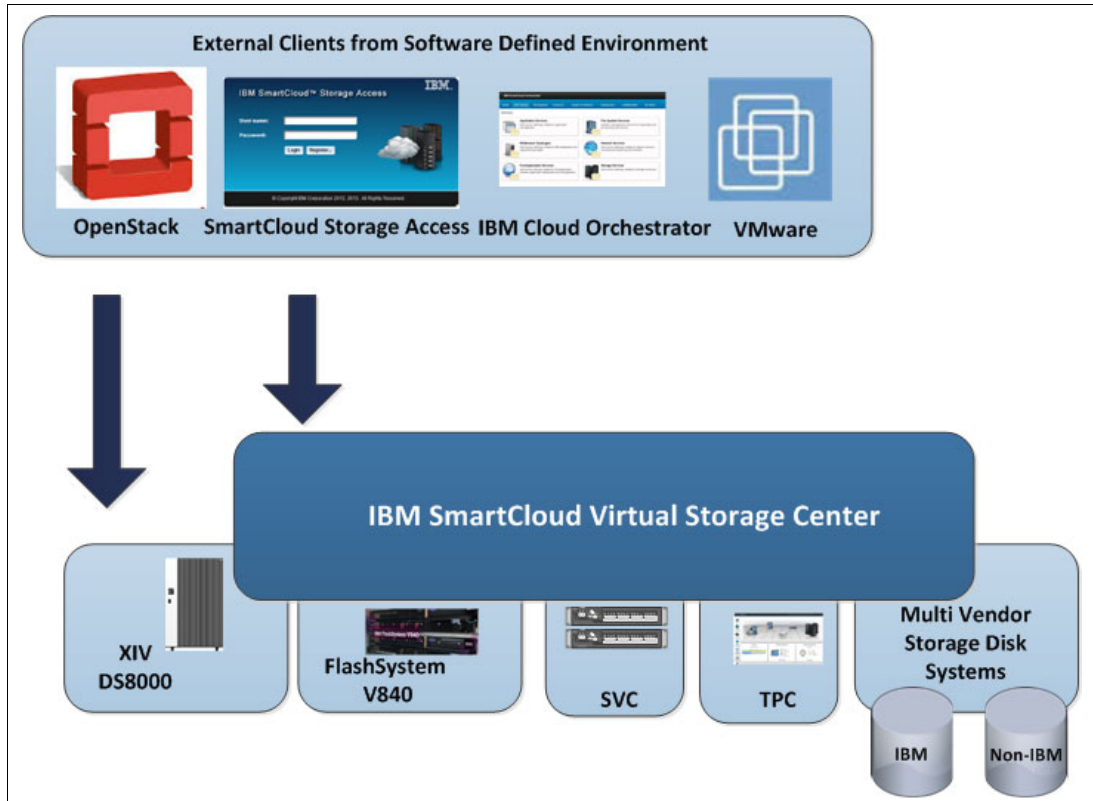


Figure 12-8 Interfaces to IBM SmartCloud Virtual Storage Center

Here are the four SDEs:

- ▶ OpenStack open source code can manage compute, network, and storage resources.
- ▶ IBM SmartCloud is based on OpenStack with added value IBM proprietary features.
- ▶ IBM Cloud Orchestrator is for storage provisioning, orchestration, and automation.
- ▶ VMware runs primarily on x86-based servers.

The interface features are summarized in Table 12-1.

Table 12-1 Comparison of virtual storage interfaces to a software-defined environment

OpenStack software	IBM SmartCloud and IBM Orchestrator	VMware
IBM is a platinum sponsor of the OpenStack Foundation.	IBM Cloud Manager with OpenStack is based on OpenStack open source code, with added value proprietary features from IBM.	VMware is entirely proprietary, but has a large market share for the x86-based server infrastructure.
OpenStack open source code can manage IBM compute, network, and storage resources.	IBM Cloud Manager with OpenStack and IBM Cloud Orchestrator support various server hypervisors and interfaces.	IBM was VMware's first OEM and joint development partner (since 1998). IBM continues this strong partnership. IBM Global Services is one of VMware's largest customers, using VMware in many of their client solutions.
IBM offers Cinder interfaces on most of its major storage products for block storage access and supports Swift interfaces for object storage access.	IBM SmartCloud Storage Access and IBM Cloud Orchestrator provide self-provisioning and orchestration capabilities.	VMware vStorage API for data protection (VADP), VMware Site Recovery Manager (SRM), VMware vSphere storage APIs: Array integration (VAAI), VMware vCenter.

An overview of the VMware VSC interface that is being used in the VersaStack setup is provided in 12.5.1, "VMware" on page 269.

12.5.1 VMware

VMware provides server virtualization on an Intel based architecture. Here are the core components of the VMware solution:

- ▶ VMware ESX and ESXi based hypervisor
- ▶ VMware vSphere vCenter for providing management capabilities
- ▶ vSphere vMotion to combat planned downtime
- ▶ VMware vCenter Site Recovery Manager to automate end-to-end recovery processes for virtual applications

Figure 12-9 shows the vSphere suite in a more comprehensive way. vSphere is a product suite that is similar to the Microsoft Office suite, which contains Microsoft Office Word, Excel, Access, PowerPoint, and so on. VMware vSphere suite includes an ESXi hypervisor, vCenter, and vSphere client. ESXi is a hypervisor that is installed on a physical machine. The vSphere client is installed on the VMware administrator's notebook or desktop computer and is used to access the ESXi server to install and manage VMs on the ESXi server. The vCenter server is installed as a VM on top of the ESXi server. The vCenter server is a vSphere component that is mostly used in a large environment where there are many ESXi servers and several VMs. The vCenter server can also be accessed by vSphere client for management purposes. So, the vSphere client is used to access the ESXi server directly in a small environment; in a larger environment, the vSphere client is used again to access the vCenter server, which ultimately manages the ESXi server.

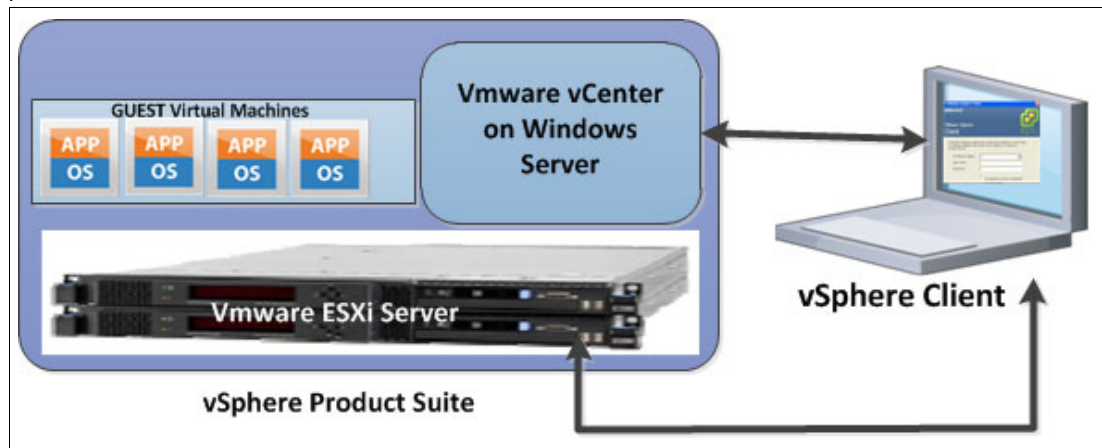


Figure 12-9 VMware vSphere Suite overview

VMware servers hypervisor

VMware ESX and VMware ESXi are hypervisors that you can use to abstract processor, memory, storage, and networking resources into multiple VMs that can run unmodified operating systems and applications. VMware ESX and VMware ESXi reduce server sprawl by running applications on VMs that consist of fewer physical servers. VMware ESX and VMware ESXi hosts can be organized into clusters. This configuration allows ESX to provide flexibility in terms of what VMs are running on what physical infrastructure.

VMware vCenter

vCenter is the management software suite that is used to manage the VMs inside an ESX or ESXi host. When you allocate resources such as memory, storage, networking, or processors to a VM, a vCenter server manages how these resources are allocated and maintained. vCenter can manage a single ESX or ESXi hosts and clusters of hosts. VMware vCenter has several features that allow for mobility of VMs between ESX hosts and storage. These features can add to the availability of the VMs running in a cluster.

VMware vMotion

vMotion is a technology that combats planned downtime. vMotion is used to move VMs between host and datastores to allow scheduled maintenance procedures to proceed without affecting VM availability or performance. It is included in the Enterprise and Enterprise Plus versions of VMware vSphere. It is shown in Figure 12-10.

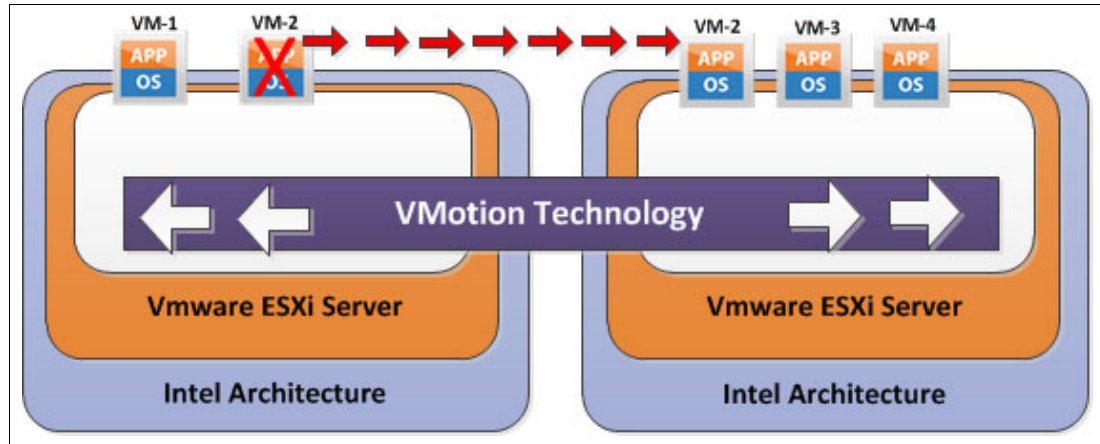


Figure 12-10 VMware vMotion

VMware Host vMotion

Host vMotion eliminates the need to schedule application downtime for planned server maintenance. It does so through live migration of VMs across servers with no disruption to users or loss of service.

This process is managed from a vCenter server, which maintains client or application access to a VM while it is moving between physical servers. In a SAN Volume Controller stretched cluster solution, this feature is useful for moving VMs between two failure domains. You might need to move VMs to load-balance across failure domains or because a failure domain needs an outage for maintenance.

VMware Storage vMotion

Storage vMotion eliminates the need to schedule application downtime because of planned storage maintenance or during storage migrations. It does so by enabling live migration of VM disks (VMDKs) with no disruption to users or loss of service. The vCenter server manages the copy of data from one datastore to another. With vStorage APIs for Array Integration (VAAI), this process can be offloaded to the storage system, saving resources on both the vCenter host and data network.

Figure 12-11 illustrates the use of VMware Storage vMotion in a SAN Volume Controller stretched cluster solution.

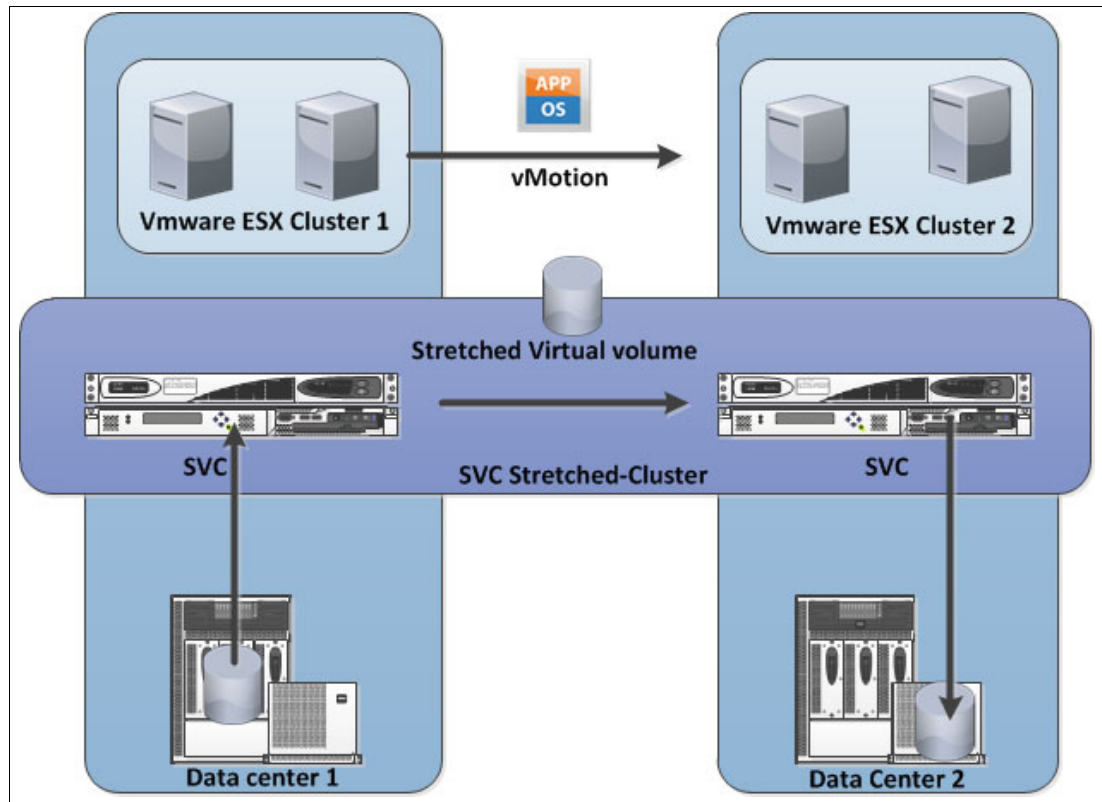


Figure 12-11 VMware Storage vMotion

In a SAN Volume Controller stretched cluster solution, this feature is useful for moving a VM's VMDK file between two systems. You might move this file to ensure that it is on the same failure domain as the VM, or to migrate off a storage device that is becoming obsolete or is undergoing maintenance, as shown in Figure 12-11.

For more information about Storage vMotion, see the following website:

<http://www.vmware.com/files/pdf/VMware-Storage-VMotion-DS-EN.pdf>

VMware vCenter Site Recovery Manager

Site Recovery Manager integrates with VMware vCenter server, and underlying storage replication products, to automate end-to-end recovery processes for virtual applications. It provides a simple interface for setting up recovery plans that are coordinated across all infrastructure layers. Recovery plans can be tested non-disruptively as frequently as required to ensure that the plan meets availability objectives. At the time of a failure domain failover or migration, Site Recovery Manager automates both the failover and failback processes. It ensures fast and highly predictable RPOs and RTOs.

For more information about vCenter Site Recovery Manager, see the following website:

<http://www.vmware.com/products/site-recovery-manager/overview.html>

VMware Distributed Resource Scheduler

Distributed Resource Scheduler (DRS) dynamically balances computing capacity across a collection of hardware resources that are aggregated into logical resource pools. It continuously monitors utilization across resource pools and intelligently allocates available resources among the VMs that are based on predefined rules that reflect business needs and changing priorities. When a VM experiences an increased load, VMware DRS automatically allocates more resources by redistributing VMs among the physical servers in the resource pool.

VMware DRS migrates and allocates resources by using a set of user-defined rules and policies. These rules and policies can be used to prioritize critical or high-performing VMs, ensure that particular VMs never run on the same storage or host, or save on power and cooling costs by powering off ESX servers that are not needed.

For more information about Distributed Resource Manager, see the following website:

http://www.vmware.com/pdf/vmware_drs_wp.pdf

VSC and VMware integration

VSC and VMware are integrated by using Tivoli Storage Productivity Center plug-ins, as shown in Figure 12-12.

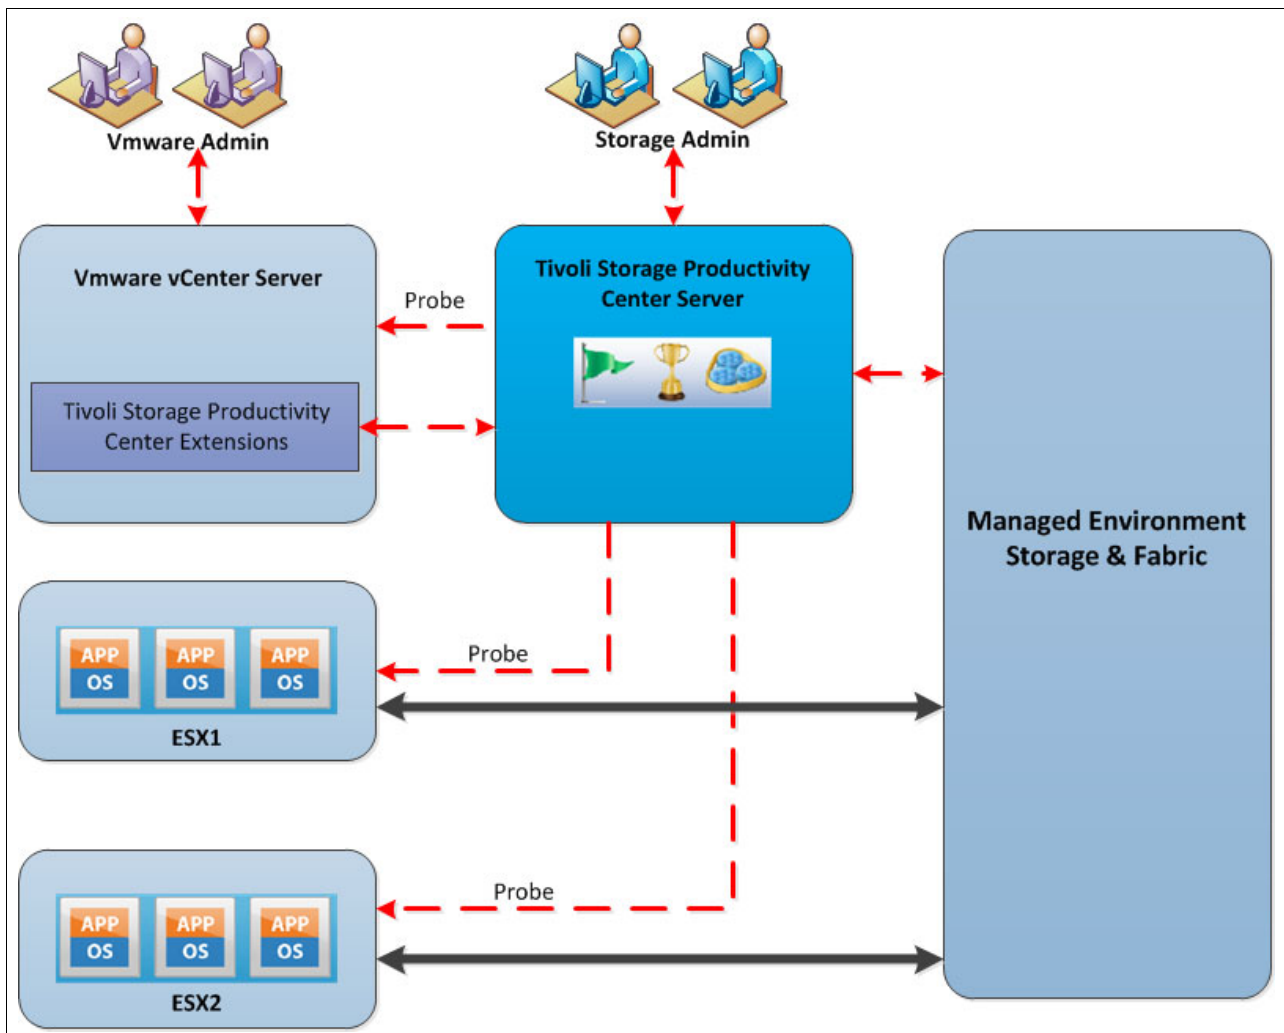


Figure 12-12 VSC and VMware integration topology

The VMware vCenter web client extension provides the following features to VMware administrators:

- ▶ Visualization of connections:
 - End-to-end, from storage volume to VM
 - Storage details, such as pools, volumes, and attributes
 - Performance charts
- ▶ Automated policy-based storage provisioning, based on the storage services catalog:
 - Block volumes
 - File shares
 - Service class characteristics (storage tier, data protection, compression, encryption, and so on)
- ▶ vSphere API for Storage Awareness (VASA):
 - Alerts for performance, errors, and capacity thresholds
 - Availability of volumes, pools, storage systems, and paths
 - Tivoli Storage Productivity Center service classes can be used as VASA capabilities
- ▶ Additional storage reports are available in the vCenter admin GUI:
 - Fabric connectivity
 - Storage performance
 - Storage mappings

12.6 IBM SmartCloud Virtual Storage Center offerings

VSC V5.2 has the following offerings:

- ▶ IBM SmartCloud Virtual Storage Center V5.2
- ▶ IBM SmartCloud Virtual Storage Center Entry V5.2
- ▶ IBM SmartCloud Virtual Storage Center for Storwize Family V5.2

The *VSC V5.2* license is an offering that is used with the SAN Volume Controller and is a software entitlement to run the external virtualization, FlashCopy, and remote copy services features. The only feature of the SAN Volume Controller that is not included in the IBM SmartCloud VSC V5.2 license entitlement is the Real-Time Compression option, which can be added as a separately priced license. This license does not include the hardware nodes that are required for a complete SAN Volume Controller implementation.

IBM SmartCloud VSC Entry V5.2 provides external virtualization, FlashCopy, and remote copy services software entitlement in smaller SAN Volume Controller configurations. Also, for deployment in midrange environments, an IBM Storwize V5000 or V7000 storage system can be used as the virtualization engine in a VSC configuration, and in this case the offering that is used is the *IBM SmartCloud Virtual Storage Center for Storwize Family V5.2*.

The versions of code that are available through IBM SmartCloud VSC 5.2 for download for the SAN Volume Controller and the Tivoli Storage FlashCopy Manager are the same as the versions available for download if these products were downloaded independently of IBM SmartCloud VSC 5.2. In the case of Tivoli Storage Productivity Center, the code is the same as the independent product, but the VSC license enables the Storage Advanced Engine functions to be used.

12.6.1 License model overview

VSC can help customers to migrate easily their storage to a virtual environment and manage storage efficiently. VSC licensing charges are based on the entire managed capacity, which is in contrast to SAN Volume Controller, where FlashCopy and Metro Mirror or Global Mirror can be licensed on virtual capacity for those functions only. The managed capacity model avoids over-provisioning, which can become expensive with SAN Volume Controller. Table 12-2 compares the VSC and Tivoli Storage Productivity Center licensing options and features. The sections after the tables have more details about each of the VSC licenses.

Table 12-2 Current 5.x Virtual Storage Center and Tivoli Storage Productivity Center licensing by offerings

Product name	Licensing usage	Tivoli Storage Productivity Center license	FlashCopy Manager license	SAN Volume Controller license ^a	Storwize license ^b
VSC	Per terabyte (greater than 250 TB or greater than two I/O groups). For example, with the VSC license, you can have 100 TB and grow to 300 TB. This is not possible with VSC Entry, which is limited to less than 250 TB.	Tivoli Storage Productivity Center Advanced	✓	✓	
VSC Entry	Per terabyte (less than 250 TB and less than two I/O groups).	Tivoli Storage Productivity Center Advanced	✓	✓	
VSC for Storwize Family	Per enclosure.	Tivoli Storage Productivity Center Advanced	✓		✓
Tivoli Storage Productivity Center	Per terabyte.	Tivoli Storage Productivity Center			
Tivoli Storage Productivity Center Select	Per enclosure.	Tivoli Storage Productivity Center Select			

a. SAN Volume Controller License includes Base, and FlashCopy and Remote Copy (Metro Mirror and Global Mirror) licenses.

b. The Storwize license that is included in VSC for Storwize Family is for external virtualization only. The base virtualization license must be configured for each Storwize enclosure as usual.

12.6.2 VSC for Storwize Family license

Are you managing a small to medium storage configuration (100 TB - 1 PB) where the storage virtualization investment is largely with Storwize V7000 or Storwize V5000 storage systems, which might manage some variety of storage systems under them? Then, consider using the VSC for Storwize Family license.

VSC for Storwize Family license features

This license offers these features:

- ▶ Restricted to deployment on Storwize V7000 and V5000 hardware.
- ▶ Per enclosure price metric.
- ▶ No restrictions on the number of enclosures.
- ▶ Includes all features of VSC (external virtualization, Mirroring, and Advanced Analytics).
- ▶ The license does not include base software license for Storwize enclosures.

The VSC for Storwize Family license aligns perfectly with the Storwize V7000 component of the VersaStack solution and enhances this offering by providing these functions.

12.7 VersaStack: Spectrum Control

This section demonstrates how we integrated VersaStack components in the example Spectrum Control environment by performing the following actions:

- ▶ Deploy the connections to the hardware infrastructure.
- ▶ Set up and use monitoring and alerting.
- ▶ Enable provisioning to the hypervisor.
- ▶ Create departments and applications to group resources.
- ▶ Monitor and protect the DB2 cluster environment.

12.7.1 Tivoli Productivity Center Virtual Storage Edition installation

The Tivoli Productivity Center Virtual Storage Edition Version 5.2.6 is deployed on a Windows Server 2012 VM running on one of the VMware ESXi hosts in the VersaStack environment.

You can install Tivoli Storage Productivity Center in single-server or multiple-server environments. In a single-server environment, all components are installed on one server.

In a single-server environment, when you install Tivoli Storage Productivity Center, the following components are installed:

- ▶ Database repository
- ▶ Tivoli Storage Productivity Center servers, which are composed of the following components:
 - Data server
 - Device server
 - Alert server
 - Replication server
 - Stand-alone GUI
 - Web-based GUI
 - Command-line interface (CLI)
 - Storage Resource agent
- ▶ Cognos Business Intelligence reports (optional)

In this example, we followed the steps that are outlined at the following website:

http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.6/com.ibm.tpc_V526.doc/fqz0_t_installing_main.html

A field guide that is published on the VSC IBM developerWorks® Wiki provides detailed installation instructions about how to deploy Tivoli Storage Productivity Center V5.2.3 on Windows. You can find it at the following website:

https://www.ibm.com/developerworks/community/wikis/form/anonymous/api/wiki/b6f0fb06-4200-4f2f-9a10-382bddf87c6f/page/f84056cf-76e7-4389-8796-907d9231b2eb/attachment/9d24b843-e00e-4790-b4b5-70e6469fedd0/media/TPC_523_Field_Install_Guide.pdf

The same instructions apply to our Tivoli Productivity Center Virtual Storage Edition V5.2.6.

After the Tivoli Productivity Center Virtual Storage Edition is deployed, you can start the main Web GUI interface, and you are presented with a window that similar to Figure 12-13, which shows the VSC Web GUI with a Storwize V7000 storage system configured.

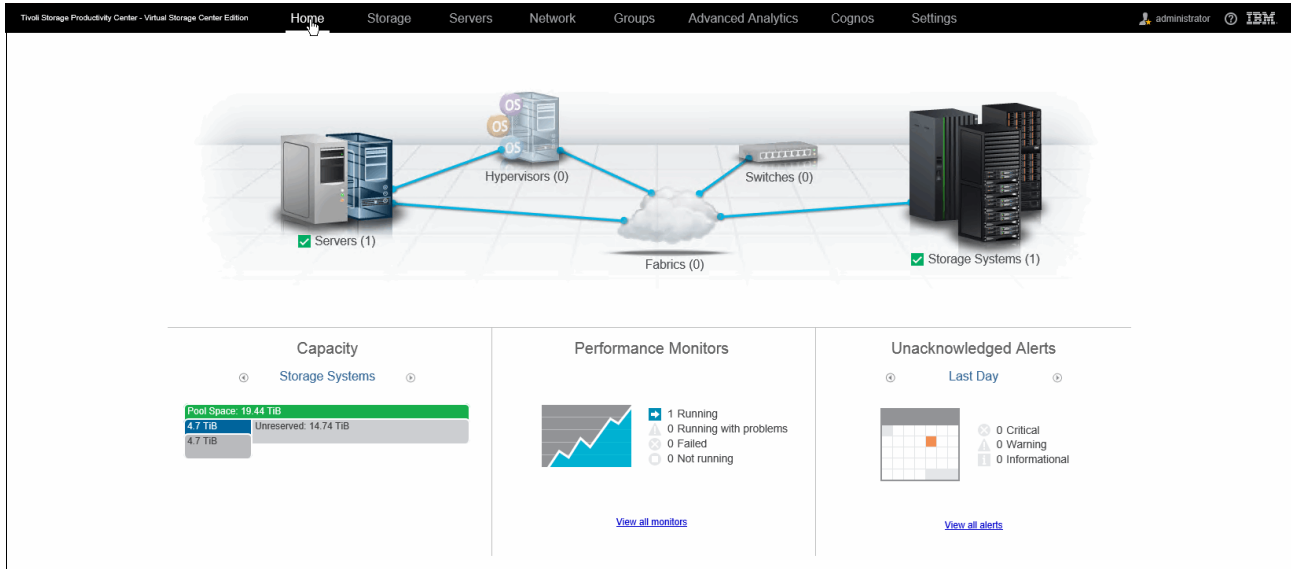


Figure 12-13 Virtual Storage Center Web GUI

By default, a Storage Resource Agent (SRA) is deployed on the system on which the VSC is installed. This agent performs initial SAN discovery.

Note: If the VSC is deployed in a VM, an SRA on physical server with access to the SAN is required to have fabric-based discovery. The configuration that is described in this book is fully virtual. As a result, the fabric and switches are not discovered automatically.

In the subsequent sections, we add the Storwize V7000 storage system and the VMware vCenter hypervisor and deploy SRAs on the cluster members.

12.7.2 Integrating the Storwize V7000 storage system with Spectrum Control

Adding the Storwize V7000 storage system as a new storage device to the VSC follows an easy to use, wizard-driven approach:

1. Within the VSC Web GUI, double-click the **Storage Systems** section, and then click **Add Storage System**.

Figure 12-14 shows the VSC Add Storage System wizard.




Figure 12-14 Add Storage System

2. Click the Storwize Family icon and enter the IP/DNS and login credentials for your Storwize V7000 storage system.

Figure 12-15 on page 279 shows VSC discovering the Storwize V7000 storage system.

Add Storage System

Discover

Storize Family

Host name or IP address: v7000.versastack.local

Authentication: User Name and Password

User name: superuser

Password:

Need Help Back Next Cancel

Figure 12-15 Discover Storize V7000 storage system

- Every device in the VSC environment must be probed at regular intervals for configuration changes. As part of the initial registration, you are prompted to schedule a probe and enable performance monitoring if it is applicable for that specific device.

Figure 12-16 shows VSC scheduling the storage system probe and enabling performance monitoring.

Add Storage System

Configure

Storize Family

Display name: Storize V7000 VersaStack

Location: San Jose

Data Collection

Probe: 16:45 PDT Every day

Run initial probe immediately

Performance monitor: Enabled Every minute

Back Configure Cancel

Figure 12-16 Schedule a probe for a V7000 storage system

- Optionally, you can specify a location where the system is, which allows for logical grouping and classification later.

Figure 12-17 shows that the Storwize V7000 storage system was successfully added through VSC.

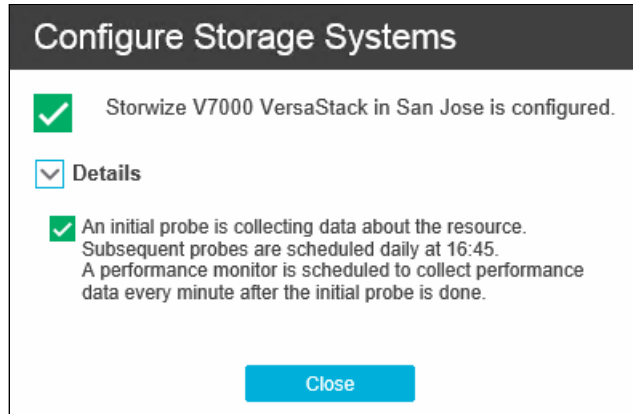


Figure 12-17 Storwize V7000 configuration successful

You are redirected to the Storage Systems section, where the Storwize V7000 storage system is now listed.

Figure 12-18 shows the VSC Storage Systems overview with the Storwize V7000 storage system present.

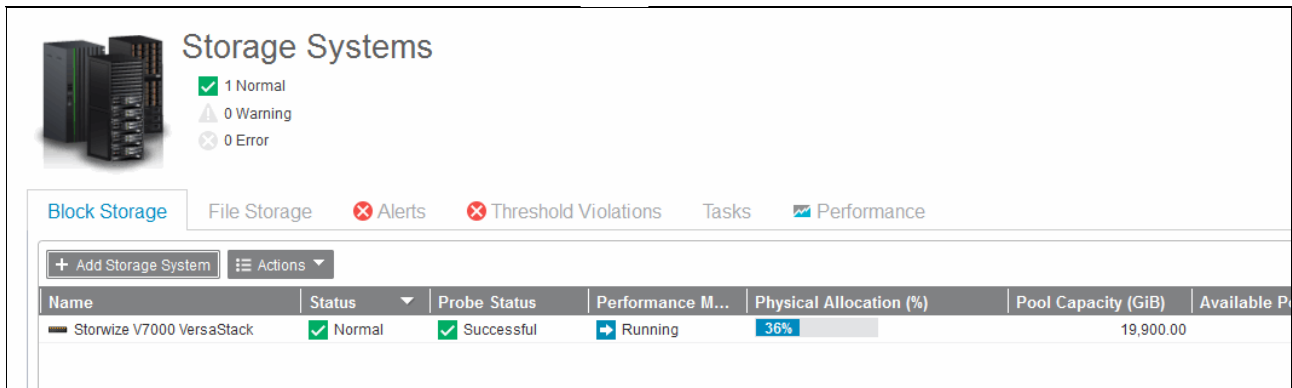


Figure 12-18 Storage Systems overview

Managing the storage infrastructure

Double-clicking the Storwize V7000 VersaStack entry in the VSC Storage Systems pane opens the Overview window, as shown in Figure 12-19 on page 281.




Figure 12-19 Storwize V7000 overview

Throughout the whole VSC GUI, a unified approach is used to chart graphical data and to group resources for the selected device into three categories:

► General

– Overview: This category brings you back to the graphical charts. These charts can be toggled and provide summarized data for the following items:

- Capacity
- Overall System Activity
- Most Active Nodes
- Most Active Volumes
- Most Active Pools
- MDisk Activity
- Space by Host
- Space by Pool
- Space by Volume
- Space by Tier

Figure 12-20 shows the VSC Storwize V7000 Overview with Active Pools, MDisk Activity, Space by Host, and Space by Pool.

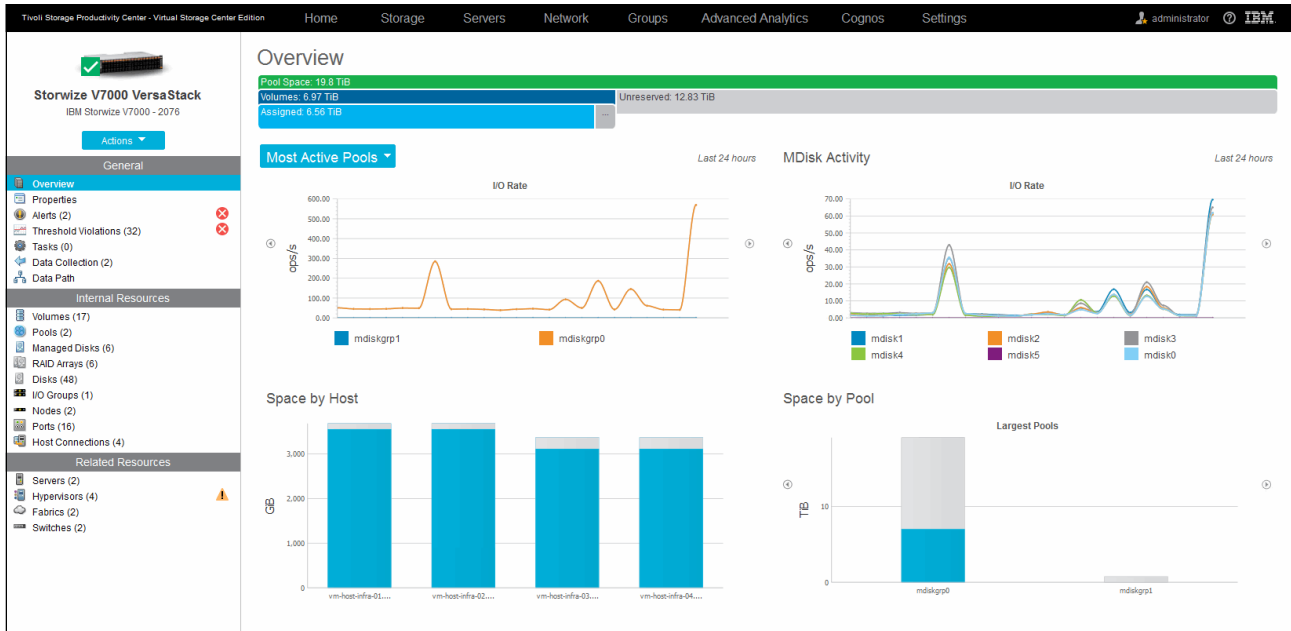


Figure 12-20 Storwize V7000 Overview - Alternative Data Graphs

Figure 12-21 shows the VSC Storwize V7000 Overview with Active Pools, MDisk Activity, Space by Volume, and Space by Tier.

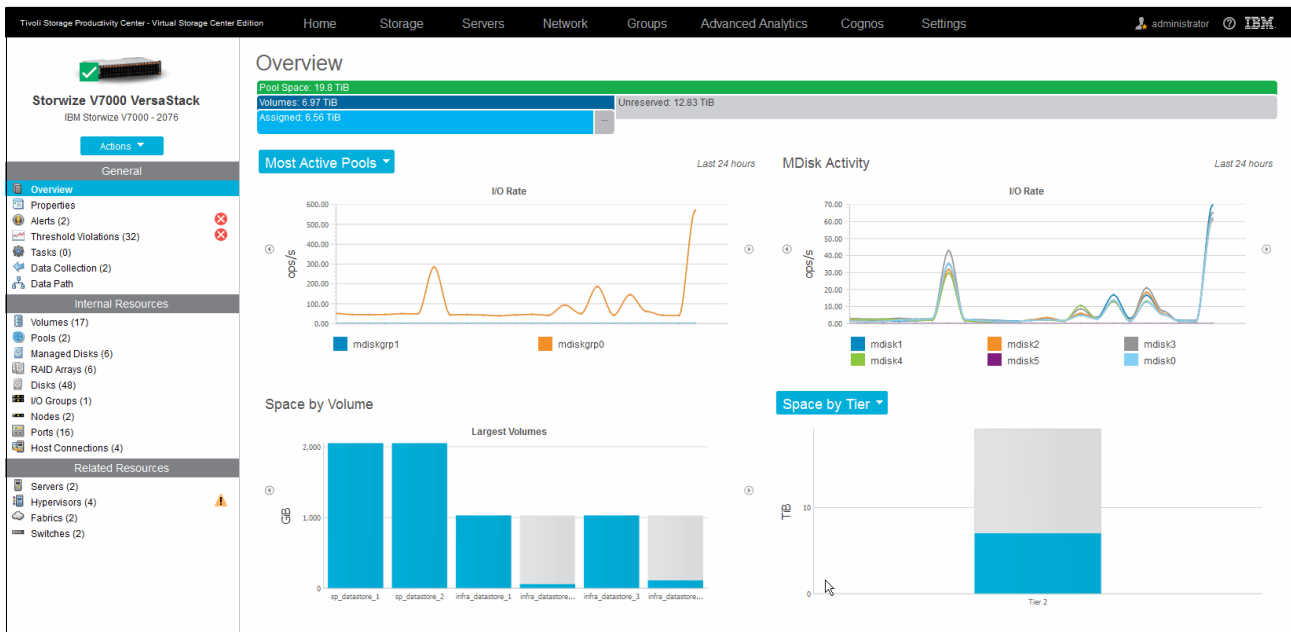


Figure 12-21 Storwize V7000 Overview - Alternative Data Graphs 2

- Properties: Provides a summarized overview of the system, including details such as model number, serial number with tabs for high-level configurations, capacity, and performance.

Figure 12-22 shows the VSC Storwize V7000 Properties for editing the custom tags.

The screenshot shows the 'Properties' page for a Storwize V7000 VersaStack. The left sidebar contains navigation options: Overview, Properties (selected), Alerts (2), Threshold Violations (32), Tasks (0), Data Collection (2), and Data Path. Below these are sections for Internal Resources (Volumes, Pools, Managed Disks, RAID Arrays, Disks, I/O Groups, Nodes, Ports, Host Connections) and Related Resources (Servers, Hypervisors, Fabrics, Switches). The main content area is titled 'Properties' and has tabs for General, Configuration, Capacity, and Performance. The 'General' tab is active, displaying various system details. At the bottom, there are 'Save' and 'Cancel' buttons.

Property	Value
Name	Storwize V7000 VersaStack
Status	Normal
Vendor	IBM
Type	Storwize V7000 - 2076
Model	24F
Serial Number	00000100208030A6
Firmware	7.5.0.0 (build 115.50.1506041858900)
Turbo Performance	Active
IP Address	192.168.10.19
Probe Status	Successful
Probe Schedule	Daily. Next run at Jun 29, 2015 16:45:00 PDT
Performance Monitor Status	Running
Performance Monitor Interval (min)	1
Time Zone	US/Pacific
Data Source Count	1
Location	San Jose
Custom Tag 1	No Custom Tag
Custom Tag 2	No Custom Tag
Custom Tag 3	No Custom Tag

Figure 12-22 Storwize V7000 Properties - Custom Tags

- Alerts: Gives you the alerts that are related to this device only, as opposed to system-wide alerts. For more information about alerts, see 12.7.3, “Monitoring and alerting” on page 290.

Figure 12-23 shows the VSC Storwize V7000 Overview Alerts.

The screenshot shows the 'Alerts' page for a Storwize V7000 VersaStack. The left sidebar is identical to Figure 12-22, with 'Alerts (2)' selected. The main content area is titled 'Alerts' and shows a summary: 1 Critical, 1 Warning, and 0 Informational. Below this is a table of alerts with columns for Condition, Severity, Last Occurrence, and Internal Resource.

Condition	Severity	Last Occurrence	Internal Resource
Total I/O rate threshold	Critical	Jun 29, 2015 10:45:48 PDT	io_qrp0
Total data rate threshold	Warning	Jun 29, 2015 10:44:46 PDT	io_qrp0

Figure 12-23 Storwize V7000 - Alerts

- **Threshold Violations:** Shows any violations for the thresholds that you filtered for this specific device. In our example, we set an aggressive warning (1500 ops/s) and critical (2000 ops/s) for the Total I/O Rate Threshold.

Figure 12-24 shows the VSC Storwize V7000 Overview Threshold Violations.

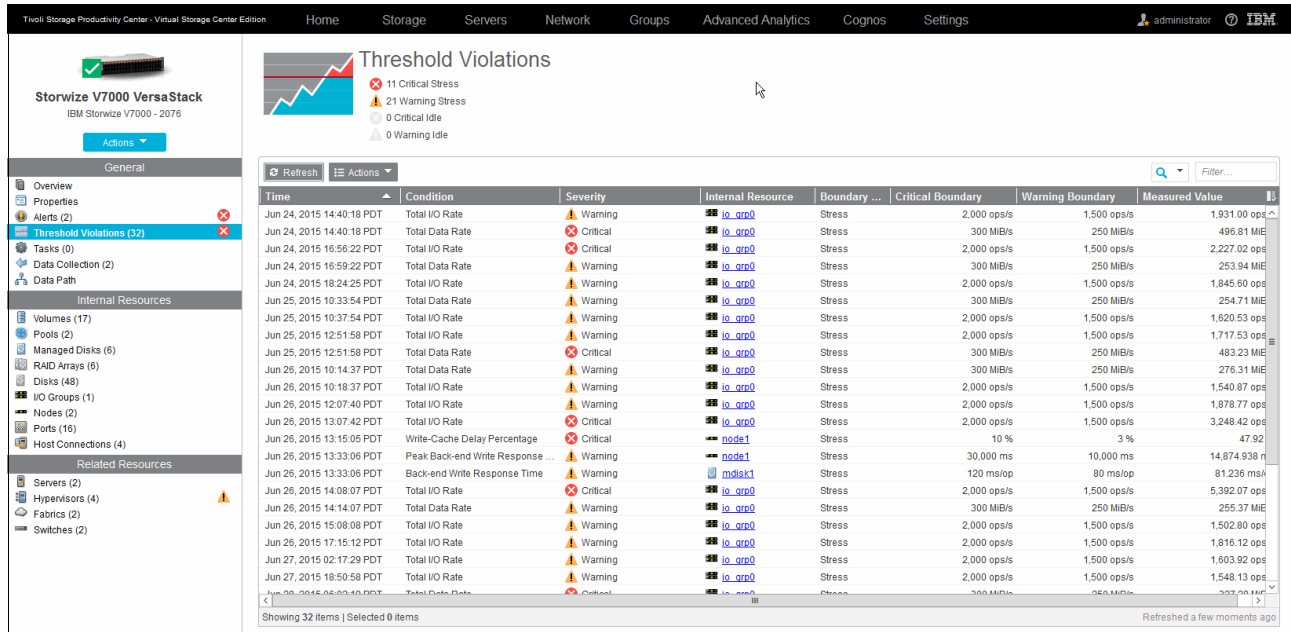


Figure 12-24 Storwize V7000 - Threshold Violations

- **Tasks:** Shows the VSC Auto-Tiering or Provisioning functions that can trigger tasks for the storage system. These tasks can be to up- or down-tier a volume, volume creation, and so on, that are either scheduled to run automatically or wait for administrator approval before execution.
- **Data Collection:** Shows the status of the Probe execution and Performance monitors for the storage subsystem. You can use it to schedule the probing, stop or start performance monitoring, and review the associated log files.

Figure 12-25 on page 285 shows the VSC Storwize V7000 status of the Data Collection engine.

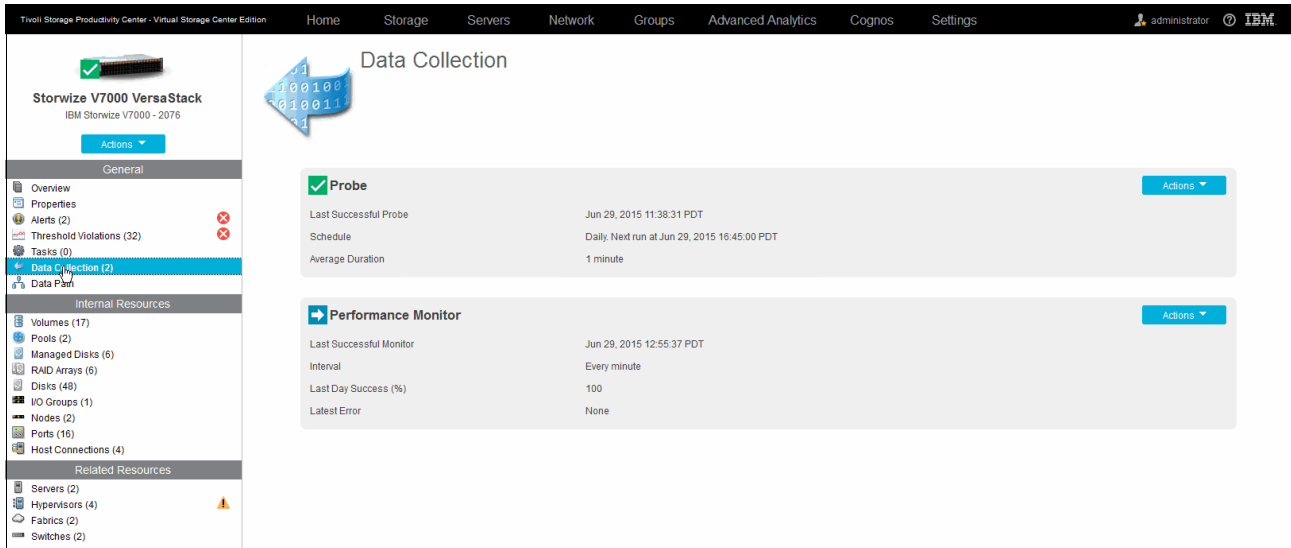


Figure 12-25 Storwize V7000 - Data Collection

- Data Path Topology View: Gives you an overview and the data path of all connected resources to the storage devices. If you right-click any of these resources, you can either open the properties or to jump directly to the overview pane of that specific device.

Figure 12-26 shows the VSC Storwize V7000 data path topology view with system summary.

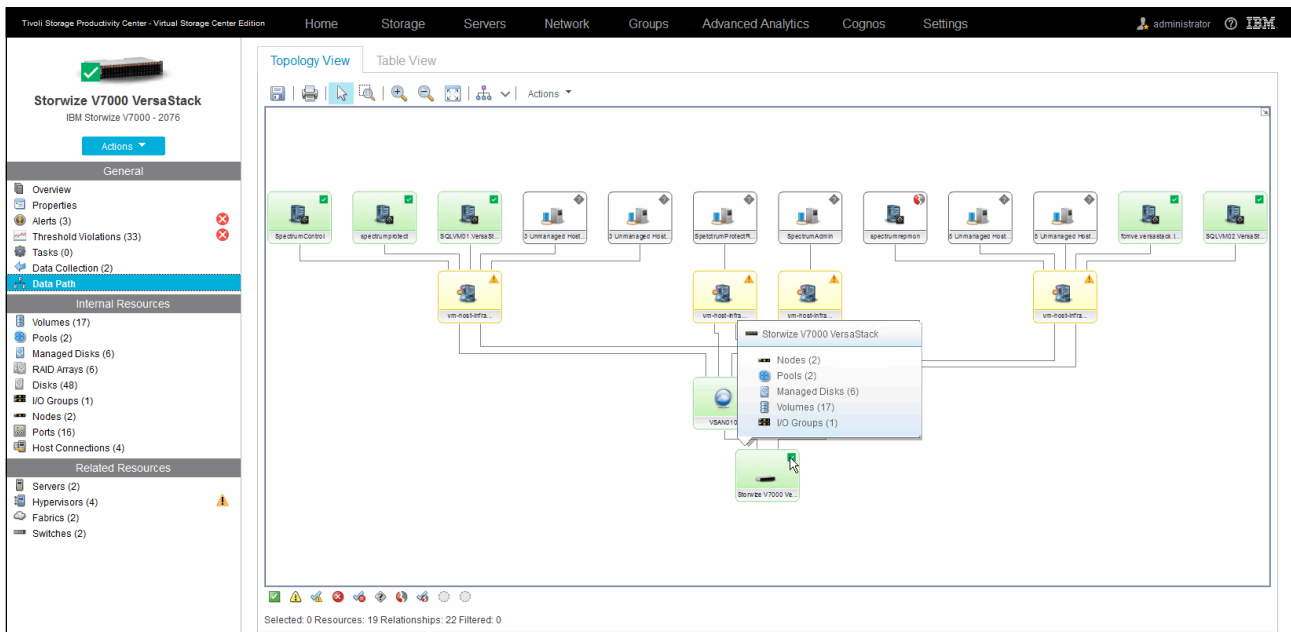


Figure 12-26 Storwize V7000 - Topology View

Note: Using the stand-alone VSC GUI, you can also create regular inventory snapshots of the whole environment that is managed by VSC, giving you point-in-time tracking of all the changes that are made to the environment. Likewise, you can perform a Configuration Analysis of your fabric that is based on industry preferred practices.

Example 12-1 shows the configuration analysis.

Example 12-1 Configuration analysis

```
2015-06-29 14:09:17.495-0700 GEN7090I: Checking/waiting for other running
analyzer(s)
2015-06-29 14:09:17.589-0700 GEN7098I: The data scope for this configuration
analysis job run: All Fabrics
2015-06-29 14:09:17.589-0700 GEN7107I: The following configuration analysis
policies have been selected:
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 1, description=Each connected computer and storage subsystem port must
be in at least one zone in the specified zone sets.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 2, description=Each HBA accesses storage subsystem ports or tape
ports, but not both.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 3, description=Each volume is accessed only by computers running the
same type and version of operating system.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 4, description=Each zone contains only HBAs from a single vendor.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 5, description=Each zone contains only a single model of storage
subsystem.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 6, description=Each zone is part of a zone set.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 7, description=Each host must be zoned so that it can access all of
its assigned volumes.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 8, description=Each computer has only HBAs of the same model and
firmware version.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 9, description=For each host type and operating system, every HBA of a
given model must have the same firmware and driver version.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 10, description=Every SAN switch of a given model must have the same
firmware version.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 11, description=Every storage subsystem of a given model must have the
same firmware version.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 14, description=Replication Plan is intact with respect to the SRG(s)
and the replication session associated during planning through SAN Planner.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 15, description=All the source volumes involved in Metro Mirror
Failover/Failback sessions are conforming to 1:4 primary to secondary LSS for a
failback direction scenario.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 16, description=Inter/intra site connectivity is valid for replication
plan deployments.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 12, description=Each fabric may have a maximum of x zones.
2015-06-29 14:09:17.589-0700 GEN7097I: The configuration analysis policy:
name=Policy 13, description=Each zone may have a maximum of x zone members.
```


2015-06-29 14:09:17.589-0700 GEN7096I: The configuration analysis job run has started.
 2015-06-29 14:09:19.009-0700 GEN7093I: No policy violations occurred during this configuration analysis job run.
 2015-06-29 14:09:19.009-0700 GEN7092I: The configuration analysis job run completed successfully.

Figure 12-27 shows the VSC Stand-alone GUI slideable configuration history.




Figure 12-27 Configuration History

- ▶ Internal Resources: Groups the corresponding device-specific resources. From within each resource, you can directly jump to the performance metrics for that specific resource. For the Storwize V7000 storage system, the following resources are shown:
 - Volumes
 - Pools
 - Managed Disks
 - RAID Arrays
 - Disks
 - I/O Groups
 - Nodes

In our example configuration, only half of the Storwize V7000 ports are actively connected to the Cisco UCS fabric interconnects. To avoid the system giving an error status, we acknowledged this status from within the Internal Resources window, as shown in Figure 12-28.

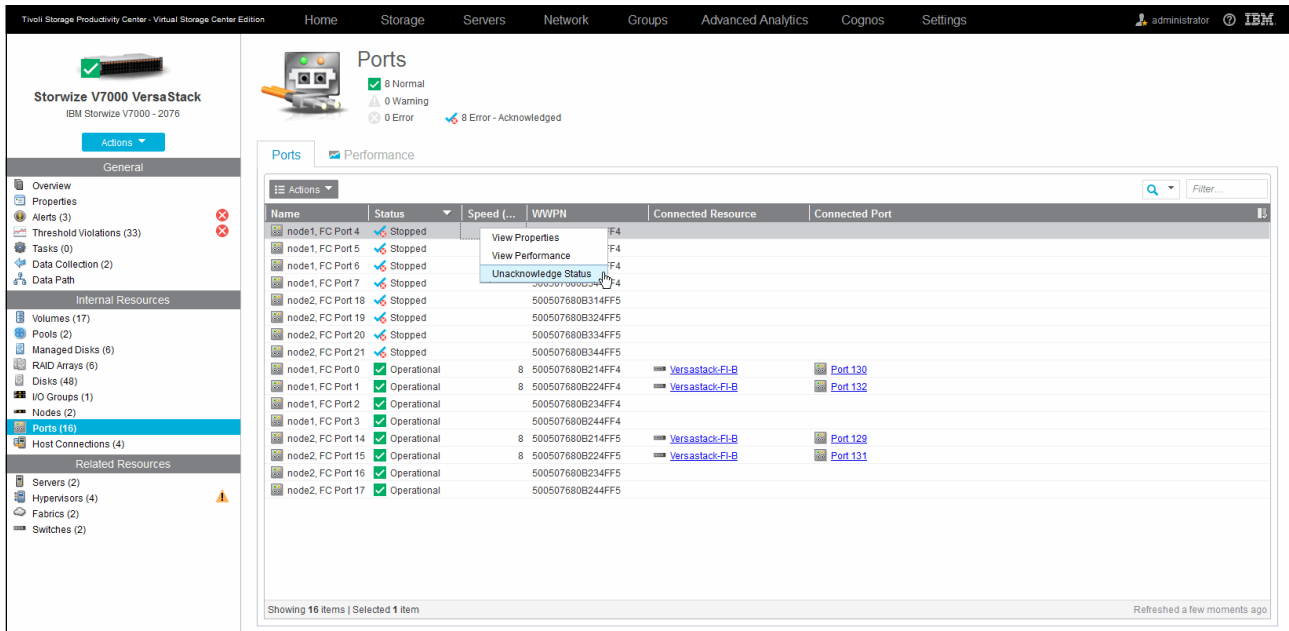


Figure 12-28 Storwize V7000 Internal Resources

Note: Every column view in the VSC GUI can be customized to show related information by selecting the column check mark in the upper right corner of the table.

Figure 12-29 shows the VSC Storwize V7000 adding additional content to the column view.

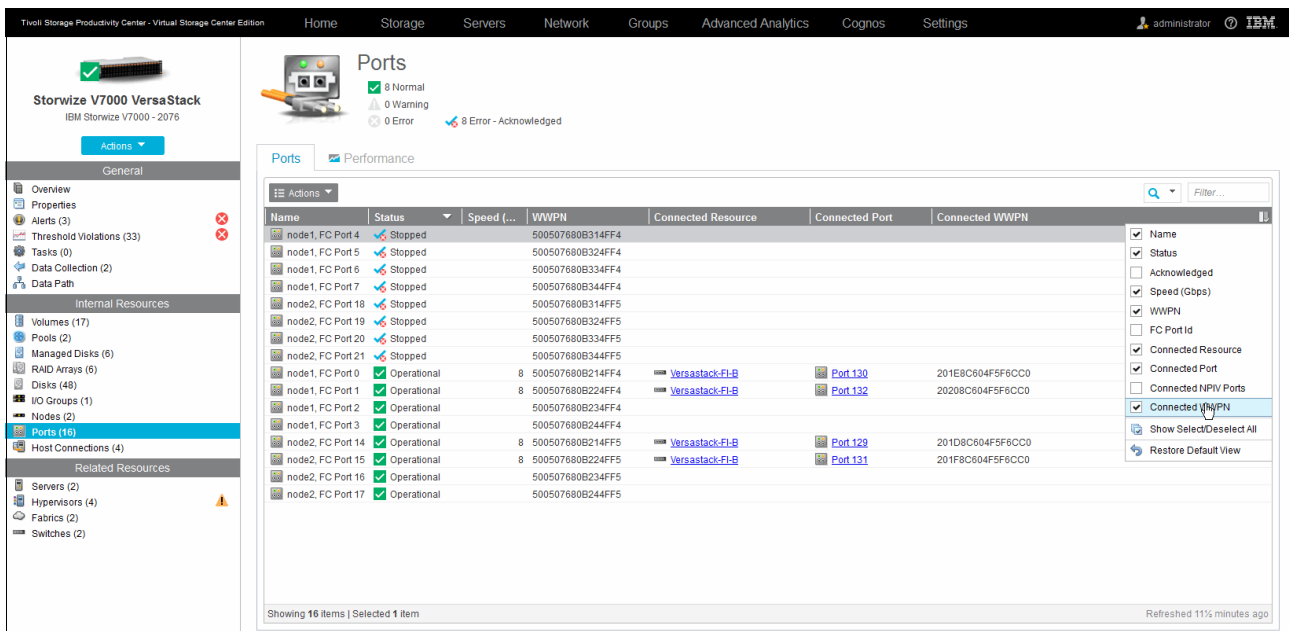


Figure 12-29 Storwize V7000 - Add Columns

After you specified the information that you want to display, you can also export this information as a CSV, PDF, or HTML file through Actions/Export in all of the column table views.

- ▶ Related Resources are similar to Internal Resources, and Related Resources provides you with information about equipment that interacts with the Storwize V7000 storage system and its resources grouped by the following categories:
 - Servers
 - Hypervisors
 - Fabrics
 - Switches

You can use the Servers and Hypervisors resources to also display co-related information from their specific detail panes and have a Disk Mapping Section outlining the disks that they use on the Storwize V7000 storage system.

Figure 12-30 shows the VSC Storwize V7000 Related Resources for Servers with an additional information column selection.

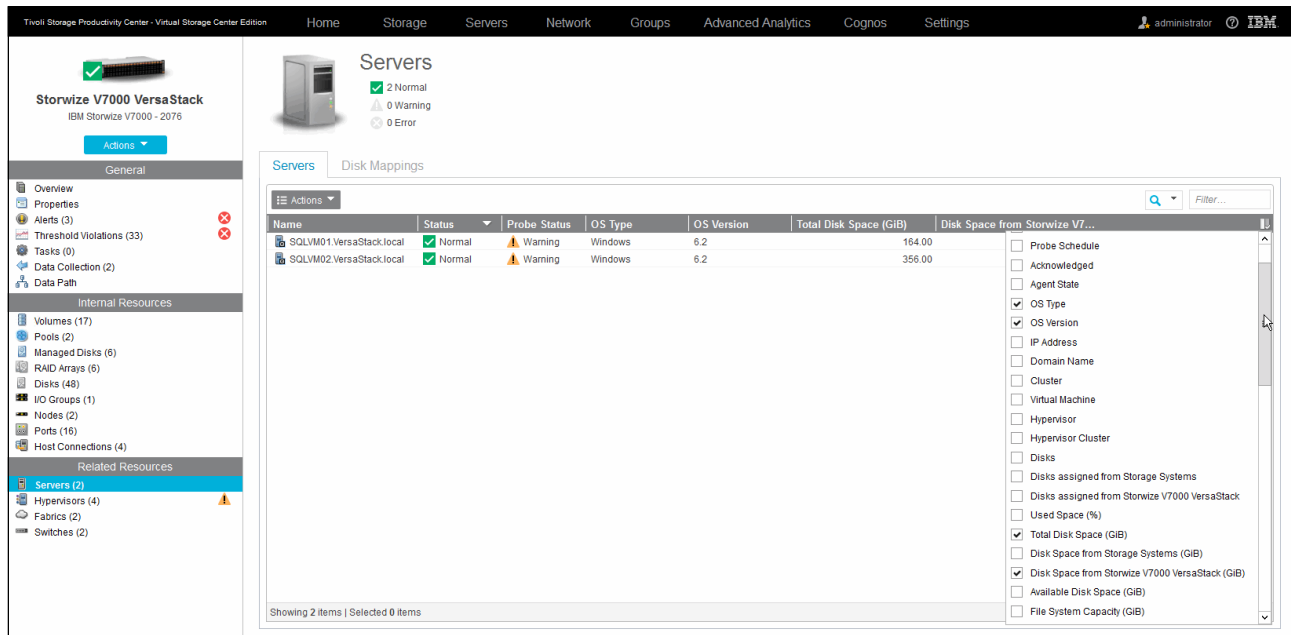


Figure 12-30 Storwize V7000 - Related Resources Servers

Fabrics shows the fabric and switch relationships for the connected Storwize V7000 storage system, where Switches also take you directly to the performance pane of the corresponding switch.

This completes the functional overview of the Storwize V7000 storage system from within the VSC Web GUI. The next section describes the performance monitoring capabilities, and you can define which alerts to be generated and thresholds to be set.

12.7.3 Monitoring and alerting

The Storwize V7000 storage system has built-in, 5-minute, and sample-based performance monitoring, as shown in Figure 12-31.

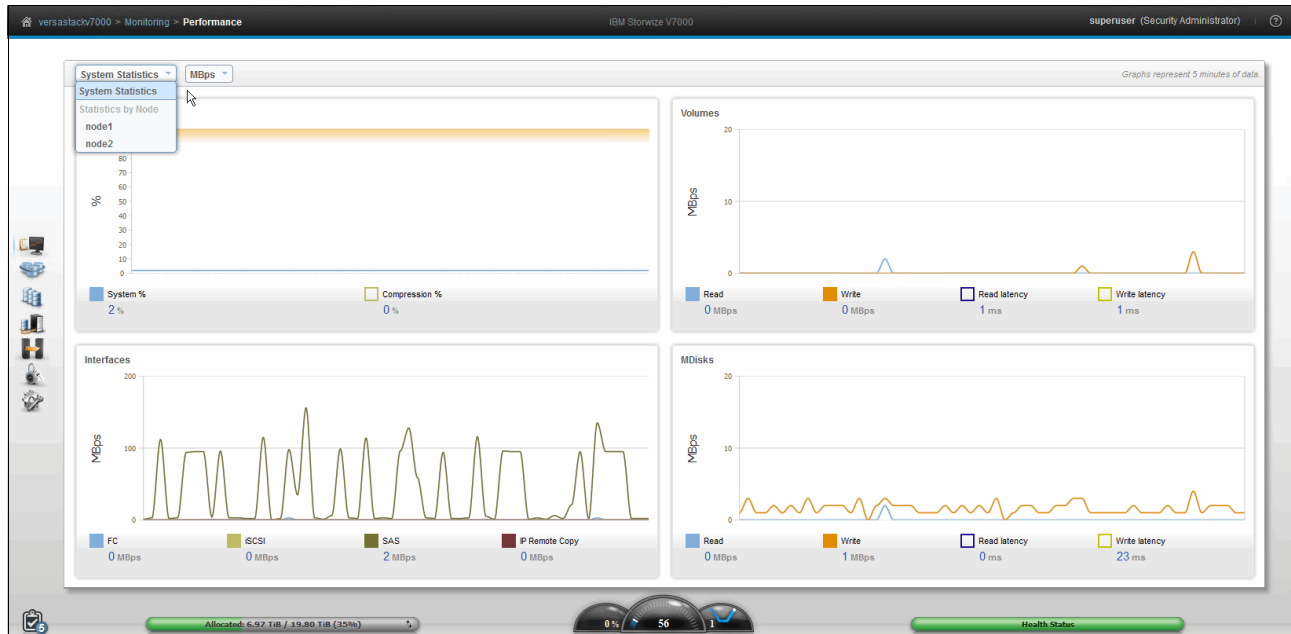


Figure 12-31 Storwize V7000 - Performance Monitoring

Tivoli Productivity Center Virtual Storage Center Edition enhances the real-time performance monitoring of the managed resources, such as the Storwize V7000 storage system, through the following features:

- ▶ Unlimited performance data capturing
- ▶ Granular performance metrics on the following levels:
 - I/O Group level
 - Node level
 - Port level
 - Pool/Volume level
 - Managed Disk/RAID Array/Disk level
- ▶ Holistic performance monitoring from server/hypervisor over fabric/switch to the storage system
- ▶ Customizable threshold settings with co-related alert triggers and actions

In this section, you perform the following tasks:

- ▶ Set the retention parameters for the performance data.
- ▶ Configure the system-wide alert notifications.
- ▶ Define performance thresholds and custom alerts.
- ▶ Correlate volume and I/O group performance.

History Retention

Review the History Retention settings by clicking the **VSC Web GUI Settings** drop-down menu, as shown in Figure 12-32 on page 291.

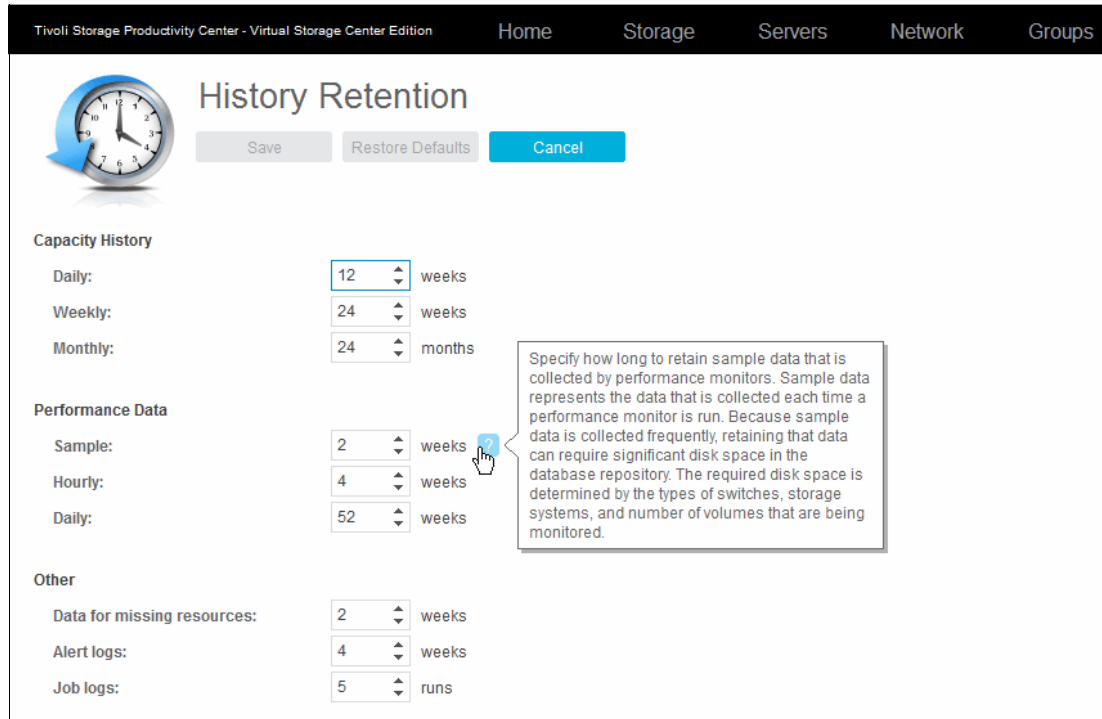


Figure 12-32 Virtual Storage Center - History Retention

The performance data is stored in the DB2 database of the VSC. You can increase the retention period if you allocated enough disk capacity on the VSC system itself. The hardware requirements are outlined in the support document that is found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg27039550>

Alert Notifications

Within the VSC, alerts and notifications can be sent to three different receivers in parallel: Email, SNMP, and IBM Netcool® / OMNIBus.

To configure the targets for the Alert Notifications, complete the following steps:

1. Click **Settings** → **Alert Notifications** in the VSC Web GUI, as shown in Figure 12-33.

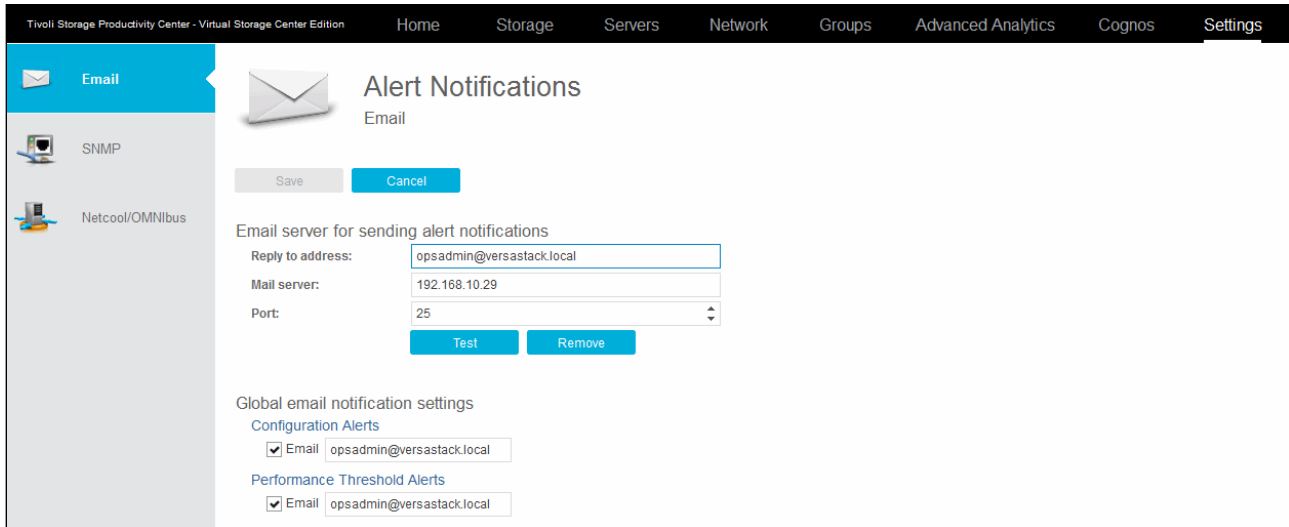


Figure 12-33 Alert Notifications - Email

2. Complete the reply to, mail server, and port settings and click the **test** button to verify email reception.

Example 12-2 shows the sample email that is generated by the Alert Notification email test.

Example 12-2 Alert email verification

SRV0785I: This email is a test of the alert notification configuration in Tivoli Storage Productivity Center. Receiving this message indicates that email notification is configured correctly.

You can specify up to two SNMP destinations by providing the community, host name, or IP address and port settings.

Example 12-3 shows the SNMP trap of a failed VSC Job.

Example 12-3 Sample VSC SNMP trap

```
20:38:51 2015/06/24 ZBXTRAP 192.168.155.18
PDU INFO:
version                0
community              public
errorstatus            0
receivedfrom           UDP: [192.168.155.18]:61914->[192.168.155.23]:162
messageid              0
notificationtype       TRAP
errorindex              0
requestid              0
transactionid          3
VARBINDS:
DISMAN-EVENT-MIB::sysUpTimeInstance type=67 value=Timeticks: (385718) 1:04:17.18
SNMPv2-MIB::snmpTrapOID.0          type=6 value=OID: TIVOLI-SRM-MIB::jobFailedTrap
TIVOLI-SRM-MIB::scheduleName        type=4 value=STRING:
"administrator.Probe_linux_tsm71lnx.ibmdemo.local"
```

```

TIVOLI-SRM-MIB::scheduleType    type=4  value=STRING: "Probe"
TIVOLI-SRM-MIB::scheduleRun     type=4  value=STRING: "3"
TIVOLI-SRM-MIB::alertType       type=4  value=STRING: "Scheduled Job Failed"
TIVOLI-SRM-MIB::alertName       type=4  value=STRING:
"administrator.probeFailedAlertConditionName_7065"
TIVOLI-SRM-MIB::serverName      type=4  value=STRING: "Data Manager server on
tpc52.ibmdemo.local"
TIVOLI-SRM-MIB::alertID         type=4  value=STRING: "6001"
TIVOLI-SRM-MIB::alertURL        type=4  value=STRING:
"https://TPC52.ibmdemo.local:9569/srm/gui#alerts?id=6001"
TIVOLI-SRM-MIB::resourceURL     type=4  value=STRING:
"https://TPC52.ibmdemo.local:9569/srm/gui#resources?type=servers&id=7065"
SNMP-COMMUNITY-MIB::snmpTrapAddress.0 type=64 value=IpAddress: 192.168.155.18
SNMP-COMMUNITY-MIB::snmpTrapCommunity.0 type=4  value=STRING: "public"
SNMPv2-MIB::snmpTrapEnterprise.0 type=6  value=OID: TIVOLI-SRM-MIB::srmServer

```

Instructions about how to configure your SNMP server and where to obtain the VSC MIB files can be found at the following website:

http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.6/com.ibm.tpc_v526.doc/fqz0_t_configuring_snmp_alerts.html

Alternatively, supply the host name or IP address for the IBM Netcool / OMNIBUS server.

Defining performance thresholds

With the system-wide notifications set, proceed with creating a performance threshold alert and apply a custom alert notification to it. Notification settings can be system-wide, device-specific, and event-specific.

Figure 12-34 shows how you can override the global notification settings for the Storwize V7000 storage system itself, and set custom notifications for the Storwize V7000 storage system.

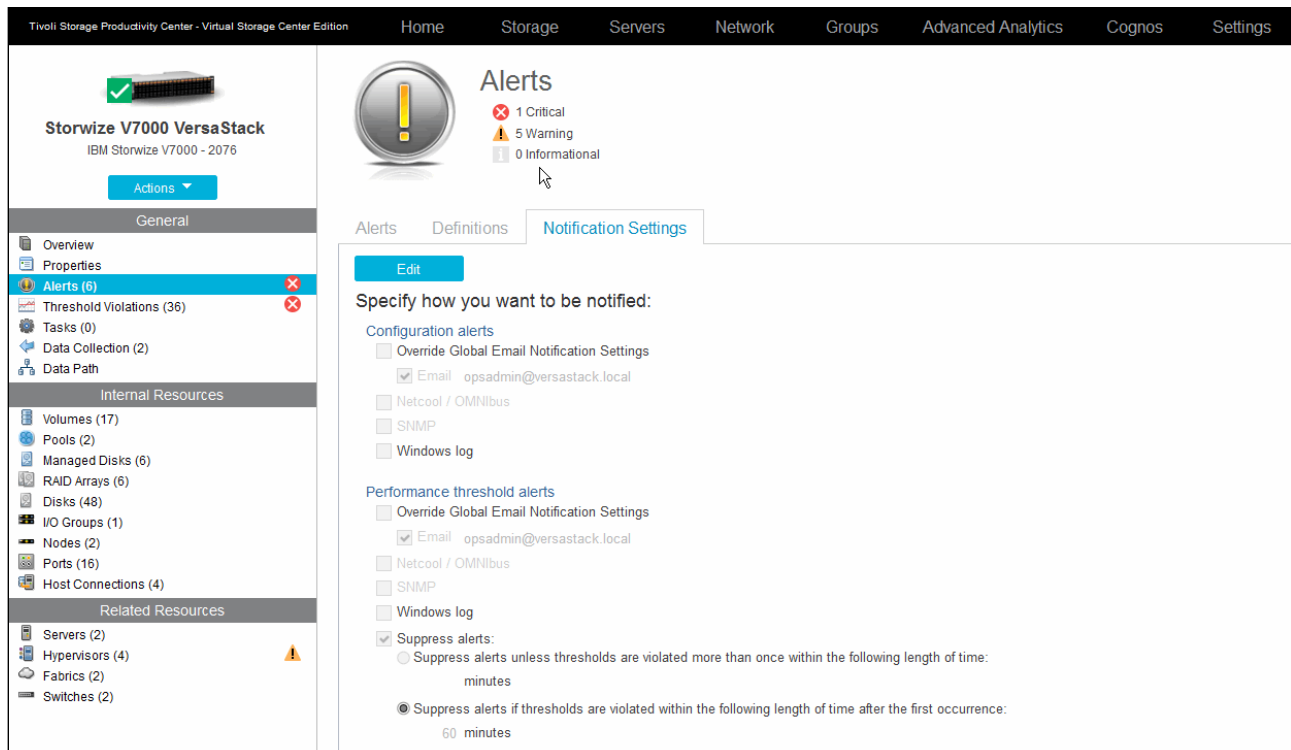


Figure 12-34 Storwize V7000 - custom notification settings

You can distinguish between configuration and performance alerts routing, for example, the configuration alerts to the infrastructure team and the performance alerts to the application team in your organization. By default, repeating performance alerts are suppressed within the first 60 minutes on subsequent occurrences.

If you switch back to the Definitions tab, you can toggle and customize alerts for the following alert types:

- ▶ Storage Systems
- ▶ Nodes
- ▶ Pools
- ▶ Volumes
- ▶ Disks
- ▶ Performance

Figure 12-35 on page 295 shows setting the I/O threshold rates for the Storwize V7000 storage system.

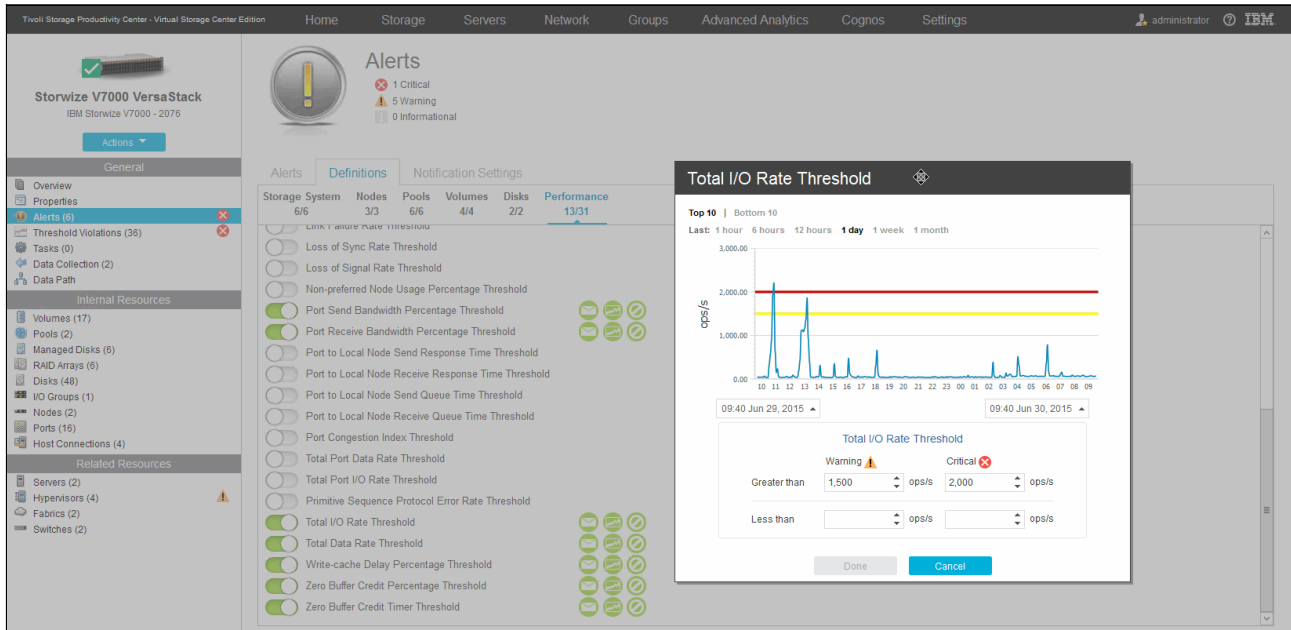


Figure 12-35 Storwize V7000 - I/O Threshold Customization

With the new thresholds defined, you can override the email notification by sending it to the storage admin team email of `storadmin@versastack.local`.

Figure 12-36 shows setting the custom I/O threshold notifications for the Storwize V7000 storage system and exploring the Run script option.

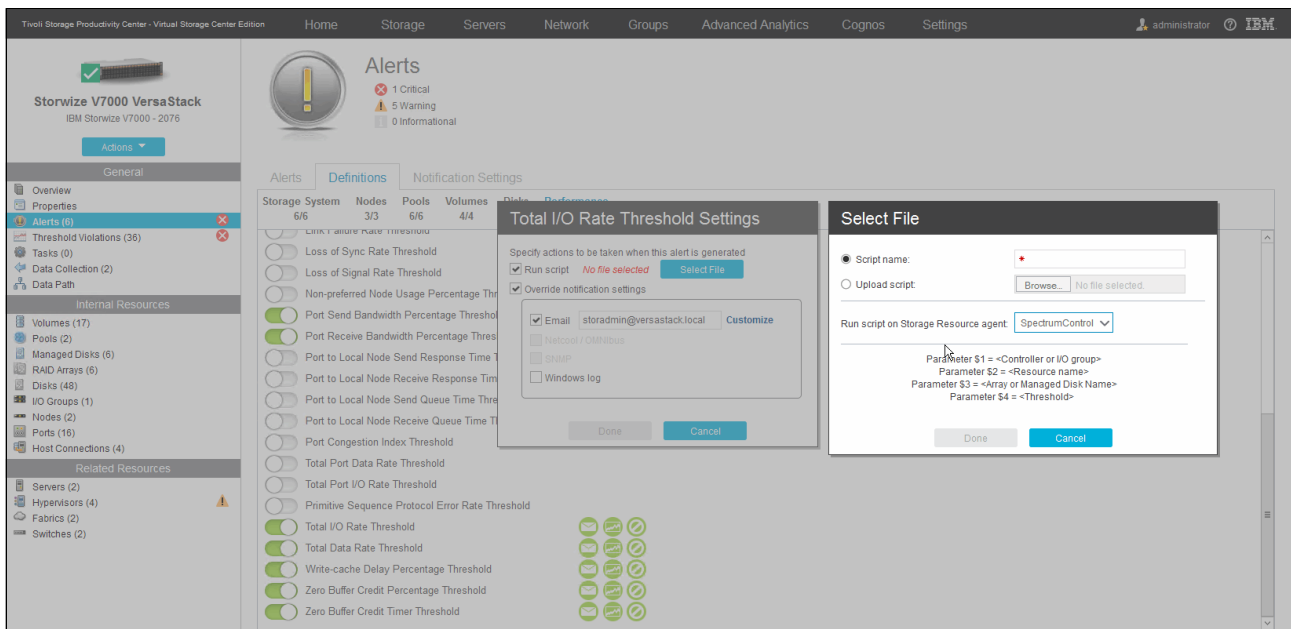


Figure 12-36 Storwize V7000 - Custom Notifications

Another option is to have a script run when the alert is being triggered. These scripts are run by the Storage Resource Agents in your environment. By default, a Storage Resource Agent is deployed on the VSC itself.

These scripts can trigger corrective actions directly against the storage system by using remote CLI or interacting with the VSC itself to create or run scheduled tasks. They can also run scripts and commands directly on the host operating system of the SRA. You can, for example, trigger the Analyze Tiering for the storage system to have VSC automatically up- or down-tier the volumes to optimize the I/O load whenever a high-level or low-level threshold is passed.

Correlating performance data

In the example environment, we set two threshold alerts for Total Data Rate and Total I/O Rate.

Figure 12-37 shows the VSC Storwize V7000 triggered threshold violations.

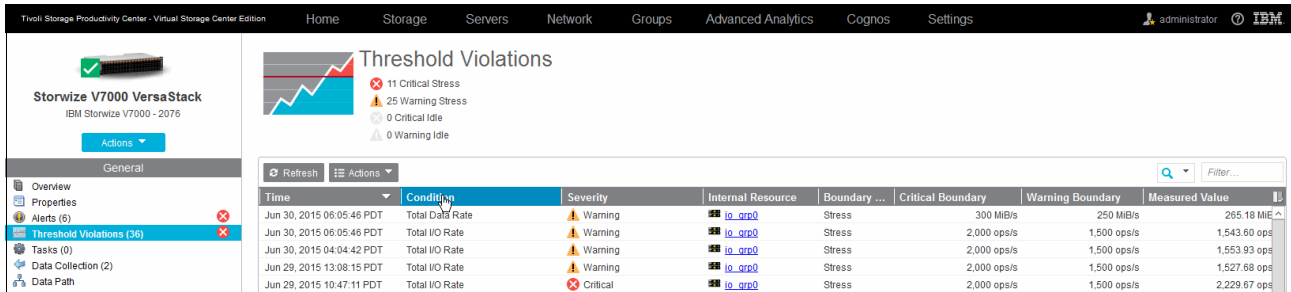


Figure 12-37 Threshold violations

As an example, we investigate what was causing the Total Data Rate alert by double-clicking the alert itself, as shown in Figure 12-38 on page 297.

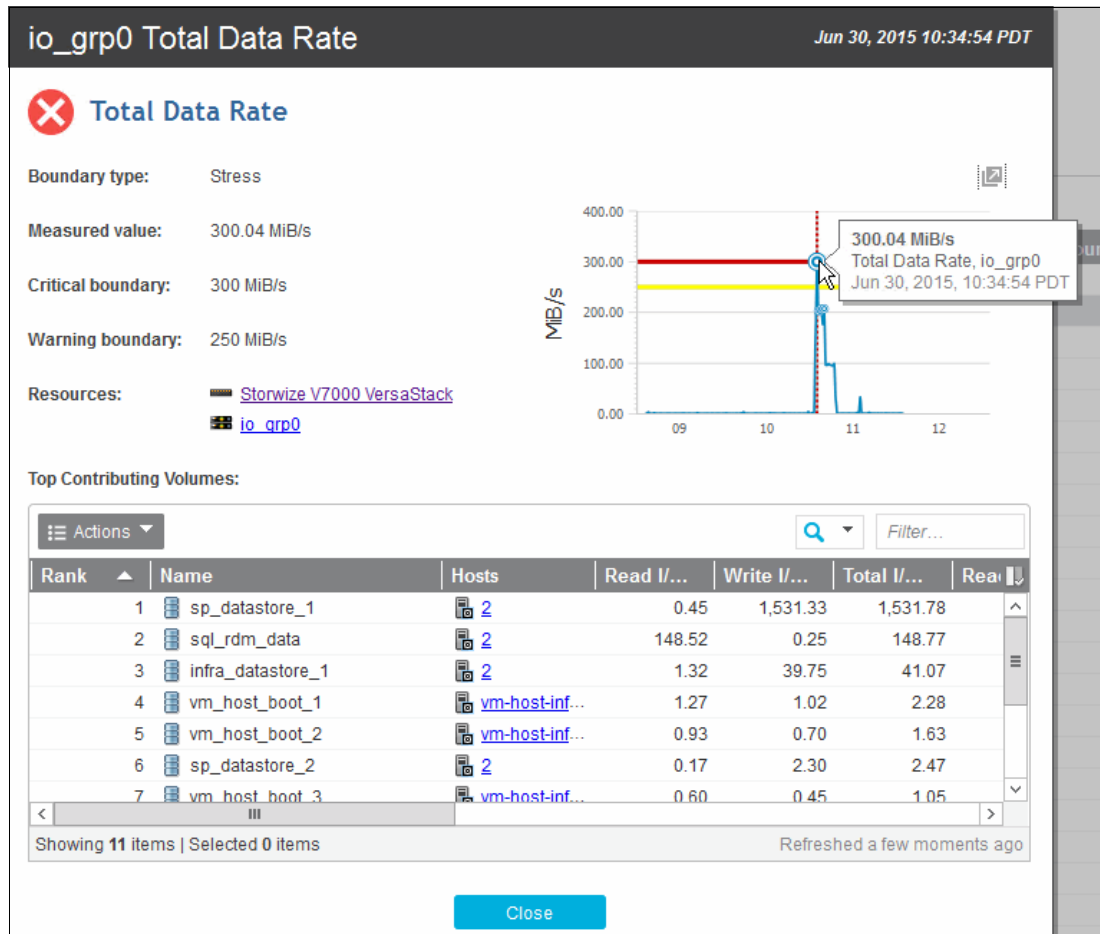


Figure 12-38 Total Data Rate Alert Detail

The details inform you about the measured value when the alert occurred. You can also hover your cursor over the chart itself to get more data samples. The column chart indicates that the `sp_datastore_1` (which hosts the IBM Spectrum Protect Tivoli Storage Manager server in our example environment) has the highest write rate with the highest read coming from the clustered data volume from the DB2 Server.

If you do not know what system is causing the load, you can double-click, for example, `sp_datastore_1` to get more information. We are interested in finding out the disk mappings on this datastore, so go to the corresponding tab, as shown in Figure 12-39.

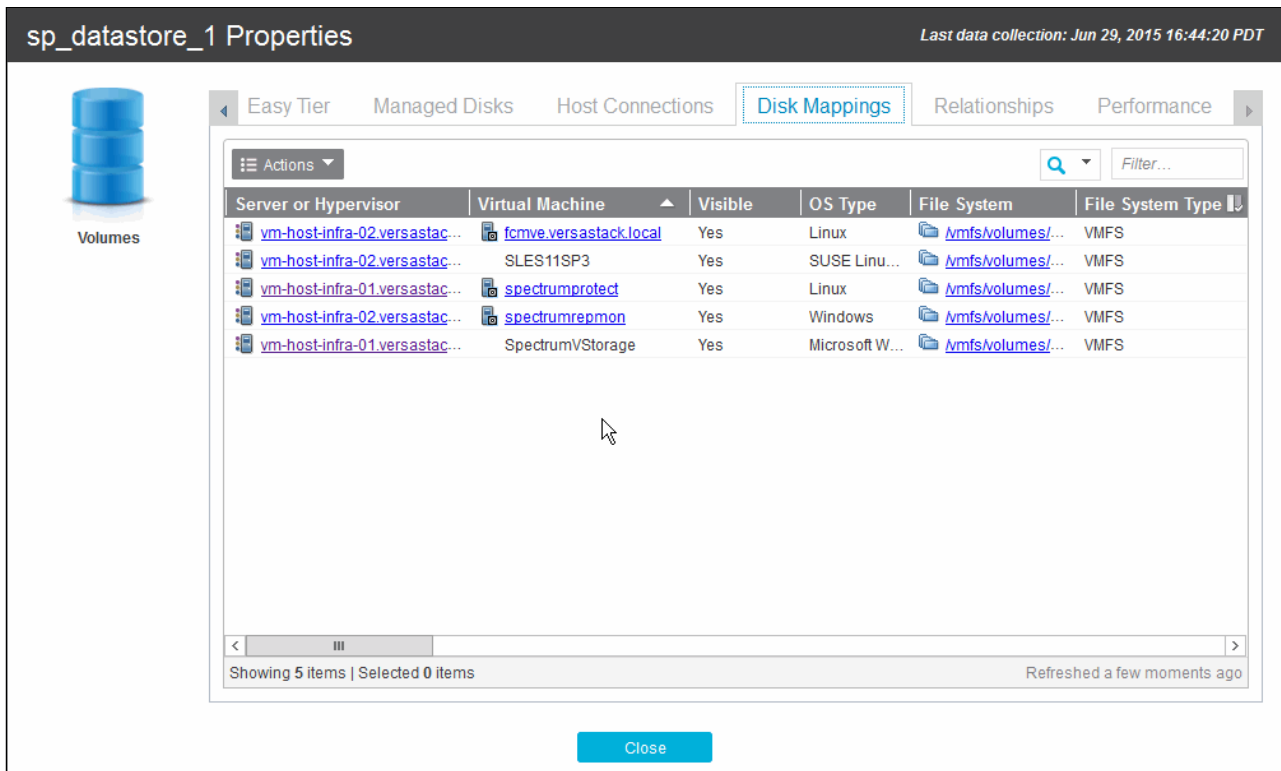


Figure 12-39 datastore volume disk mappings

The IBM Spectrum Protect VM is hosted on the `vm-host-infra-01` hypervisor. Clicking the link takes you directly to the overview pane of that system, as shown in Figure 12-40.

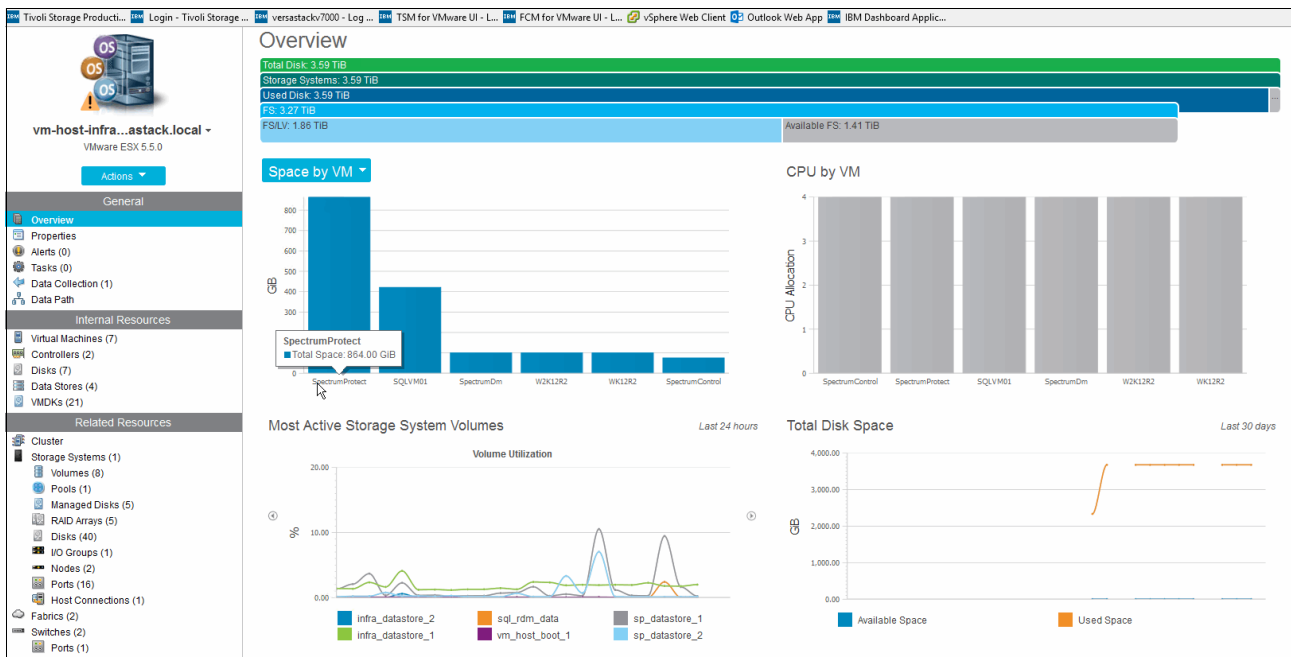


Figure 12-40 Hypervisor overview

The sp_datastore_1 is one of the most active volumes, followed by the sp_datastore_2. Going to the volumes by using the related sources on the left side gives an overview of all the volumes that are related to this hypervisor.

Figure 12-41 shows the VSC Hypervisor Infra-01 volume performance.

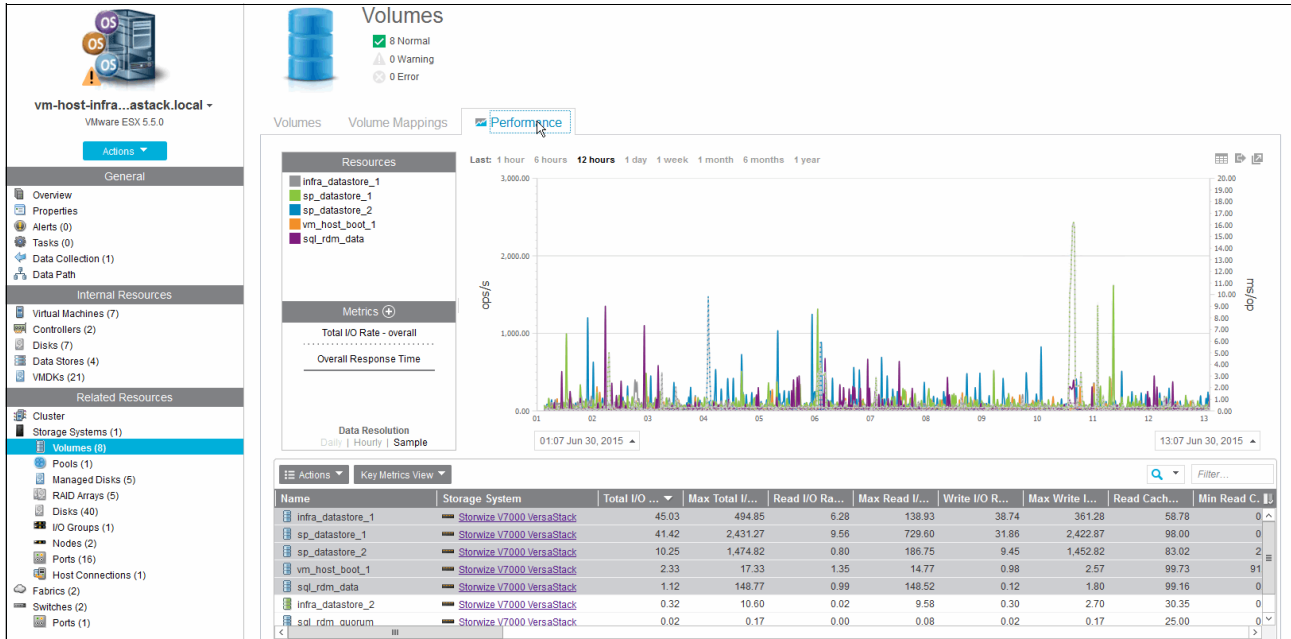


Figure 12-41 Hypervisor Volume Performance

The VSC selected the top five volumes that are grouped by Total I/O performance and shows the key metrics that are related to these volumes. You can customize this view and select the metric that you need for both the table view and for the performance graph independently, as shown in Figure 12-42.

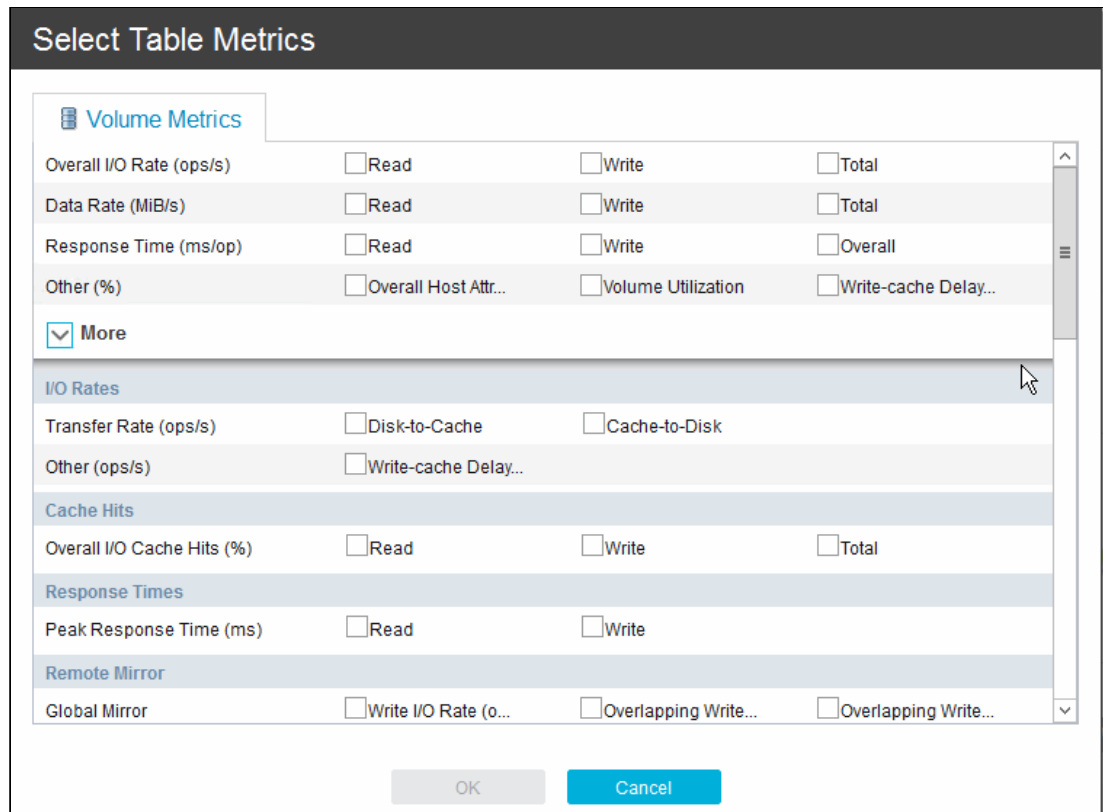


Figure 12-42 Hypervisor volume key metrics

We want to investigate the Data Rate Response Time for the volume hosting the IBM Spectrum Protect Tivoli Storage Manager server and the data volume of the DB2 cluster in our VersaStack environment and see how it evolved over the last month in the performance graph. Selecting **1 month** and using the **Metrics +** button shows the required information, as shown in Figure 12-43 on page 301.

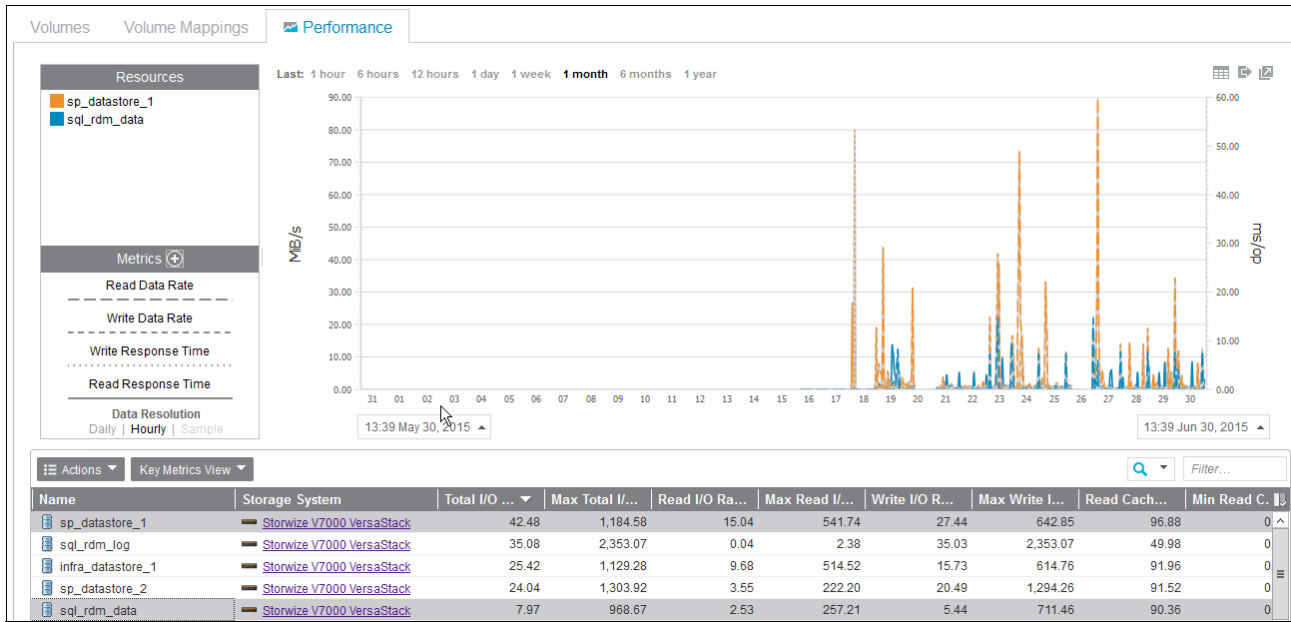


Figure 12-43 One-month volume performance

You can split this window to evaluate the performance of multiple sources one by one by using the open in a new window arrow button in the upper right corner of the graph. These resources can be dissimilar, showing, for example, volume performance, storage system FC port performance, and SAN fabric performance in separate windows with different metrics. After you find the specific spot that you want to investigate in more detail, you can synchronize all the windows by using the Synchronize Time clock icon, as shown in Figure 12-44.

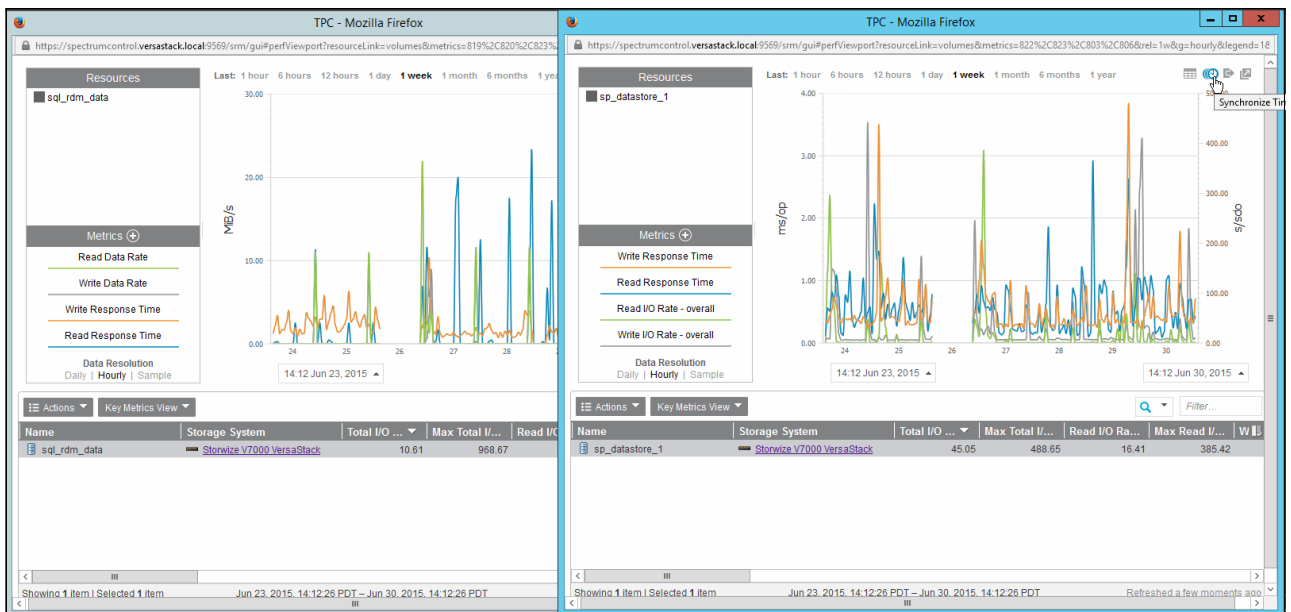


Figure 12-44 Time synchronized multiple performance windows

This completes the section about how to integrate the Storwize V7000 storage system with the Spectrum Control VSC. You explored the main VSC interfaces that are related to the Storwize V7000 storage system, defined and examined alerts, and viewed performance metrics on the Storwize V7000 resources.

In the next section, you add the hypervisors to the Spectrum Control VSC, register the vCenter Web Client extension, and explore the alerting and performance monitoring from within the VSC Web GUI and through the vSphere Web Client interface.

Integrating the VMware vCenter Hypervisor with Spectrum Control

Integrating the VMware vCenter and the ESXi Hypervisors that are used in the example VersaStack environment follows a similar approach as adding the Storwize V7000 storage system to the Tivoli Productivity Center SmartCloud Virtual Storage Edition.

Start the VSC Web GUI and click the **Hypervisors** section to start the registration process, as shown in Figure 12-45.

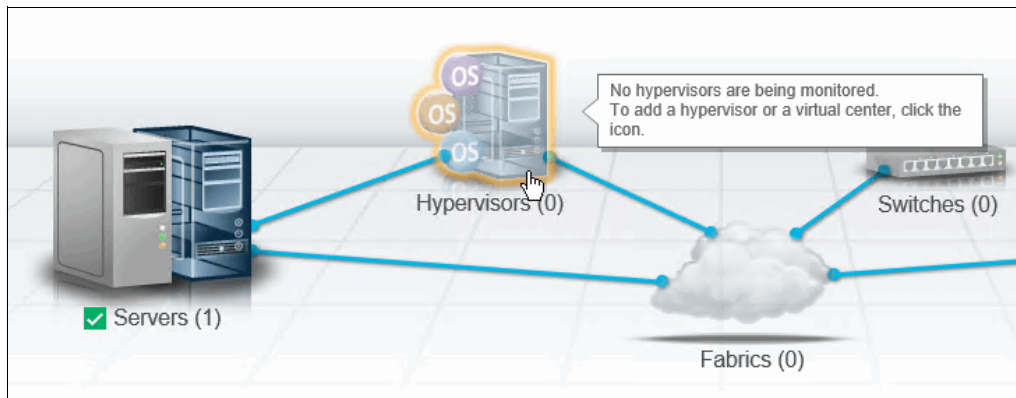


Figure 12-45 Initiate Add Hypervisor wizard

Figure 12-46 shows the VSC Hypervisor Add VMware vCenter wizard.

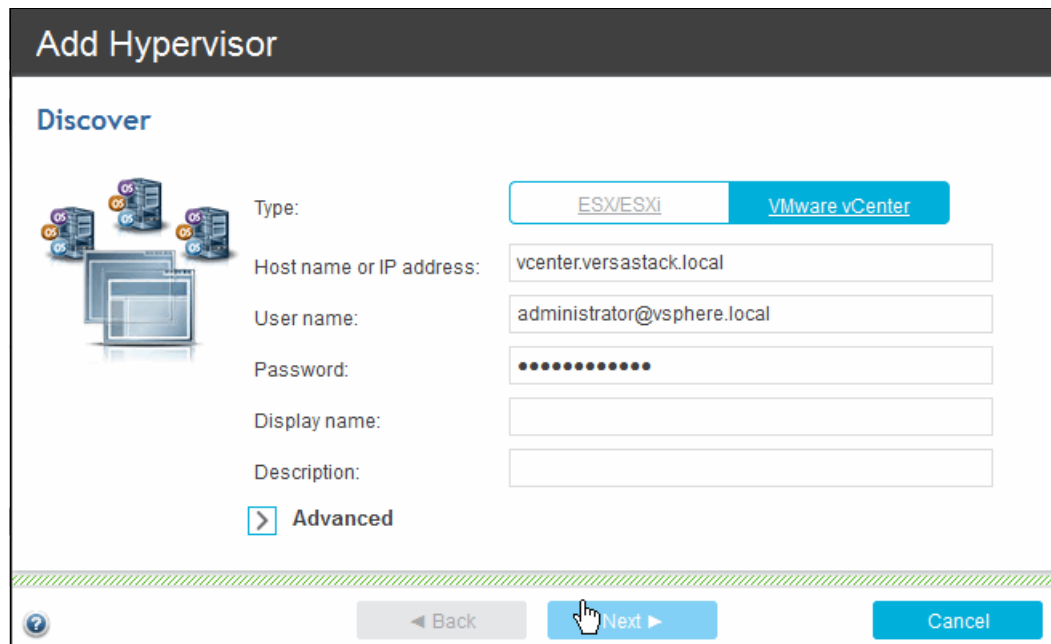


Figure 12-46 Add Hypervisor

VSC now connects to the vCenter and obtains a list of registered clusters and hypervisors. After this task completes, you can deploy the vSphere Web Client extension, as shown in Figure 12-47.

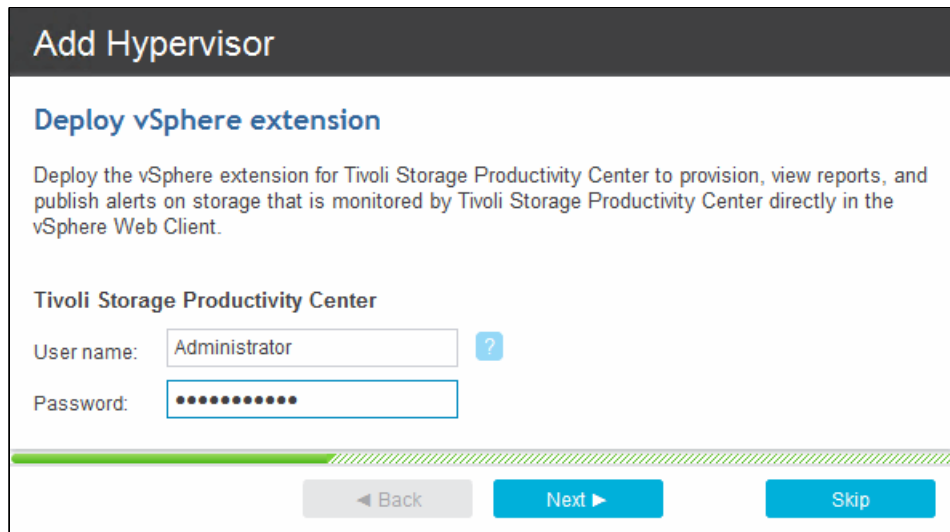


Figure 12-47 Deploy vSphere extension

VSC uses a probing mechanism to check for configuration changes at regular intervals. After the Hypervisors are discovered, the system proposes the creation of a daily probe schedule, as shown in Figure 12-48.

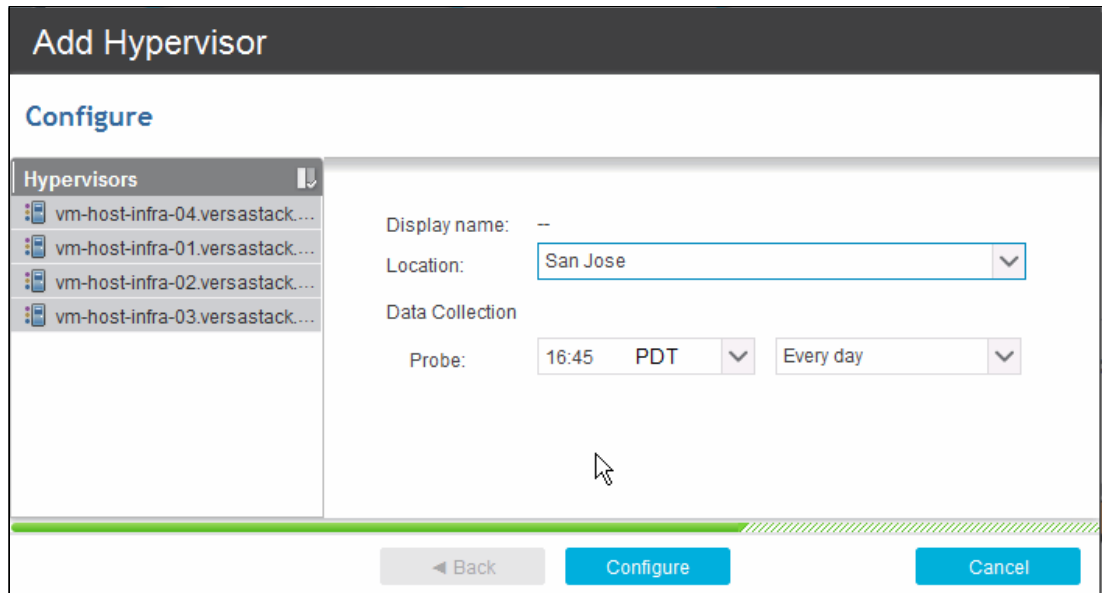


Figure 12-48 Hypervisor probe schedule

Back in the VSC Hypervisors section, you can immediately start the probe of the discovered hypervisors, as shown in Figure 12-49.

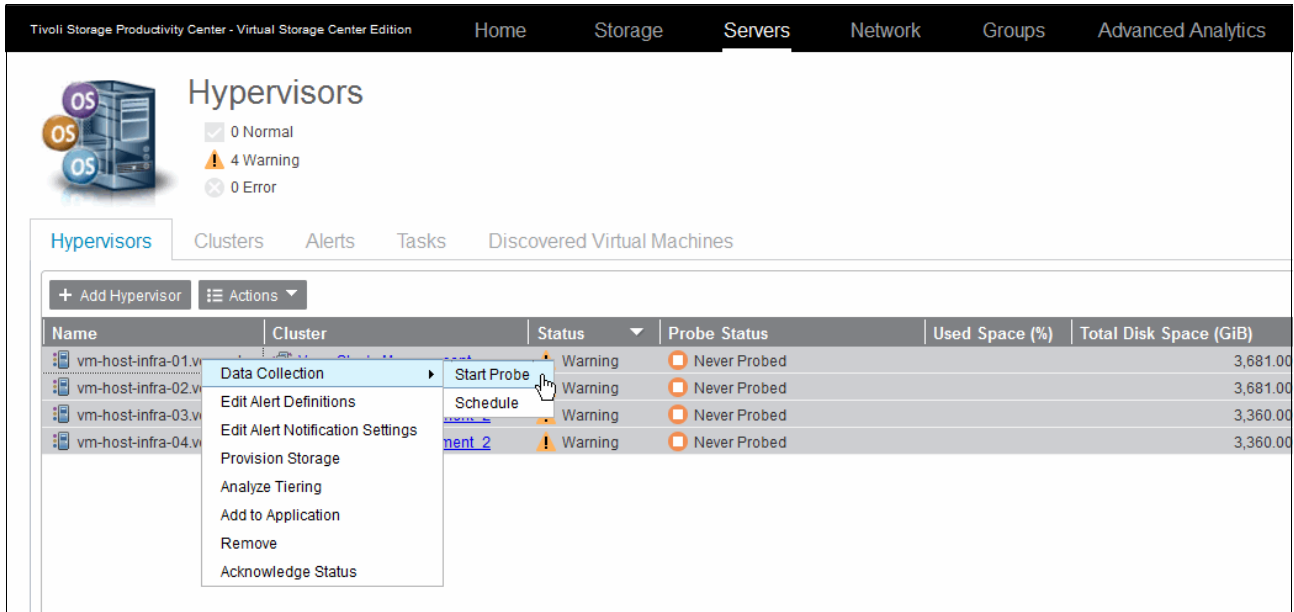


Figure 12-49 Start probe manually

After the probe is started, you can follow the progress by opening the probe logs, as shown in Figure 12-50.

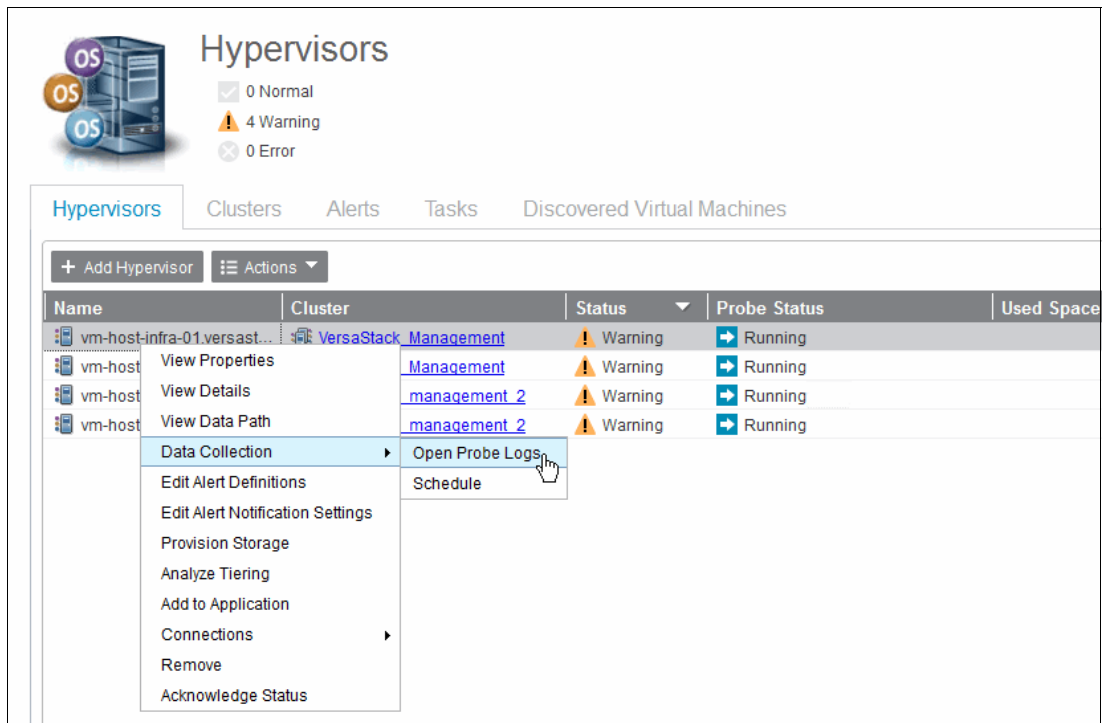


Figure 12-50 Check probe logs

Depending on the resource on which the probe is run there are several stages of the probing to be run, as shown in Figure 12-51.

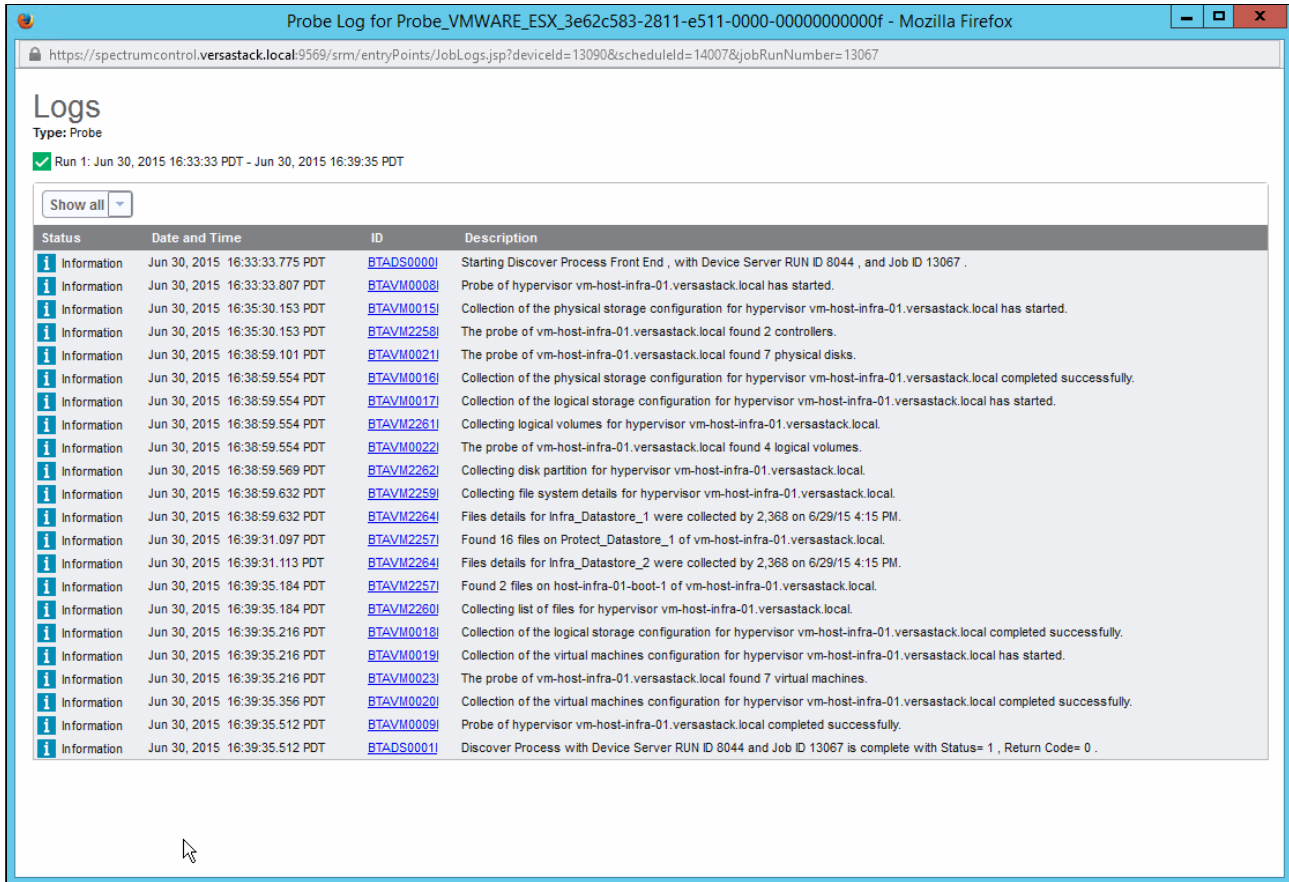


Figure 12-51 Probe results

The ESXi hypervisors of the VersaStack infrastructure are now added to the VSC environment.

Spectrum Control hypervisor overview

Similar to the Storwize V7000 storage system that you registered before, you explore the different panes and information that VSC provides in the Web GUI. Start the VSC Web GUI and select **Servers** → **Hypervisors** from the main menu, as shown in Figure 12-52.

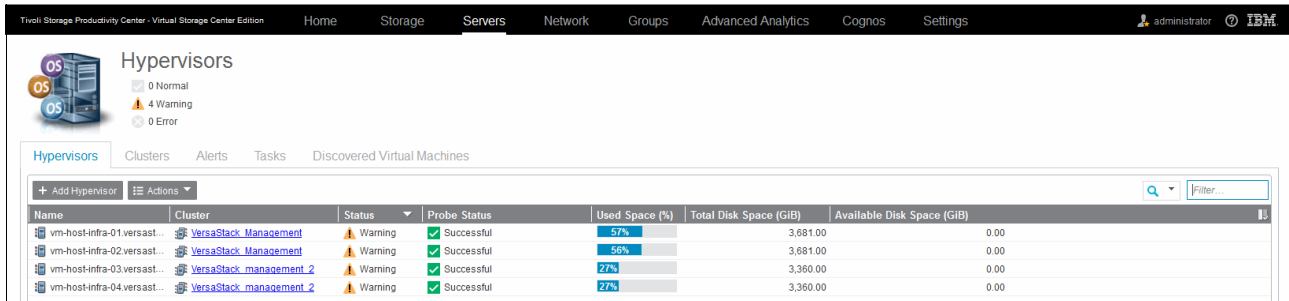


Figure 12-52 Hypervisor overview

Throughout all the components in the Web GUI, a similar approach is taken to outline the information by using tabs. The tabs for the Hypervisor are grouped into the following categories:

- ▶ Hypervisors: Lists all the discovered hypervisors and allows you to open the individual hypervisor's overview windows.
- ▶ Clusters: Groups the hypervisor per cluster if deployed that way in the vCenter.

Figure 12-53 shows the VSC Hypervisor overview of discovered clusters and associated resources.

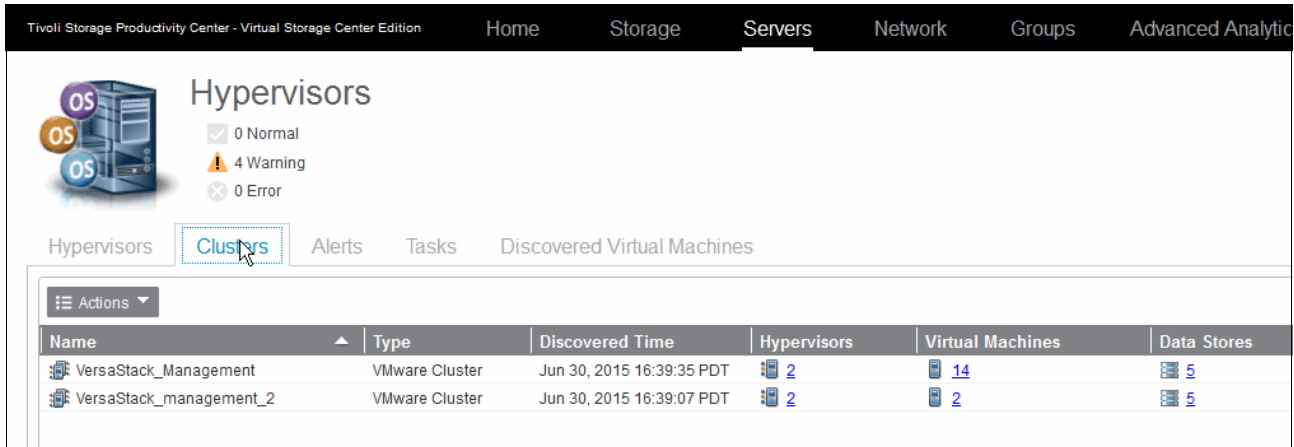


Figure 12-53 Hypervisor cluster overview

- ▶ Alerts: Filters all alerts that are related to the hypervisors.
- ▶ Tasks: Shows tasks such as provisioning and storage tiering for the affected hypervisors.
- ▶ Discovered Virtual Machines: Lists all VMs that were discovered since the last probe, which allows you to perform agentless registration of these VM for logical grouping and reporting purposes.

Figure 12-54 shows that the VSC Hypervisor discovered the VMs and sorted them by name.

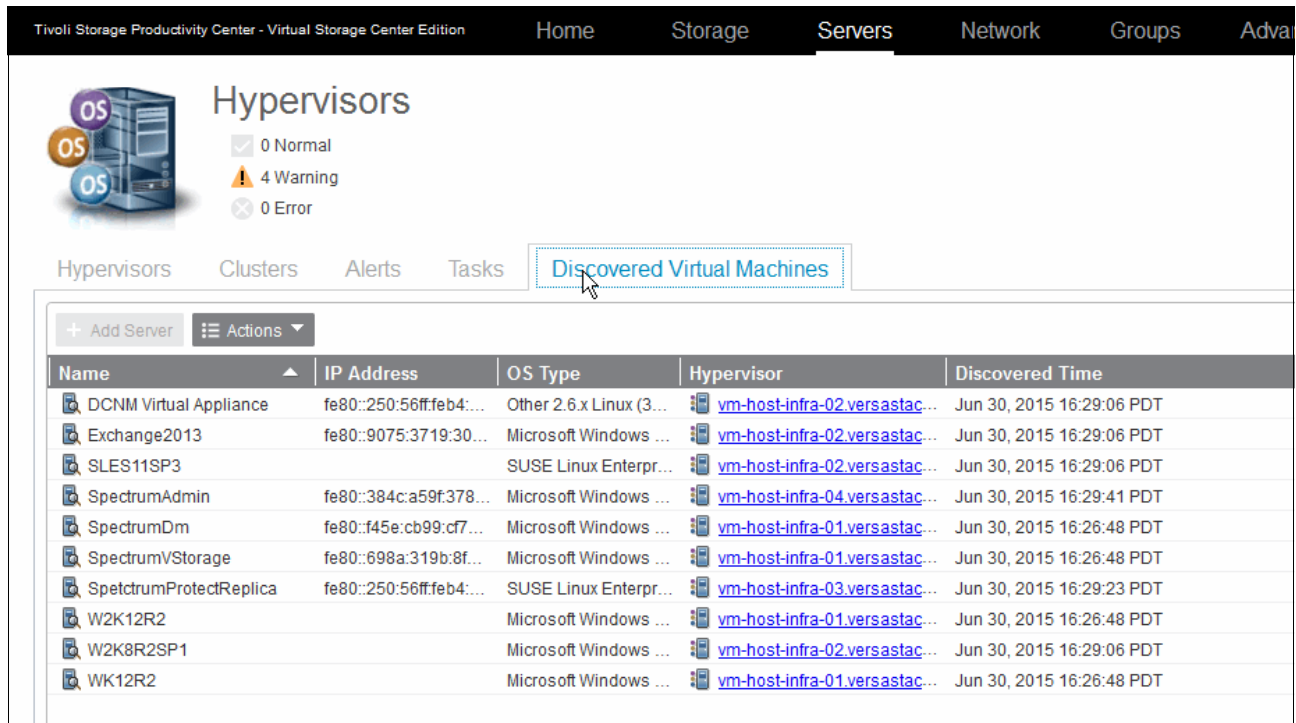


Figure 12-54 Discovered virtual machines

The hypervisors in Figure 12-52 on page 305 are in a warning state. However, no alerts are triggered from the VSC perspective. Checking the properties from the General Resources menu indicates that the system was in a warning status. Connecting to the vCenter environment shows that the hypervisors are indeed in a warning state because the SSH services were enabled, as shown in Figure 12-55.




Figure 12-55 ESX SSH warning

For a case like this one where system warnings are received for conditions that you are aware of and that you want to ignore, you can acknowledge the alerts for those specific resources. Here, acknowledge the ESXi warning state, as shown in Figure 12-56.

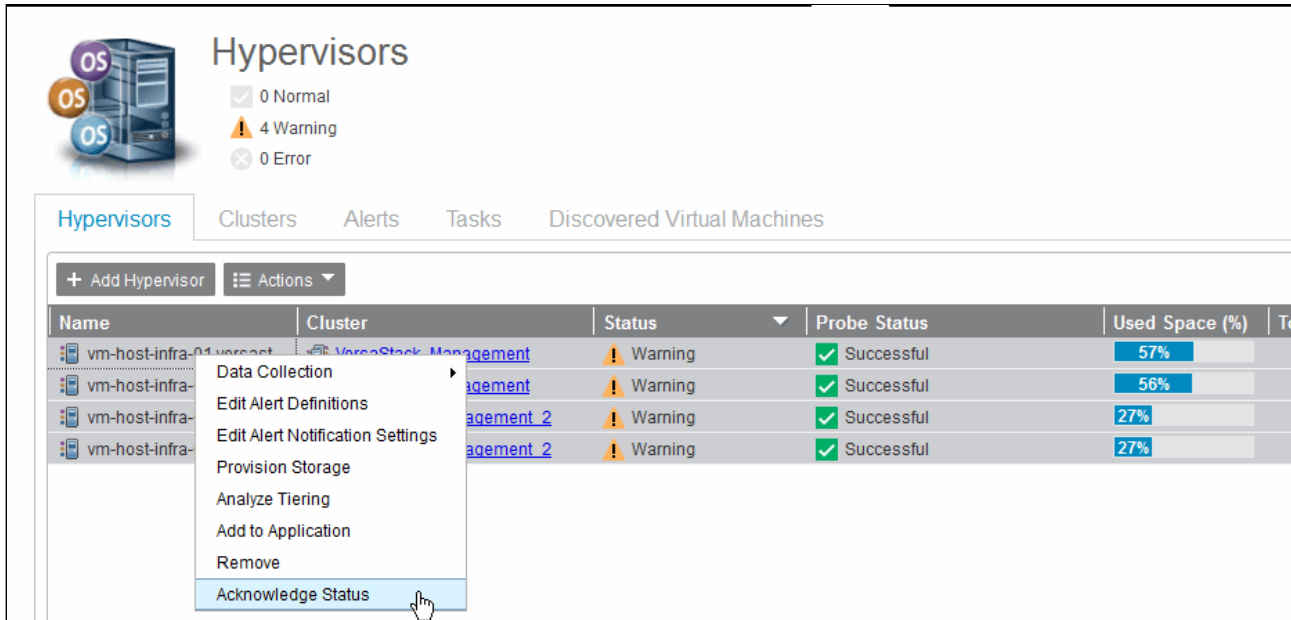


Figure 12-56 Acknowledge warning status

With the hypervisors correctly configured, open the overview of vm-host-infra-01 by double-clicking its entry from the Hypervisors overview, as shown in Figure 12-57.

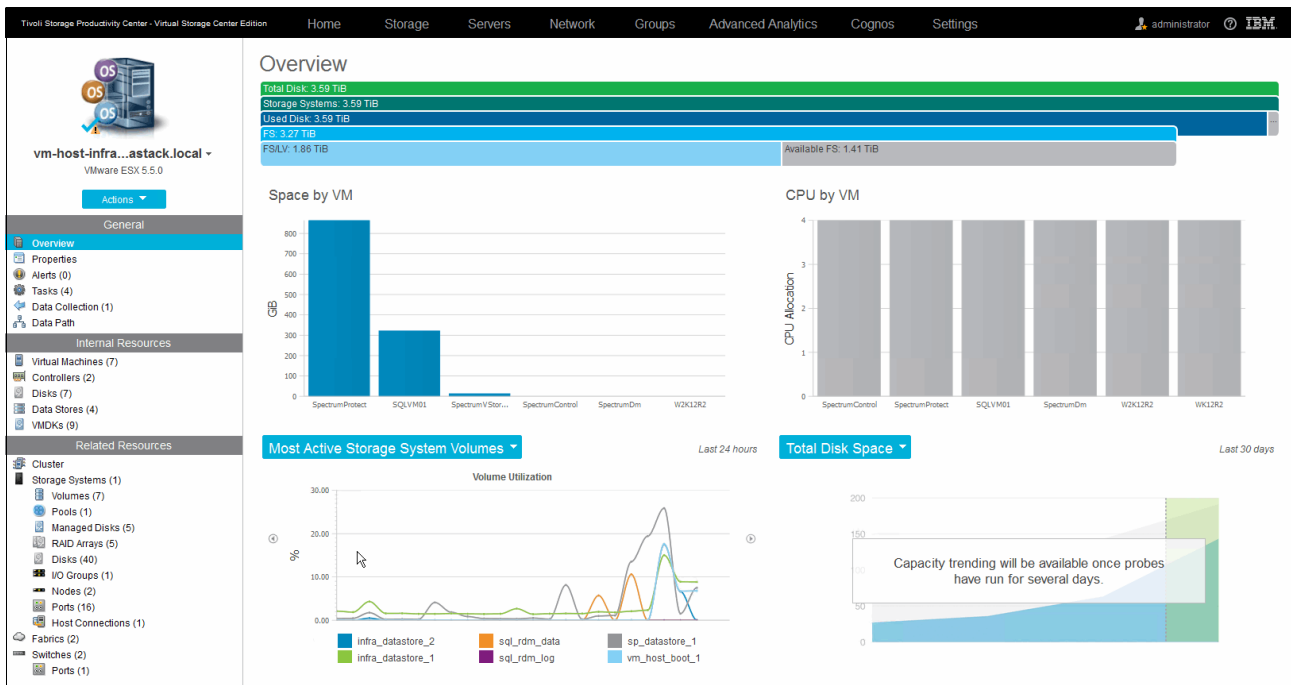


Figure 12-57 Hypervisor overview

Throughout the whole VSC GUI, a unified approach is used to chart graphical data and to group resources for the selected device into three categories: General, Internal Resources, and Related Resources. As we did for the Storwize V7000 storage system, we go over the sections individually and highlight some of them that are of interest for the current setup.

► **General:**

- Overview: Brings you back to the graphical charts. These charts can be toggled and provide summarized data for the following categories:
 - Space by VM
 - CPU by VM
 - Most Active Storage System Volumes
 - Total Disk Space
 - Space from Storage Systems
 - Memory by VM
 - Most Active Switch Ports
- Properties: Provides a summarized overview of the hypervisor, including details such as VMware ESX version, hardware serial number, and model and storage capacity.

Figure 12-58 shows the VSC Hypervisor overview showing the Cisco UCS B200-M4 serial information.

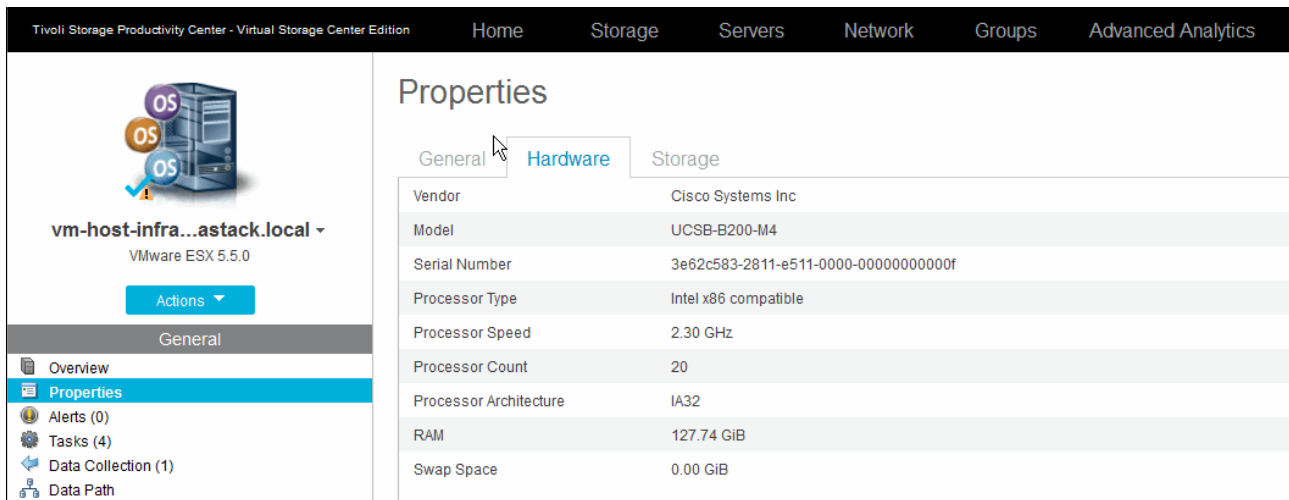


Figure 12-58 Hypervisor hardware properties

- Alerts: Groups hypervisor-related alerts here for this specific hypervisor. For examples of configuring some sample alerts and general notifications overrides, see “Spectrum Control hypervisor monitoring and alerting” on page 317.

- Tasks: Shows provisioning tasks that are completed, are scheduled to be run, or are awaiting execution approval.

Figure 12-59 shows the VSC Hypervisor overview of tasks for the selected hypervisor.

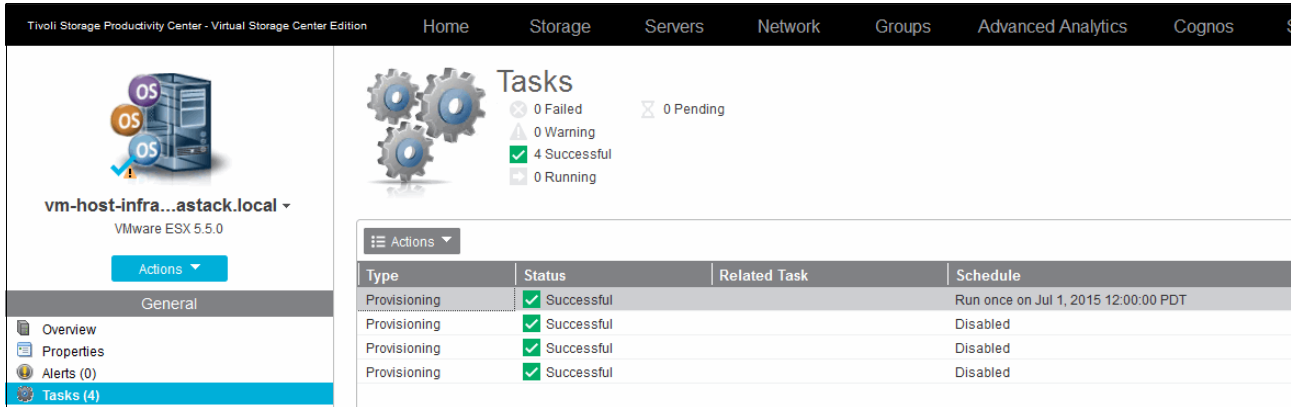


Figure 12-59 Hypervisor tasks overview

- Data Collection: Allows you to verify and control the probe settings for this specific hypervisor. You can modify existing probe schedules or start an immediate probe.

Figure 12-60 shows the VSC Hypervisor data collection options.

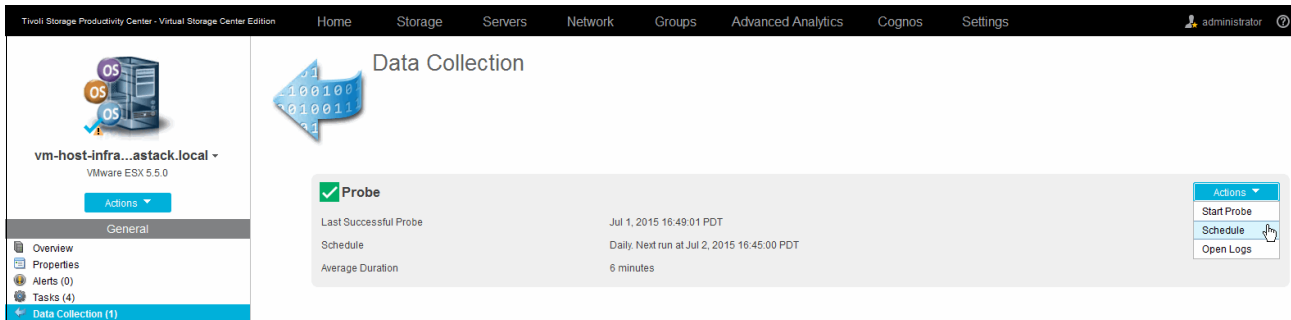


Figure 12-60 Hypervisor data collection

- Data Path: Outlines the data path for all related resources to the hypervisor. For more information, see the Data Path Topology View bullet on page 285.

Figure 12-61 on page 311 shows the VSC Hypervisor data path overview for the selected system.

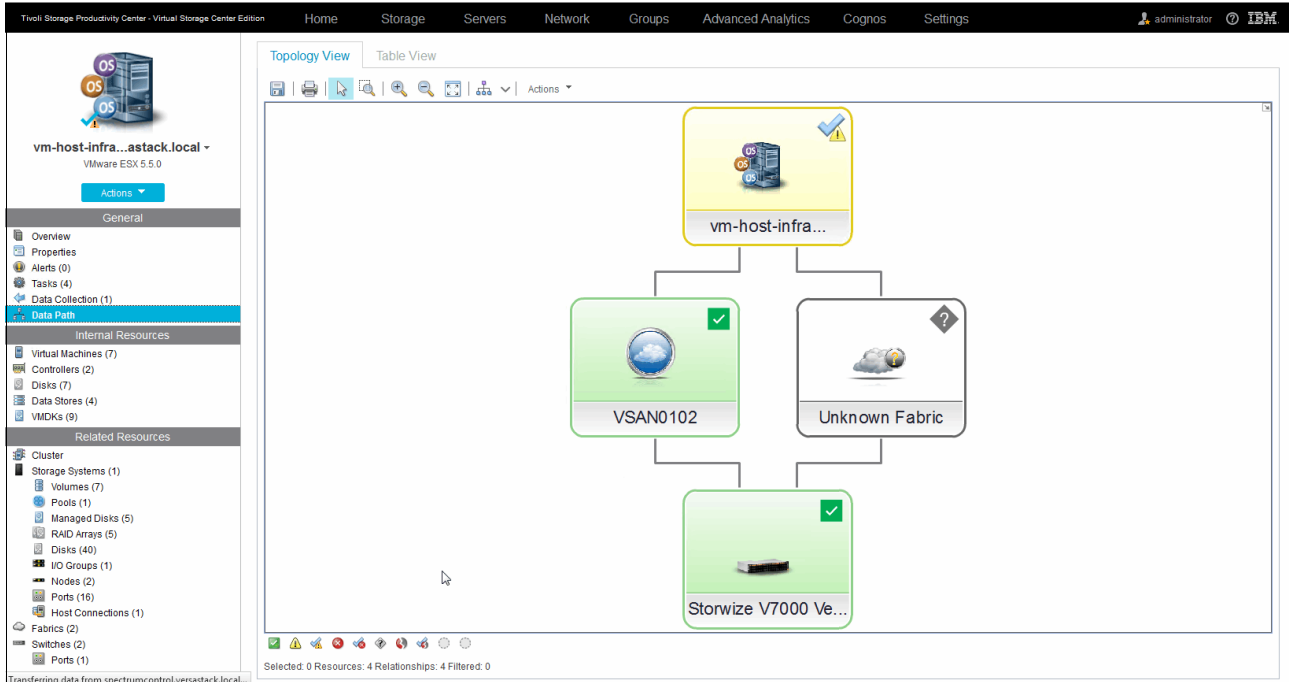


Figure 12-61 Hypervisor data path

► Internal Resources:

- Virtual Machines: Lists all VMs for the selected hypervisor, including metrics that are retrieved from the vCenter/Hypervisor, such as configuration files, number of vCPUs, and assigned RAM.

Figure 12-62 shows the VSC Hypervisor VMs that are grouped for the selected system.

Name	Status	Agent S...	Data Stores	Configuration File	OS Type	Processor Count	RAM (GiB)	Capacity (...)
SpectrumControl	Normal	Normal		SpectrumControl.vmx	Microsoft Windows Server 2012 (...)	4	4	
SpectrumDm	Normal			SpectrumDm.vmx	Microsoft Windows Server 2012 (...)	4	4	
spectrumprotect	Normal	Normal	Protect_Datast...	SpectrumProtect.vmx	SUSE Linux Enterprise 11 (64-bit)	4	4	
SpectrumVStorage	Normal		Protect_Datast...	SpectrumVStorage.vmx	Microsoft Windows Server 2008 ...	2	4	
SQLVM01.VersaStack.local	Normal	Normal		SQLVM01.vmx	Microsoft Windows Server 2012 (...)	4	4	
W2K12R2	Normal			W2K12R2-Template.vmx	Microsoft Windows Server 2012 (...)	4	4	
WK12R2	Normal			WK12R2.vmx	Microsoft Windows Server 2012 (...)	4	4	

Figure 12-62 Hypervisor virtual machines

The column view can be modified to display only a subset of data, as shown in Figure 12-62.

Note: This view is updated every time the probe for the selected hypervisor is run. If you have an environment where the VMs migrate often between hosts of a DRS-enabled cluster, as in our example setup, you might want to increase the frequency of the probing or run an *ad hoc* probe through the Data Collection entry in the General section.

- Controllers: Shows the internal storage controllers for the hypervisor and their data, such as the HBA WWN and associated disks.

Figure 12-63 shows the VSC Hypervisor controllers overview window, which lists the HBA WWNs.




Figure 12-63 Hypervisor controllers

- Disks: Provides an overview of all physical disks being used by this hypervisor and the data paths towards them. The Storwize V7000 storage virtualization engine groups these disks within MDisk arrays, which are themselves grouped into pools. In these pools, volumes are created and mapped to hosts. From the Data Path tab, you can observe the correlation between the virtual volumes and the underlying physical disks and how they are shared across multiple volumes.

Figure 12-64 shows the VSC Hypervisor disk, datastore, capacity, and other metrics.

The screenshot displays the 'Disks' page in the Tivoli Storage Productivity Center. The top navigation bar includes 'Home', 'Storage', 'Servers', 'Network', 'Groups', 'Advanced Analytics', 'Cognos', and 'Settings'. The left sidebar shows a tree view of resources, with 'Disks (7)' selected under 'Internal Resources'. The main content area shows a summary of disk status (7 Normal, 0 Warning, 0 Error) and a table of disk details.

Name	Status	Paths	Vendor	Mo...	Firmware	Serial Number	Capacity (GiB)	Available Disk Space (GiB)	Used Space (GiB)	Data Store
naa.6005076...	Connected	Z	IBM	2145	0000	0000000000000007	2,048.00	0.00	2,048.00	Protect_Data...
naa.6005076...	Connected	Z	IBM	2145	0000	0000000000000006	1.00	0.00	1.00	
naa.6005076...	Connected	Z	IBM	2145	0000	0000000000000005	64.00	0.00	64.00	
naa.6005076...	Connected	Z	IBM	2145	0000	0000000000000004	256.00	0.00	256.00	
naa.6005076...	Connected	Z	IBM	2145	0000	0000000000000002	1,024.00	0.00	1,024.00	Infra_Datasto...
naa.6005076...	Connected	Z	IBM	2145	0000	000000000000000b	256.00	0.00	256.00	Infra_Datasto...
0200000000...	Connected	Z	IBM	2145	0000	0000000000000000	32.00	0.00	32.00	host-infra-01...

Figure 12-64 Hypervisor disks

- Datasets: Lists the datasets for the hypervisor together with data such as Used Space %, Available dataset Space, and VMDKs. Double-clicking a dataset brings you to the Properties notebook for that dataset, where you can see more information about the VMDKs, such as Virtual Machine, Volume, and Hypervisor. This correlated click-through is consistent within the whole VSC Web GUI, allowing you to explore the data from your environment in an intuitive manner.

Figure 12-65 shows the VSC Hypervisor datasets overview with three available data set options.

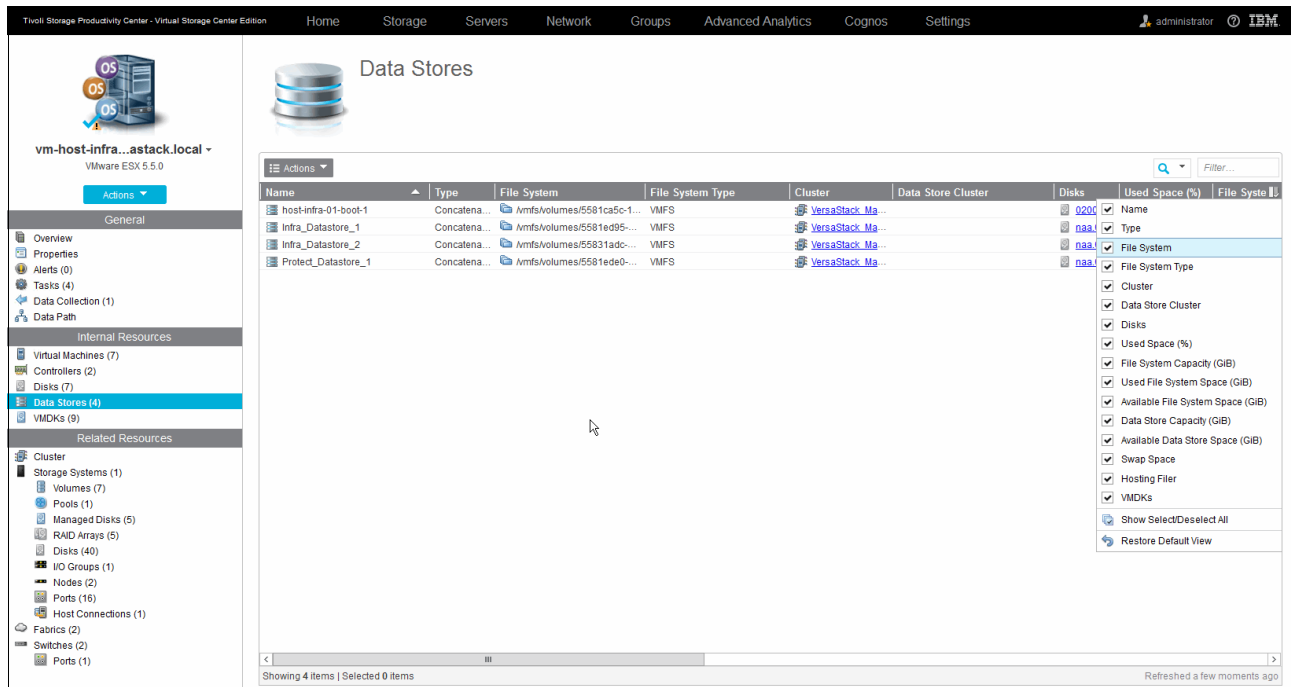


Figure 12-65 Hypervisor datasets

- VMDKs: Lists the VMDKs for the VMs running at the specific hypervisor at the time of the data collection probe.
- ▶ Related Resources:
 - Cluster: Opens the cluster notebook with several tabs.
 - General: Name of cluster and last probe time stamp.

Figure 12-66 shows the VSC Hypervisor overview at a cluster level.

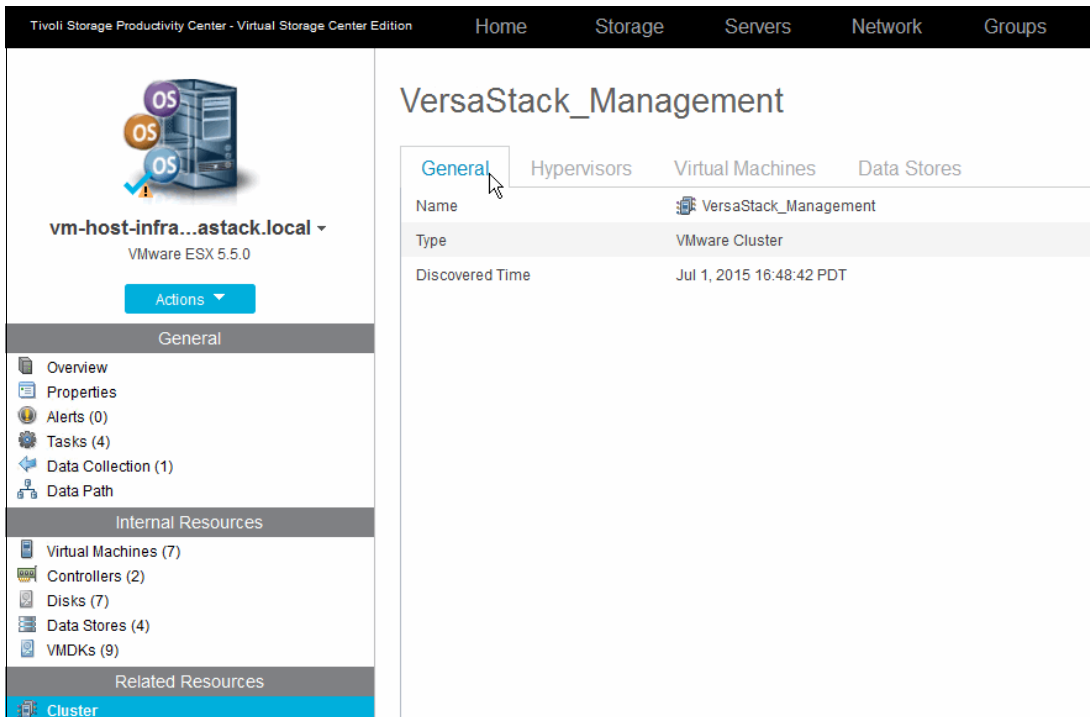


Figure 12-66 Hypervisor cluster overview

Figure 12-67 shows the VSC Hypervisor overview of members of the hypervisor cluster.

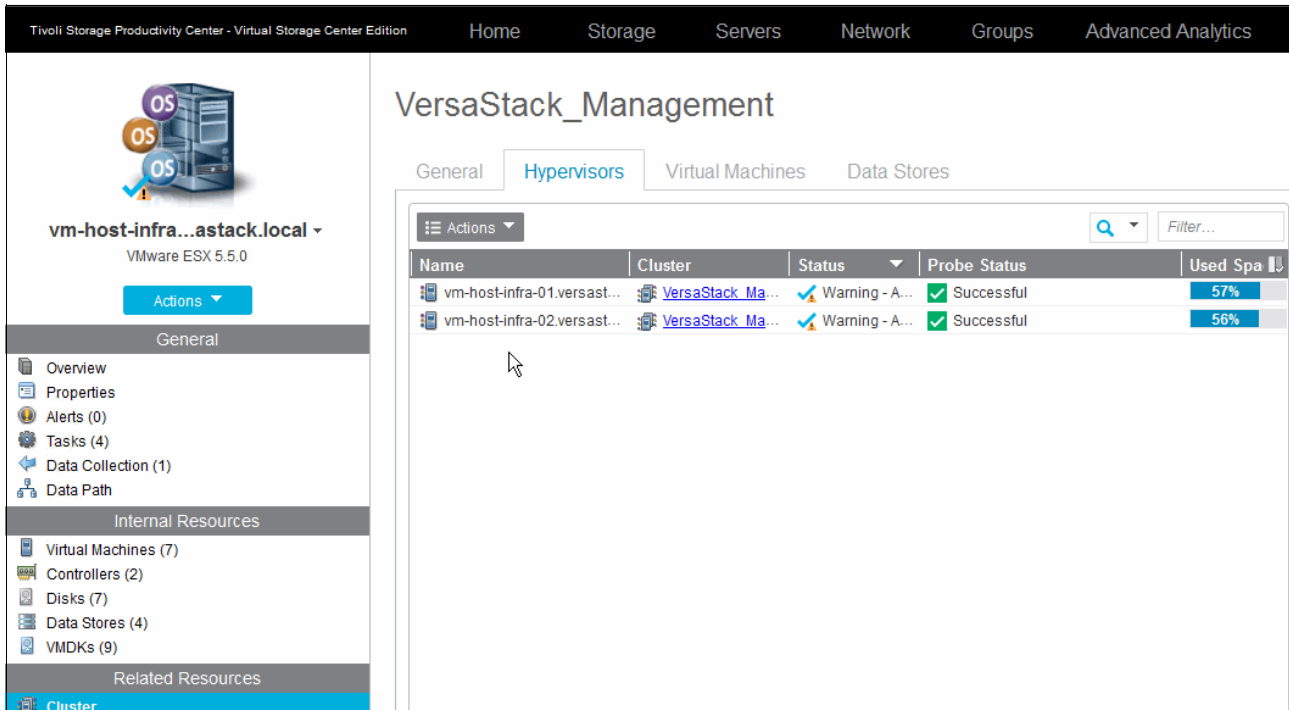


Figure 12-67 Hypervisor cluster hypervisor members

Figure 12-68 shows the VSC Hypervisor overview of VMs that are associated with the cluster.

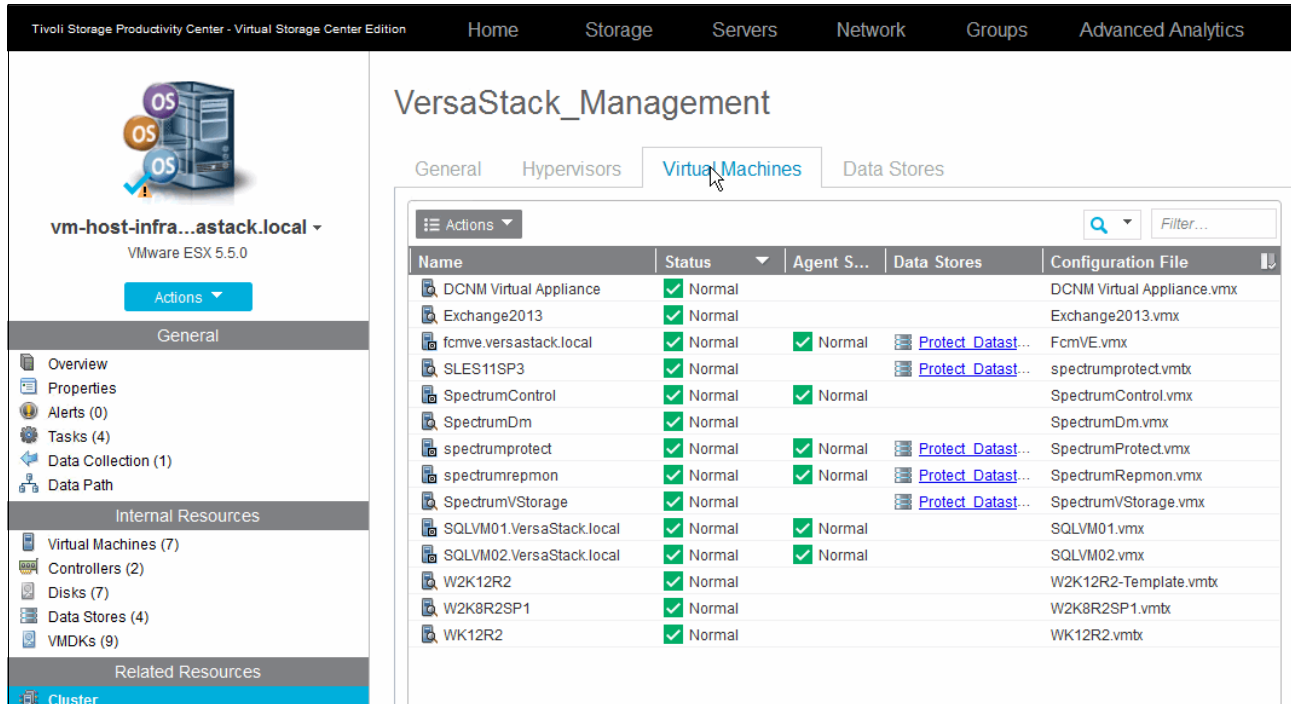


Figure 12-68 Hypervisor cluster VM overview

Figure 12-69 shows the VSC Hypervisor Cluster overview of datastores that are attached to the cluster.

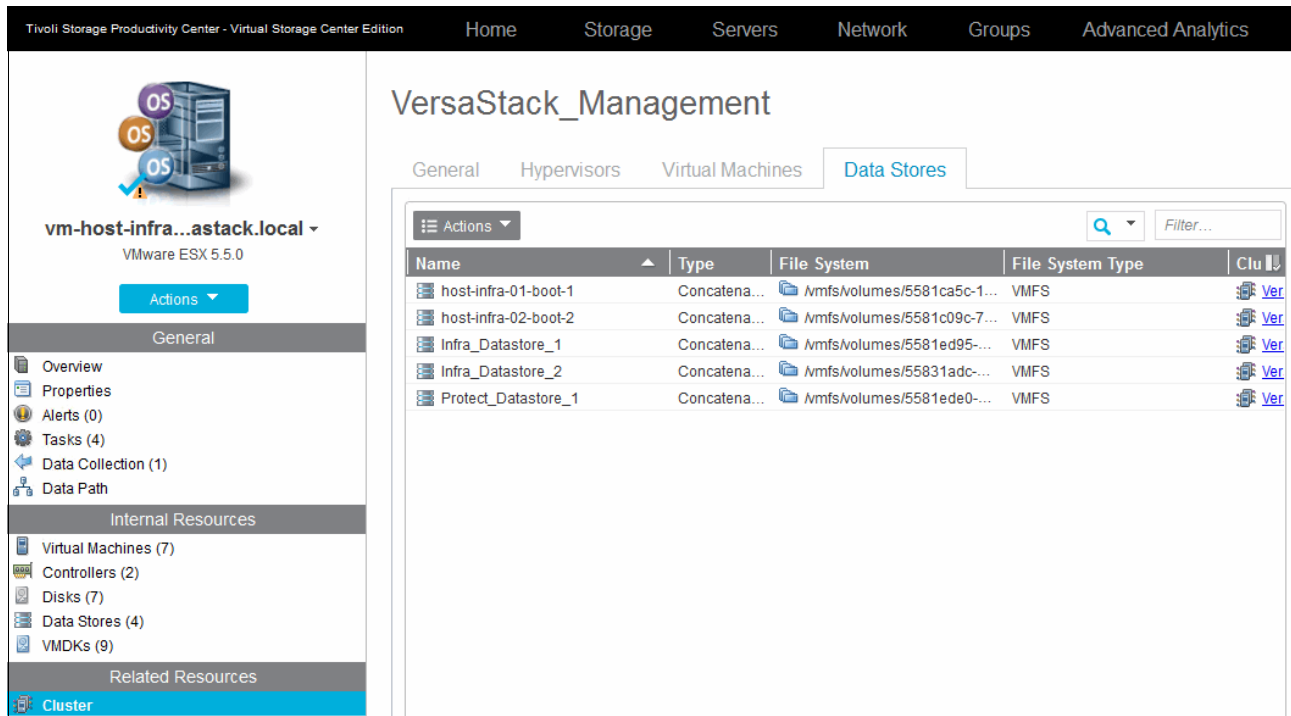


Figure 12-69 Hypervisor cluster datastores

- Storage Systems: Groups the storage resources for this hypervisor per storage system. For more information, see “Managing the storage infrastructure” on page 280.
- Fabrics: Shows the fabrics to which the hypervisor is connected.
- Switches: Shows the fabric member switches and the ports to which the hypervisor is connected.

Figure 12-70 shows the VSC Hypervisor switch port overview.

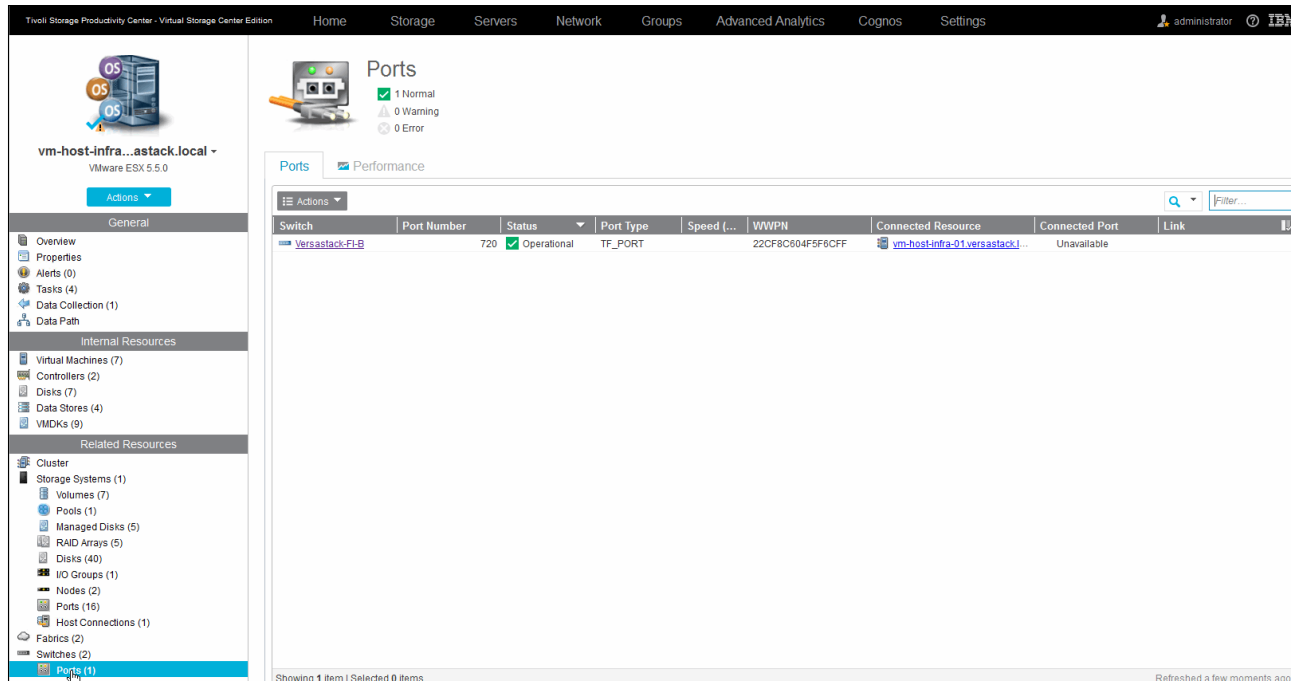


Figure 12-70 Hypervisor switch ports

Spectrum Control hypervisor monitoring and alerting

The Tivoli Storage Productivity Center Virtual Storage Edition Web GUI provides direct insights into the VMware Hypervisor environment that is deployed on the VersaStack infrastructure.

Through the screen captures that are shown in “Integrating the VMware vCenter Hypervisor with Spectrum Control” on page 302, the VSC shows you metrics such as the following ones:

- ▶ Storage capacity that is assigned to the hypervisors, and storage space that is used by the VMs
- ▶ Most active volumes and switch ports for these hypervisors
- ▶ Performance of the datastore volumes on the Storwize V7000 storage system and the Fibre Channel switch ports

This data is captured and stored in the underlying DB2 data warehouse that is integrated in the Spectrum Control VSC and used by the Cognos Business Intelligence Reporting engine. It can then be used for *ad hoc* or scheduled reporting, as described in 12.8.4, “Reporting for departments and applications” on page 350.

Monitoring performance

From within the VMware vCenter hypervisor management console, you will monitor performance and capacity aspects from a cluster, hypervisor, or single VM perspective. These performance metrics focus primarily on CPU, Memory, Network, and Disk.

At a disk level, you can review, for example, read/write rate and latency for the underlying physical disks for a specific hypervisor, or at physical disk, the datastore or VDMK level for individual VMs.

However, you can have only the most detailed performance metrics in real time from within the VMware vCenter.

Figure 12-71 shows the VMware vCenter sample real-time performance chart for `infra_datastore_1`.

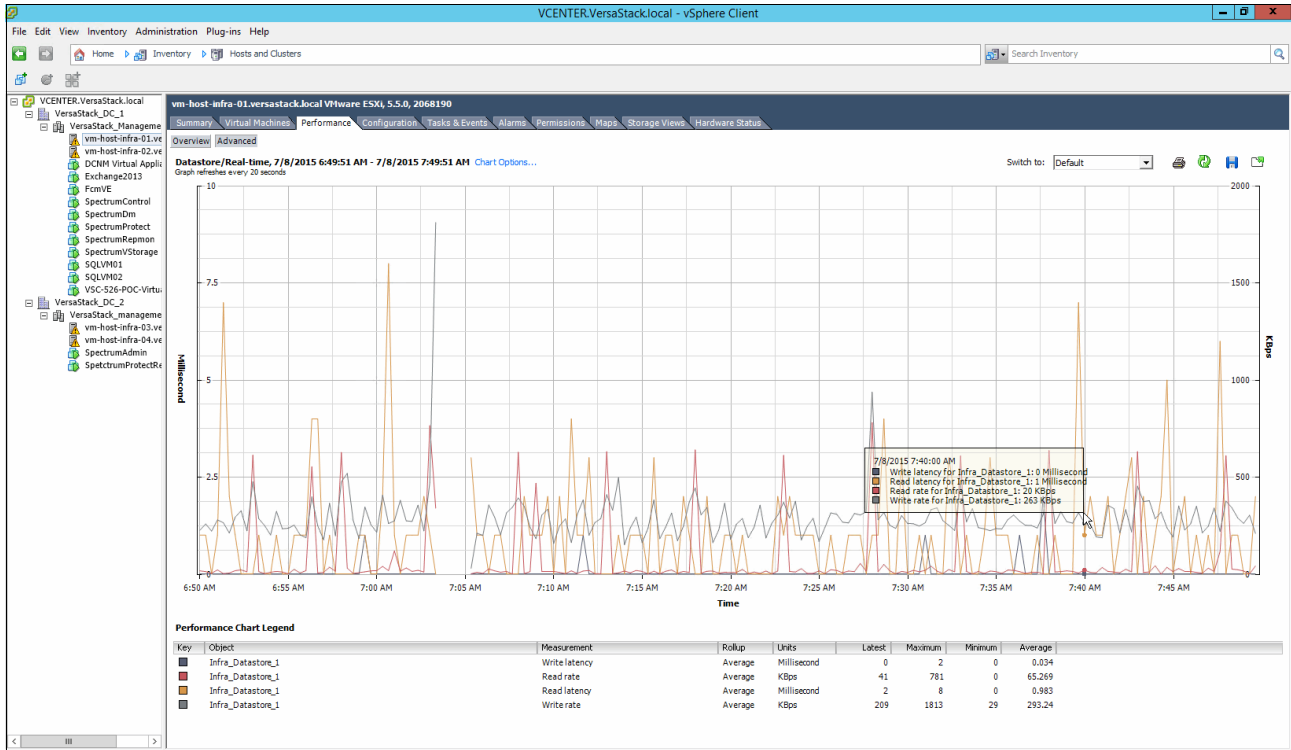


Figure 12-71 vCenter datastore real-time performance

Within VSC, you can look at the same performance in a view that encompasses both real-time and historical data, as shown in Figure 12-72.

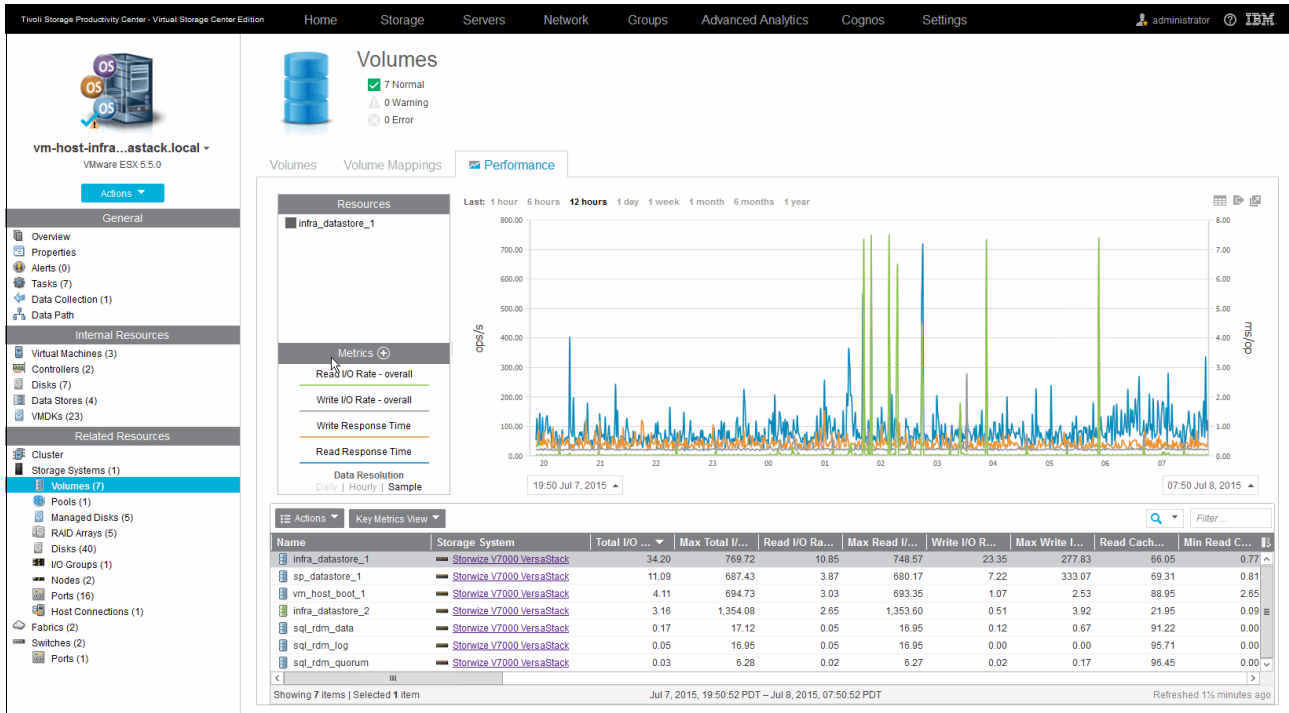


Figure 12-72 VSC datastore real-time performance

You can dynamically narrow or expand the scope from 1 minute to up to the maximum period that you specified in the VSC retention settings. In this view, you can easily toggle between 1/6/12 hour periods or day/month/6 months/year.

You also have access to additional storage hardware-related metrics, such as Cache to Disk and Disk to Cache, and you can select multiple volumes to be overlaid in a single graph. Figure 12-73 is the same graph as in Figure 12-72 on page 319, but with a second datastore added and the scope expanded to 1 month.

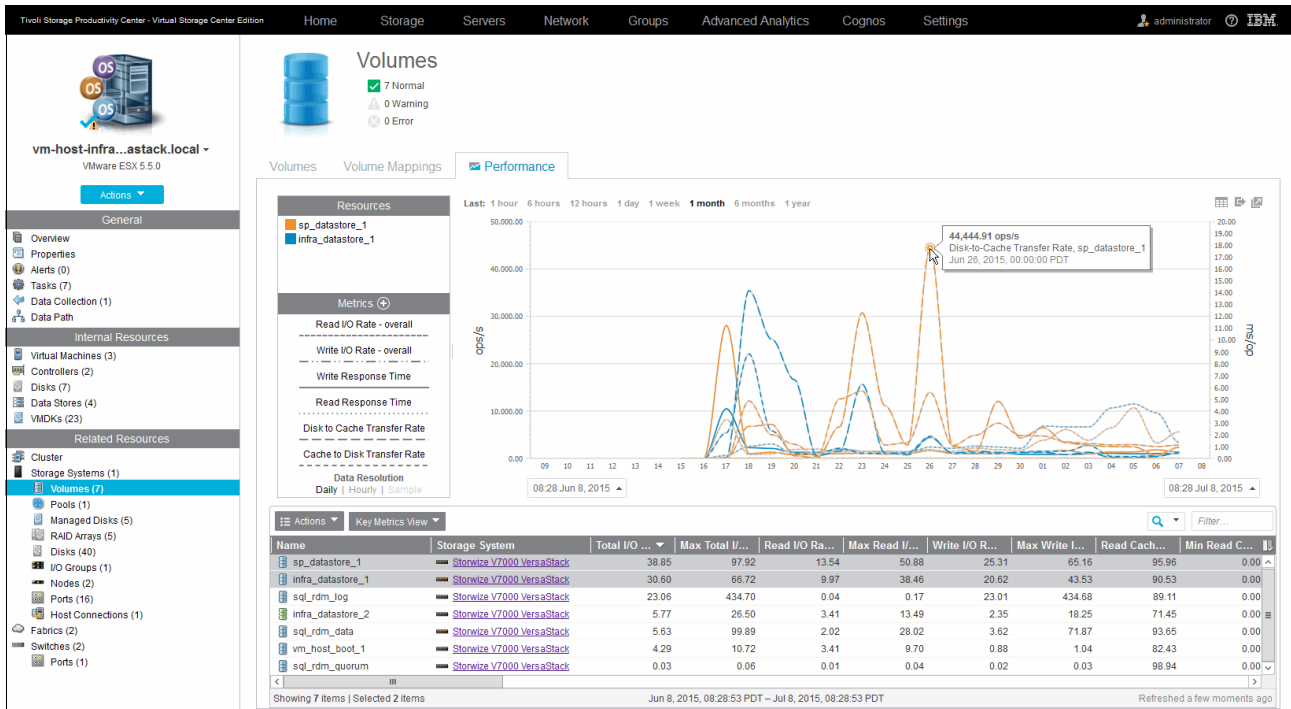


Figure 12-73 VSC multiple datastores 1-month overview

You can also use the vSphere Web Client extension to monitor performance by using storage system metrics, which gives the VMware administrator access to this information from a familiar working environment. However, the granularity is limited to a fixed 1-hour, 1-day, or 1-week interval when you access the information.

Figure 12-74 shows the VMware vSphere Web Client VASA that is provided by the Storage System Metrics performance chart.

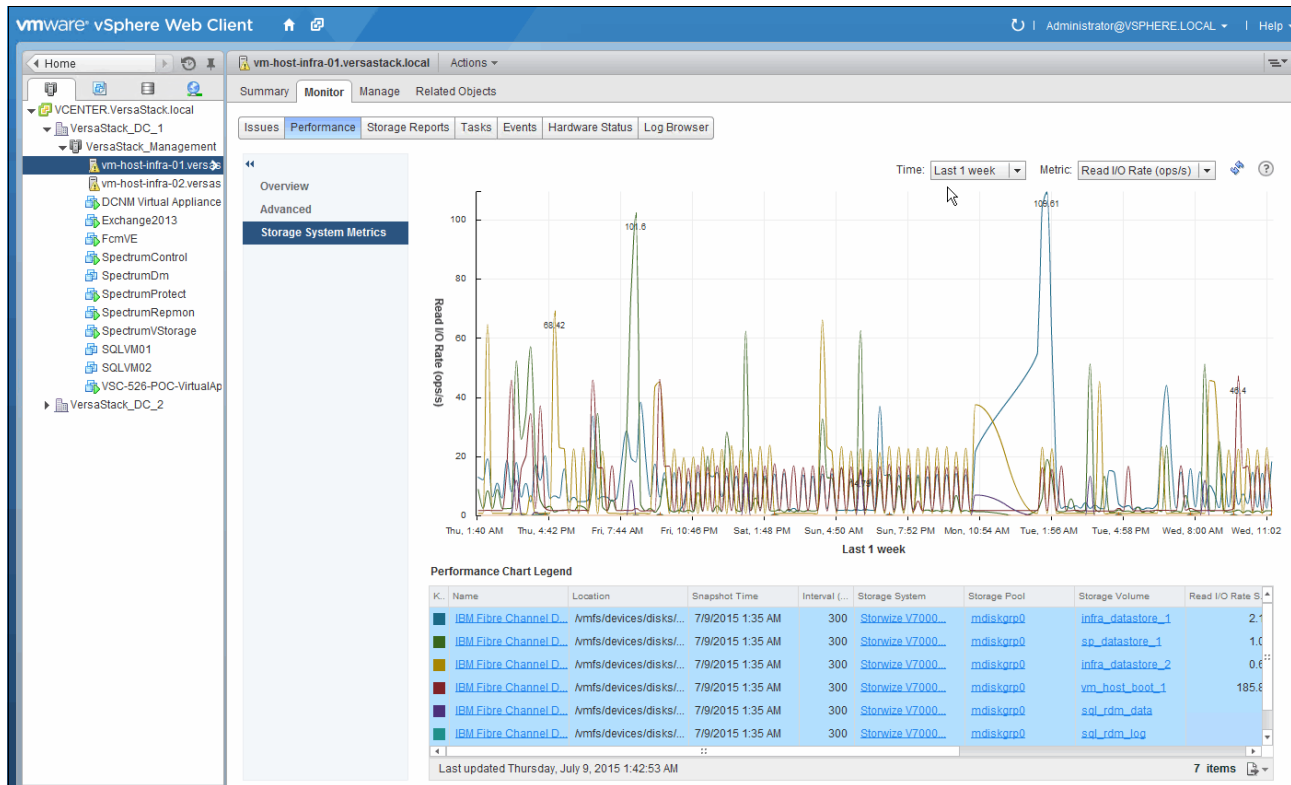


Figure 12-74 vSphere Web Client storage systems metrics performance

In addition to monitoring, you can also view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems.

In our example, we registered the VSC as a VASA provider to the vSphere Web Client, which provides the following vSphere reports to view information about your virtual resources and the back-end storage systems:

- Fabric Connections

This storage report displays fabric information that includes zone and switch details in the vSphere Web Client.

- Storage Mapping

This Storage Mapping report displays end-to-end mappings between back-end storage that is monitored by Tivoli Storage Productivity Center and the virtual resources that are monitored by vSphere.

- Storage System Metrics

This performance report displays performance metrics that include the total I/O rate and response time for the back-end storage systems that are monitored by Tivoli Storage Productivity Center, and that performance is running on the storage system.

► **SCSI Volumes (LUNs)**

This volume report displays block storage information that is provided by Tivoli Storage Productivity Center. The information includes the following details:

- Space that is committed to the volume
- Thin-provisioning status
- System capability
- Storage array name
- Volume identifier on the array
- Namespace of the VASA provider

► **Datstores**

The Datstores report includes file system information that is provided by Tivoli Storage Productivity Center, including the system capability and the namespace of the VASA provider.

► **Capacity**

The Hypervisor Overview pane in the VSC shows the capacity that is used by the VMs on the specific hypervisor.

Figure 12-75 shows the VSC Hypervisor Overview listing space by VM.

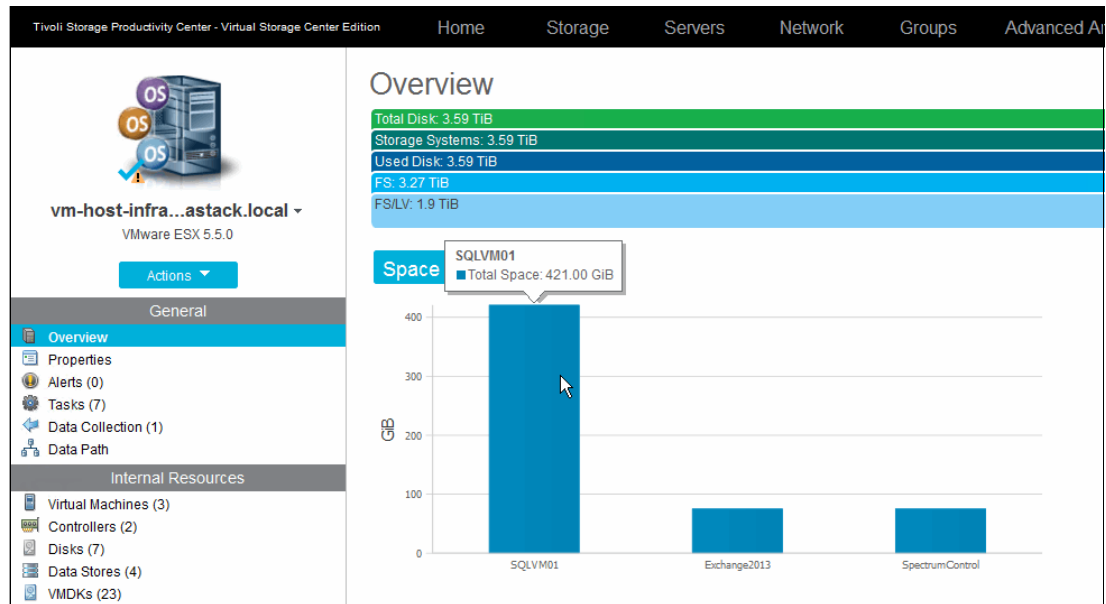


Figure 12-75 VSC Hypervisor Space by VM

Similarly, this information can be obtained through the VSC extension to the vSphere Web Client.

Figure 12-76 shows the VMware vSphere Web Client VASA-provided storage reports.

SCSI ID	Datastore	Capacity	System Capability *	Storage Array *	Identifier on Array *
020004000060050764008180c...	Protect_Datastore_1	2.00 TB	EasyTier,Replication	Stonwize V7000 VersaStack	sp_datastore_1
020002000060050764008180c...	Infra_Datastore_2	256.00 GB	EasyTier,Thin,Replication	Stonwize V7000 VersaStack	infra_datastore_2
020001000060050764008180c...	Infra_Datastore_1	1.00 TB	EasyTier,Replication	Stonwize V7000 VersaStack	infra_datastore_1
020000000060050764008180c...	host-infra-02-boot-2	32.00 GB	EasyTier,Replication	Stonwize V7000 VersaStack	vm_host_boot_2
020000000060050764008180c...	host-infra-01-boot-1	32.00 GB	EasyTier,Replication	Stonwize V7000 VersaStack	vm_host_boot_1
020006000060050764008180c...		1.00 GB	EasyTier,Replication	Stonwize V7000 VersaStack	sql_rdm_quorum
020003000060050764008180c...		256.00 GB	EasyTier,Replication	Stonwize V7000 VersaStack	sql_rdm_data
020005000060050764008180c...		64.00 GB	EasyTier,Replication	Stonwize V7000 VersaStack	sql_rdm_log

Figure 12-76 vSphere Web Client storage reports

Alert configuration

For the hypervisors, you can define the following alert triggers:

- ▶ Hypervisor
 - Server Status Change Offline
 - Server Status Change Online
 - Hypervisor Missing
 - Probe Failed
- ▶ File Systems
 - File System Discovered
 - File System Low on Free Space
- ▶ Disks
 - Disk Discovered
 - Disk Defect Discovered
 - Disk Failure Predicted

Figure 12-77 shows the VSC Hypervisor where we define a File System Low on Free Space alert. It shows how to trigger an alert when the file system has 5% free space left. Instead of a percentage value, fixed data sizes can be used.

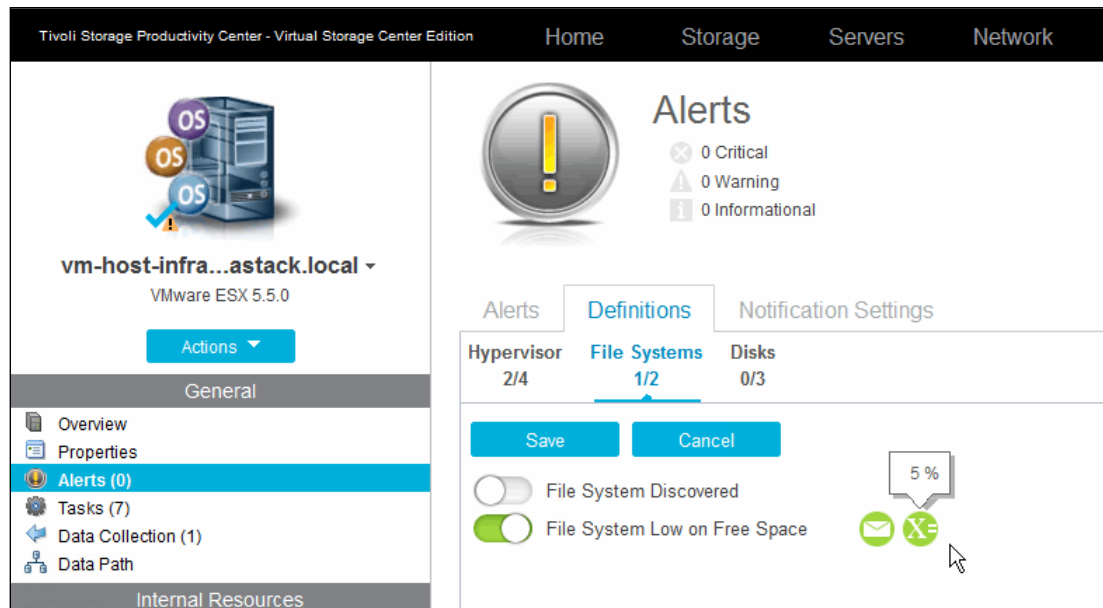


Figure 12-77 VSC Hypervisor alert definitions

12.8 Advanced Analytics

The Advanced Analytics feature that is built into Tivoli Productivity Center Virtual Storage Edition transforms the VersaStack environment on a cloud-enabled environment in the following ways:

- ▶ Defines service classes for the storage requirements
- ▶ Provides self-provisioning to servers and hypervisors
- ▶ Optimizes the placement of new volumes at volume creation and during the data lifecycle of that volume

12.8.1 Cloud Configuration

In the Cloud Configuration tab of the VSC Web GUI, you assign storage to tiers, define service classes, and create capacity pools. This section uses a *learn the concepts* overview within the GUI itself that you can use to become familiar with this function. The overview outlines the required steps.

Figure 12-78 shows the VSC Advanced Analytics Learning the Concepts built-in tutorial.

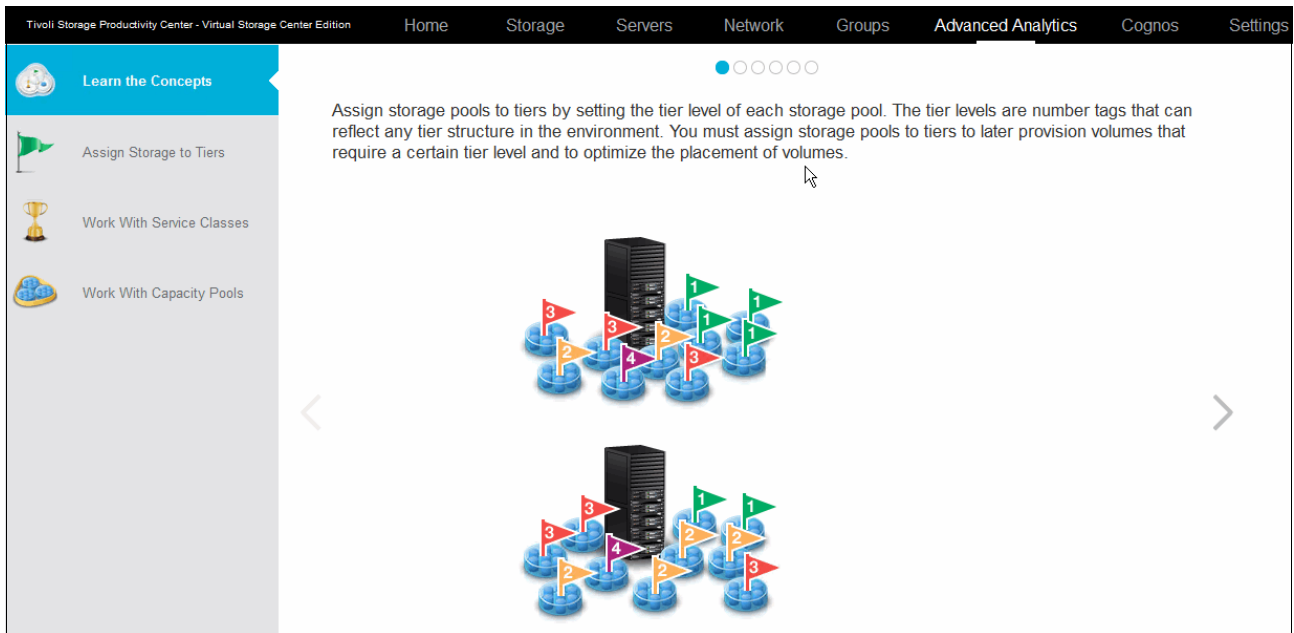


Figure 12-78 VSC Advanced Analytics - Learning the Concepts

Storage tiers are defined at a Storwize V7000 pool level. In our example setup, we have SSD- and SAS-based storage pools that are assigned to tier 1 and tier 2.

Figure 12-79 shows the VSC Advanced Analytics defined storage tiers in the VersaStack environment.

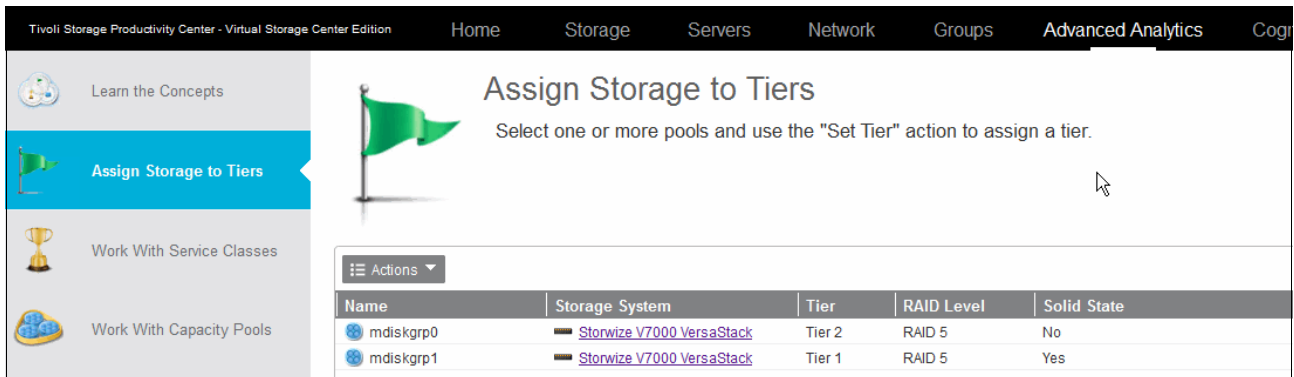


Figure 12-79 VSC Advanced Analytics - storage tiers

After you set the tiers, you can create service classes and define which tiers are used for the specific service class, as shown in Figure 12-80.

Figure 12-80 VSC Advanced Analytics - SQL Service Class

Apart from the tier selection, you can define whether this volume must be mirrored to an auxiliary storage system, whether to use thin provisioning, or to enforce redundant fabrics, as shown in Figure 12-81.

Name	Type	Used Space (%)	Total Capacity (GiB)	Available Space (GiB)	Unavailable Space (GiB)	Description
Bronze	Block		0.00	0.00	0.00	Standard storage for non-mission-critical applications.
Enhancedisolation	File		0.00	0.00	0.00	Enhanced isolated file storage.
Gold	Block	38%	1,186.00	738.00	0.00	Highest-performing storage for mission-critical applications.
Normalisolation	File		0.00	0.00	0.00	Normal isolated file storage.
Silver	Block	0%	12,044.00	12,044.00	0.00	High-performing storage for applications in production.
VersaStack_SQL	Block	0%	12,782.00	12,782.00	0.00	

Figure 12-81 VSC Advanced Analytics - Service Classes

As part of the service class creation, you can restrict the service class to specific capacity pools and define which users can provision from this service class. You can allow non-admin users to provision from these service classes without additional approval to run the provisioning request.

Figure 12-82 shows the VSC VersaStack and VersaStack_SSD capacity pools.

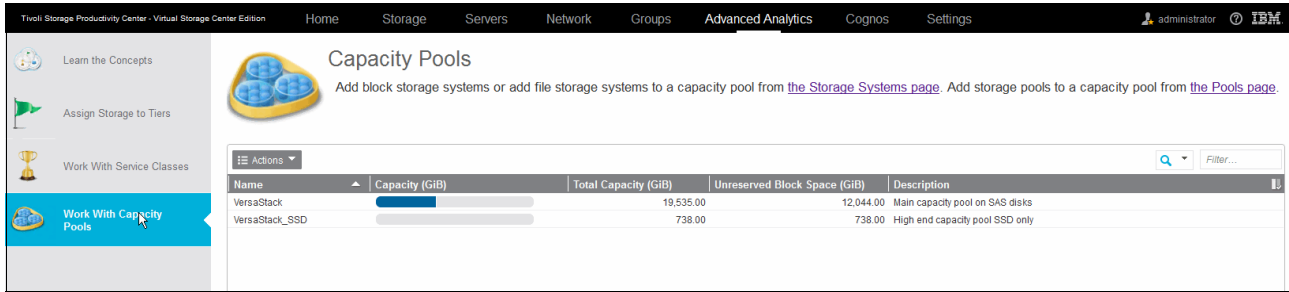


Figure 12-82 VSC Advanced Analytics - Capacity Pools

Similar to the Cisco UCS service profiles, you can use the VSC service classes to have a uniform deployment of your storage resources to the hypervisors or applications running on the hypervisors, such as this DB2 cluster in this environment.

In the VersaStack_SQL service class that we created, we assigned both the Tier1 (SSD) and the Tier2 (SAS) tiers. When creating volumes based on this service class, the VSC analyzes the load on the pools that are associated with these tiers by using the historically captured performance data for optimal volume placement at the creation of the volume.

12.8.2 Provisioning

With the service classes defined, you can now provision volumes. First, provision an additional data volume for the DB2 Servers from within the VSC Web GUI, and then provision a new datastore by using the vSphere Web Client VSC extensions.

Provisioning LUNs and volumes to the DB2 cluster

To provision LUNs and volumes to the SQL cluster, complete the following steps:

1. Start the VSC Web GUI and go to the Servers section. Select the SQLVM01 and SQLVM02 servers, right-click them, and select **Provision Storage**.

Figure 12-83 shows VSC Provisioning starting the Provision Storage wizard to provision additional LUNs to the SQL Servers.

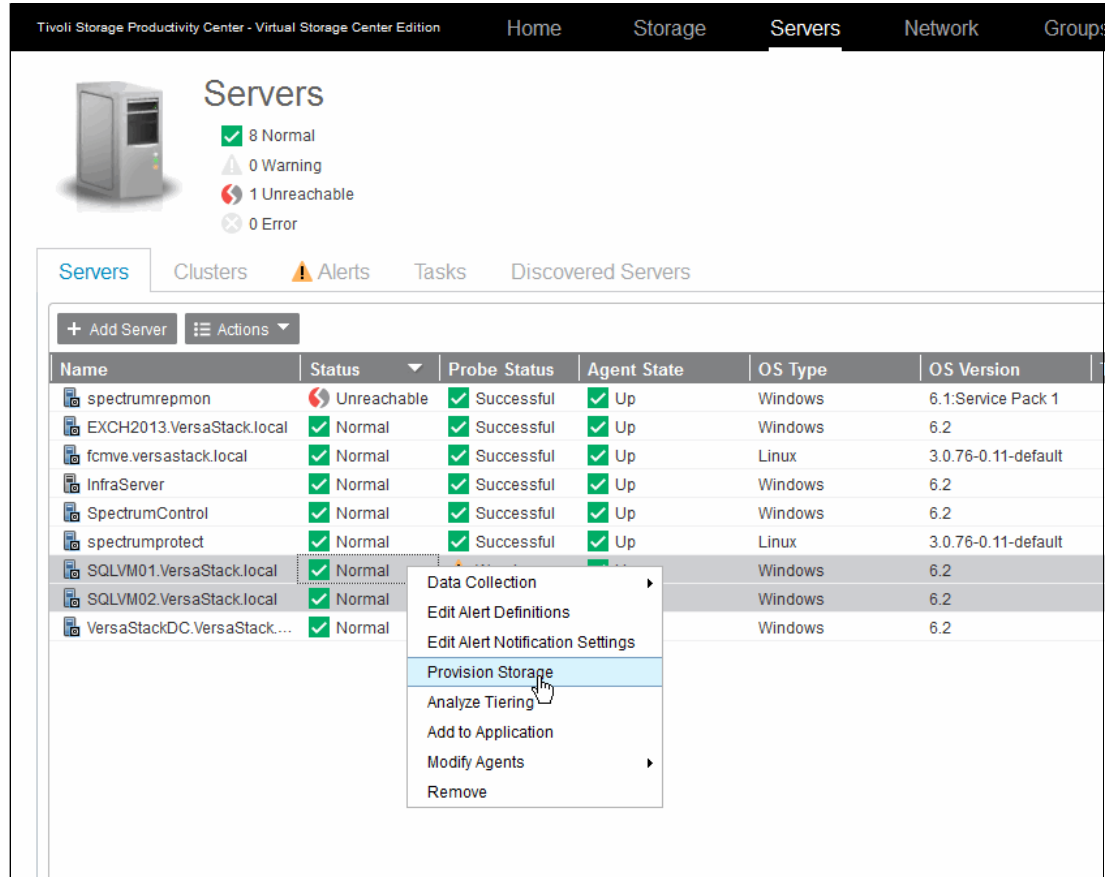


Figure 12-83 VSC Provisioning - storage

2. Choose between block or file volume provisioning. Click **Block** and click **Next**.

Figure 12-84 on page 329 shows the VSC Provisioning defining the required volumes and selecting the VersaStack_SQL service class.

Provision Storage

Define Volumes
Specify the characteristics of the volumes to create.

[BPCUI01581](#)
Volumes are assigned to the hypervisors that host virtual machines. Volumes are not assigned directly to virtual machines.

Name	Capacity	Service Class View	Capacity Pool View
1. sql_rdm_data_2	256 GIB	VersaStack_SQL	All available storage
2. sql_rdm_log_2	64 GIB	VersaStack_SQL	All available storage

+ Add More 1

Ticket:

[Back](#) [Next](#) [Cancel](#)

Figure 12-84 VSC Provisioning - Storage - Define Volumes

Note: The volumes that you want to use as raw device mapped volumes for the SQL Servers are created and assigned to the hypervisors and not to the SQL Servers directly. You must create the Raw Device Mapping (RDM)-based disk later.

When you click **Next**, the built-in analytics engine determines the ideal storage pool placement for the volume based on the historical data that it captured from the storage environment.

Figure 12-85 shows the VSC Provisioning analytics engine calculating the ideal storage pool in which to place the volume.

Provision Storage

Determining storage placement recommendation. This process might take several minutes.

Figure 12-85 VSC Provisioning - storage placement

After the placement recommendation is defined, an overview window opens and shows the provisioning task with the option to either run it immediately or to schedule it for processing during your regular maintenance window.

Figure 12-86 shows the VSC Provisioning overview of the tasks to be run.

Provision Storage

Service Class: [VersaStack_SQL](#) User: administrator

Storage

Status	Action	Volume Name	Pool	Storage System	Capacity (GiB)
Not running	Create	sql_rdm_data_2	mdiskgrp0	Storwize V7000 Ver...	256.00
Not running	Create	sql_rdm_log_2	mdiskgrp0	Storwize V7000 Ver...	64.00

Host Connections

Status	Action	Host Name	Host Type	Port	Volume
Not running	Assign	vm-host-infra-02.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_lo...
Not running	Assign	vm-host-infra-01.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_da...
Not running	Assign	vm-host-infra-01.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_lo...
Not running	Assign	vm-host-infra-02.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_da...

Buttons: Back, Execute, Schedule, Delete, Close

Figure 12-86 VSC Provision Storage Task Summary

3. Click **Execute** to start the job immediately and monitor the progress.

Figure 12-87 shows that all jobs completed.

Provisioning

VersaStack_SQL_sql_rdm_data_2-201 Rename Started: Jul 13, 2015 09:31:19 PDT
Average Duration: N/A [Open Logs](#)

Service Class: [VersaStack_SQL](#) User: administrator

Storage

Status	Action	Volume Name	Pool	Storage System	Capacity (GiB)
Successful	Create	sql_rdm_data_2	mdiskgrp0	Storwize V7000 Ver...	256.00
Successful	Create	sql_rdm_log_2	mdiskgrp0	Storwize V7000 Ver...	64.00

Host Connections

Status	Action	Host Name	Host Type	Port	Volume
Successful	Assign	vm-host-infra-02.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_lo...
Successful	Assign	vm-host-infra-01.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_da...
Successful	Assign	vm-host-infra-01.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_lo...
Successful	Assign	vm-host-infra-02.versastack.lo...	VMware ESX	20000025B501...	sql_rdm_da...

Buttons: Need Help, Execute, Schedule, Delete, Close

Figure 12-87 VSC Provisioning - job results

4. To create the RDM mapping to the SQLVM01 and SQLVM02, complete the following steps:
 - a. Log in to the VMware vCenter.
 - b. Go to **Hosts and Clusters**.
 - c. Select vm-host-infra-01, select **Configuration**, and click **Storage Adapters**.
 - d. Click **Rescan All**.

- e. Select vm-host-infra-02, select **Configuration**, and click **Storage Adapters**.
- f. Click **Rescan All**. Two new devices (256 GB and 64 GB) show up in the devices list.
- g. Go to **VMs and Templates**.
- h. Select the SQLVM01 and **Edit Settings**.
- i. Click **Add** → **Hard Disk** → **Raw Device Mappings**, and click **Next**.

Figure 12-88 shows VSC Provisioning adding the new LUNs as RDMs to the SQL VMs.

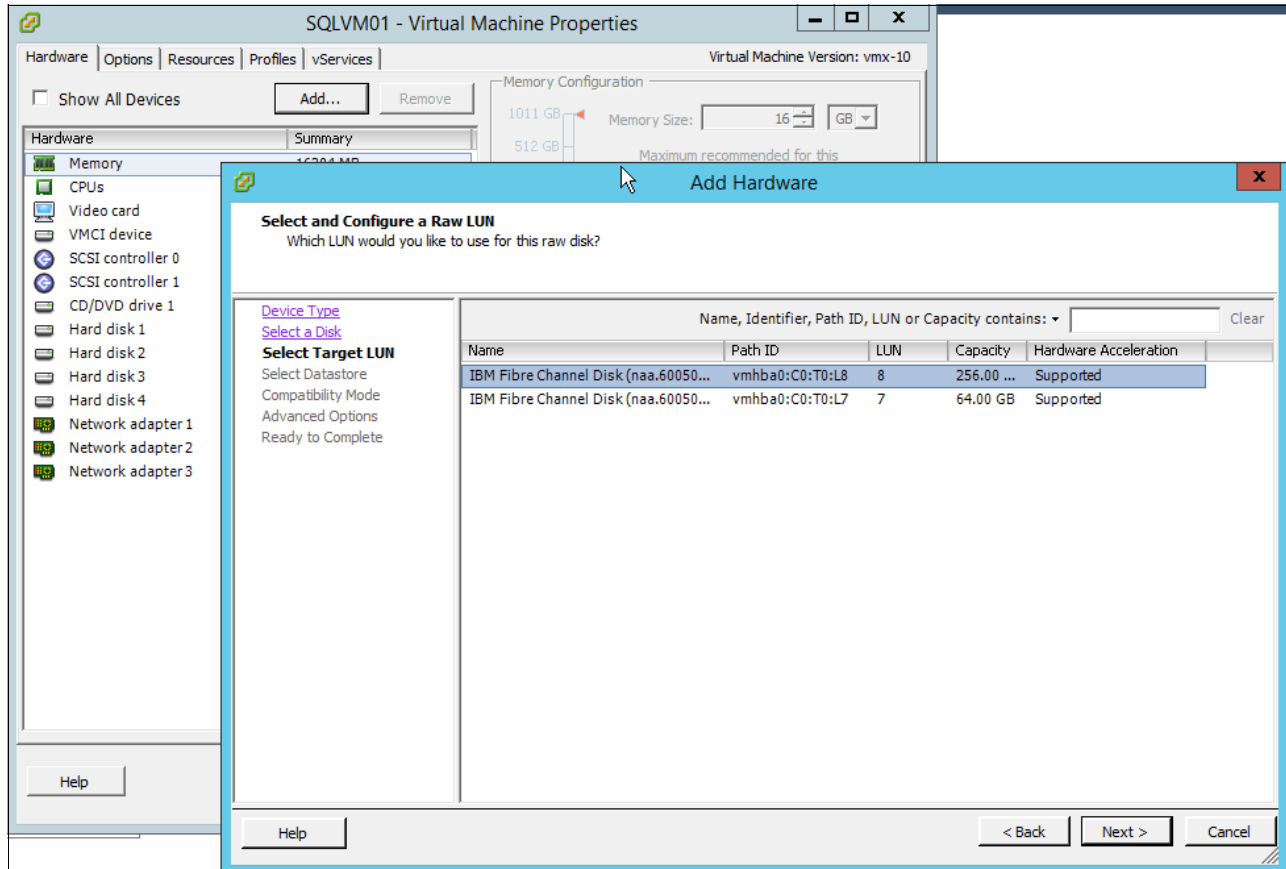


Figure 12-88 VSC Provisioning - add RDM

- j. Select the new data LUN and click **Next**.
- k. Store the LUN mapping with the VM and click **Next**.
- l. Choose **Physical compatibility** for the LUN mapping and click **Next**.
- m. Choose the virtual device node and click **Next** and then **Finish**.

Repeat steps 4a on page 330 to 4m for the log LUN on SQLVM01 and for both the data and log LUN on the SQLVM02.

Provisioning a new datastore to the hypervisor

You can provision a new volume to the hypervisors from within the VSC Web GUI and have a datastore assigned to it. The same action can be performed from within the vSphere Web Client.

This function gives VMware administrators the flexibility to foresee their own provisioning needs. Storage allocation is controlled by the use of service classes. The service class defines the target tier, the capacity pool, and whether the provisioning action can be carried out immediately or must be approved by the storage admin from within the VSC Web GUI first.

Figure 12-89 shows the VSC vSphere Web Client extension connection status.

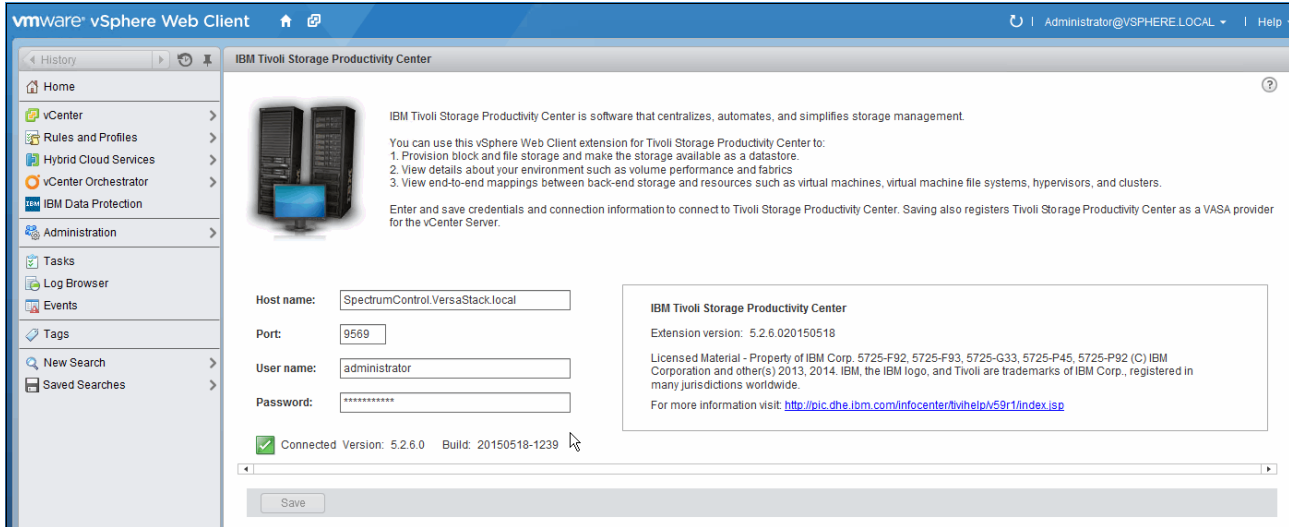


Figure 12-89 VSC vSphere - Web Client Extension

To provision a new datastore to the hypervisors, complete the following steps:

1. Log in to the vSphere Web Client.
2. Go to the **Hosts and Clusters**. Right-click **vm-host-infra-01** and select **Provision Block Storage** from the **All TPC Actions** menu.

Figure 12-90 shows the VSC vSphere Web Client provisioning of block storage.

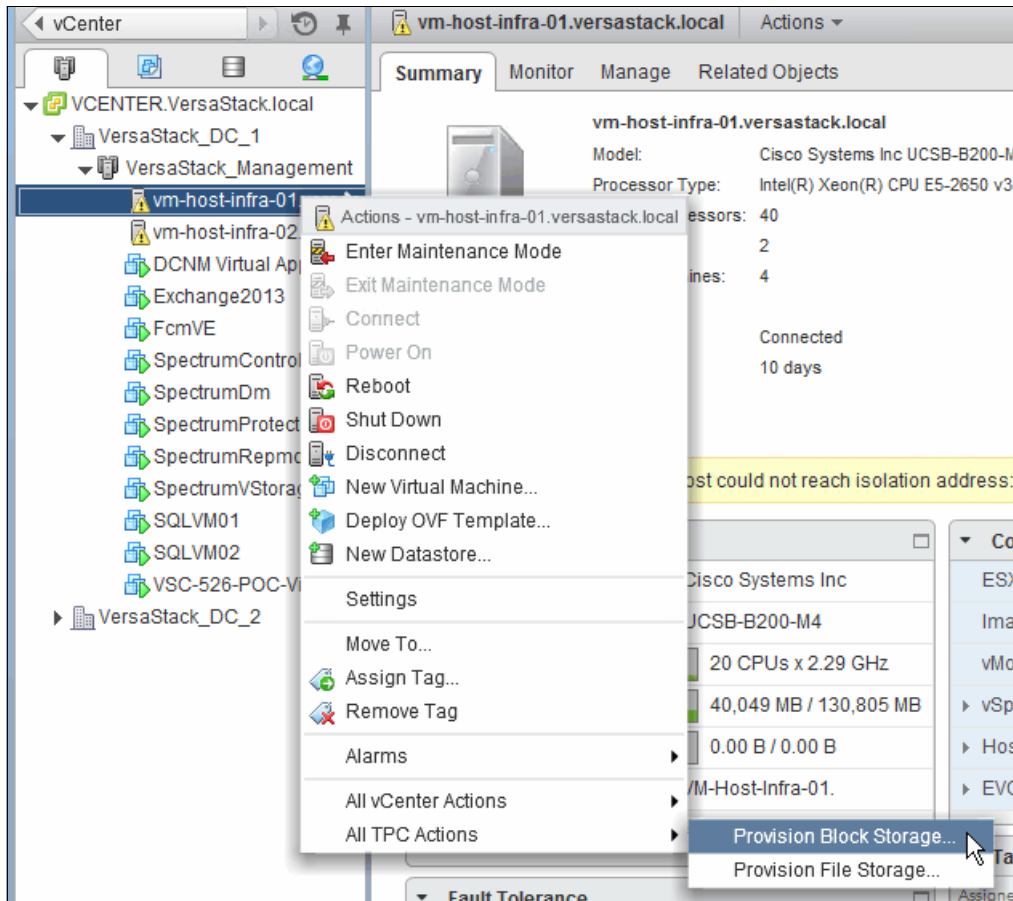


Figure 12-90 VSC vSphere - Provision Block Storage

3. Set the size to 1 TiB, choose the **Silver** service class, select the **Create datastore** check box, name the data `Infra_Datastore_3`, and click **OK** to start the provisioning.

Figure 12-91 shows the VSC vSphere Web Client defining of the block storage provisioning parameters.



Figure 12-91 VSC vSphere Block Storage Definition

The vSphere Web Client extension communicates with the VSC and runs a provisioning task that is similar to the task that was created for the SQL RAW device mapping provisioning.

Figure 12-92 shows the task results.

Provisioning

Silver_V_150713_094238-20150713-09 Rename Completed: Jul 13, 2015 09:41:41 PDT
Duration: 1 minute [Open Logs](#)

Service Class: **Silver** User: administrator

Storage

Status	Action	Volume Name	Pool	Storage System	Capacity (GiB)
✓ Successful	Create	V_150713_094238	mdiskorp0	Storwize V7000 Ver...	1,024.00

Host Connections

Status	Action	Host Name	Host Type	Port	Volume
✓ Successful	Assign	vm-host-infra-01.versastack.io...	VMware ESX	20000025B501...	V_150713 ...
✓ Successful	Assign	vm-host-infra-02.versastack.io...	VMware ESX	20000025B501...	V_150713 ...

Multipath

Status	Action	Host Name	Multipathing Policy
i Informational	...	vm-host-infra-01.versastack.io...	Default
i Informational	...	vm-host-infra-02.versastack.io...	Default

[Need Help](#) Execute Schedule Delete Close

Figure 12-92 VSC vSphere - Provisioning Task

With the vSphere Web Client itself, you can monitor the progress through the Tasks and Events subtabs from the vm-host-infra-01 monitor tab.

Figure 12-93 shows the results of datastore provisioning.

Description	Type	Date Time	Task	Target	User
The creation of the VMFS completed.	Information	7/13/2015 9:43 AM		vm-host-infra-01...	administrator
Created VMFS datastore Infra_Datastore_3	Information	7/13/2015 9:43 AM		vm-host-infra-01...	
Discovered datastore Infra_Datastore_3	Information	7/13/2015 9:43 AM		vm-host-infra-01...	
File system [Infra_Datastore_3, 55a3ea5a-6b76...	Information	7/13/2015 9:43 AM		vm-host-infra-01...	
Task: Create VMFS datastore	Information	7/13/2015 9:43 AM	Create VMFS datastore	vm-host-infra-01...	VSPHERE.LOCAL\Administrator
The rescan of the HBA completed.	Information	7/13/2015 9:43 AM		vm-host-infra-01...	administrator
Task: Rescan HBA	Information	7/13/2015 9:43 AM	Rescan HBA	vm-host-infra-01...	VSPHERE.LOCAL\Administrator
Task: Rescan HBA	Information	7/13/2015 9:43 AM	Rescan HBA	vm-host-infra-01...	VSPHERE.LOCAL\Administrator
Task: Provision storage for LUN	Information	7/13/2015 9:42 AM	Provision storage for LUN	vm-host-infra-01...	administrator

Figure 12-93 VSC vSphere - Provisioning Events

Optimization

VSC facilitates uniform business-aligned storage allocation. As outlined in 12.4.2, “Self-optimizing” on page 266, VSC can also analyze your storage environment either at scheduled intervals or when triggered by performance monitor events to perform storage optimization. Figure 12-94 shows the tiered storage optimization.

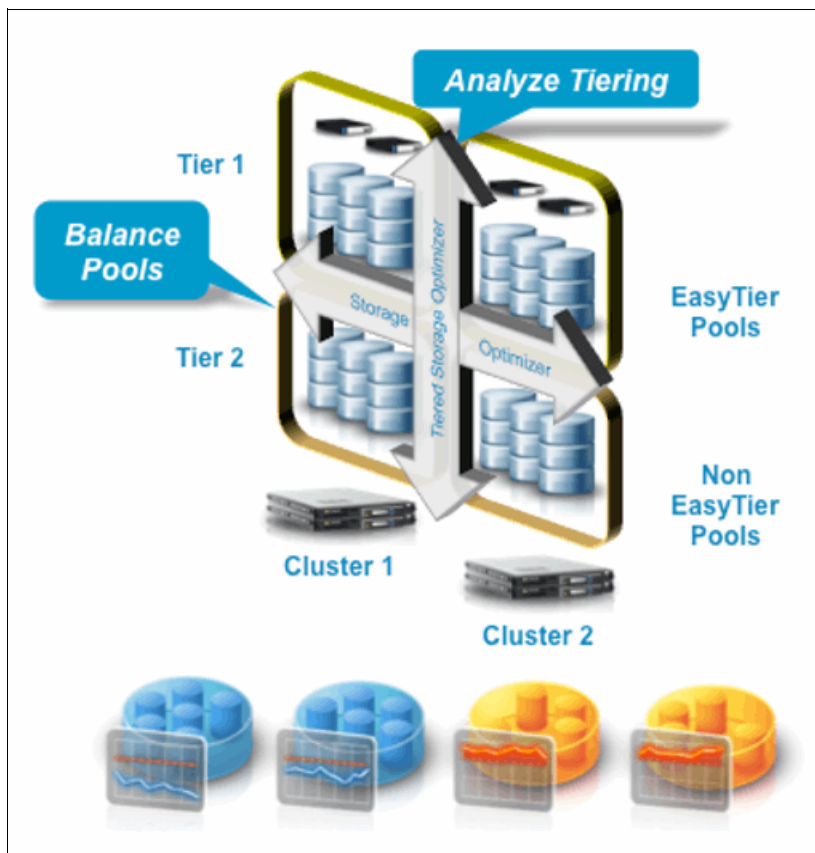


Figure 12-94 VSC tiered storage optimization

Storage optimization can consist of retiering a volume across tiers or balancing volumes within storage pools. This section outlines the logical flow for both actions.

Retiering optimizes storage performance by moving volumes to different storage tiers. You can choose the set of source volumes that you want to analyze for retiering. In this example, the volumes in a tier 2 pool are analyzed to see whether they require retiering to a set of three tier 1 pools, as shown in Figure 12-95.




Figure 12-95 VSC optimization

One of the volumes in the tier 2 pool is overutilized. If the overutilized volume is moved to a tier 1 pool with sufficient performance capacity, then the performance of the volume can be improved (Figure 12-96).




Figure 12-96 Volume should be moved to tier 1

The performance of the target pools on tier 1 is analyzed and recommendations are generated. The recommendations involve up-tiering the overutilized volume from the tier 2 pool to the tier 1 pool. You can review the recommendations and automatically move the volume to the tier 1 pool (Figure 12-97).




Figure 12-97 Volume moved to tier 1

You can also down-tier volumes. In this example, one of the volumes in the tier 1 pool is underutilized, which means that the volume is occupying more expensive storage than is necessary (Figure 12-98).




Figure 12-98 Volume that is identified to be down-tiered

By analyzing the performance of the tier 2 pool, a recommendation is generated to down-tier the volume (Figure 12-99).




Figure 12-99 Volume down-tiered

A single tiering analysis can result in multiple volume movements in which volumes are moved to both lower and higher tiers of storage. You can schedule an analysis task to run at specified intervals for a selected set of source volumes and target pools so that you can regularly monitor opportunities to retier volumes (Figure 12-100).




Figure 12-100 Multiple volume movements

Another form of optimization is balancing. An environment can contain pools with low and high activity levels. To identify pools that have high activity levels, look at the values that are shown in the Activity Deviation (%) column. The value in the Activity Deviation (%) column shows the difference between the activity level of the pool and other pools on the same tier and storage system. Pools with values greater than 10% are candidates for balancing (Figure 12-101).




Figure 12-101 Balancing

Tivoli Storage Productivity Center can analyze pools on the same tier and identify opportunities to move volumes such that the activity deviation percentage of the pool falls below 10% (Figure 12-102).




Figure 12-102 Deviation

The volumes with the most I/O activity in our VersaStack environment are the DB2 clustered volumes and the datastore hosting the IBM Spectrum Protect server. Within the VersaStack capacity pool that we defined, we have two tiers of storage available: Tier 1 using SSDs and Tier 2 using SAS for back-end storage.

Use the Analyze Tiering function of the VSC to evaluate whether the DB2 volumes require up-tiering to the SSD-based Tier 1 by completing the following steps:

1. Log on to the VSC Web GUI.
2. Click **Advanced Analytics** → **Optimization** → **Optimize Volumes**.

3. Select the `sql_rdm_data` and `sql_rdm_log` volumes, right-click, and select **Analyze Tiering**, as shown in Figure 12-103.

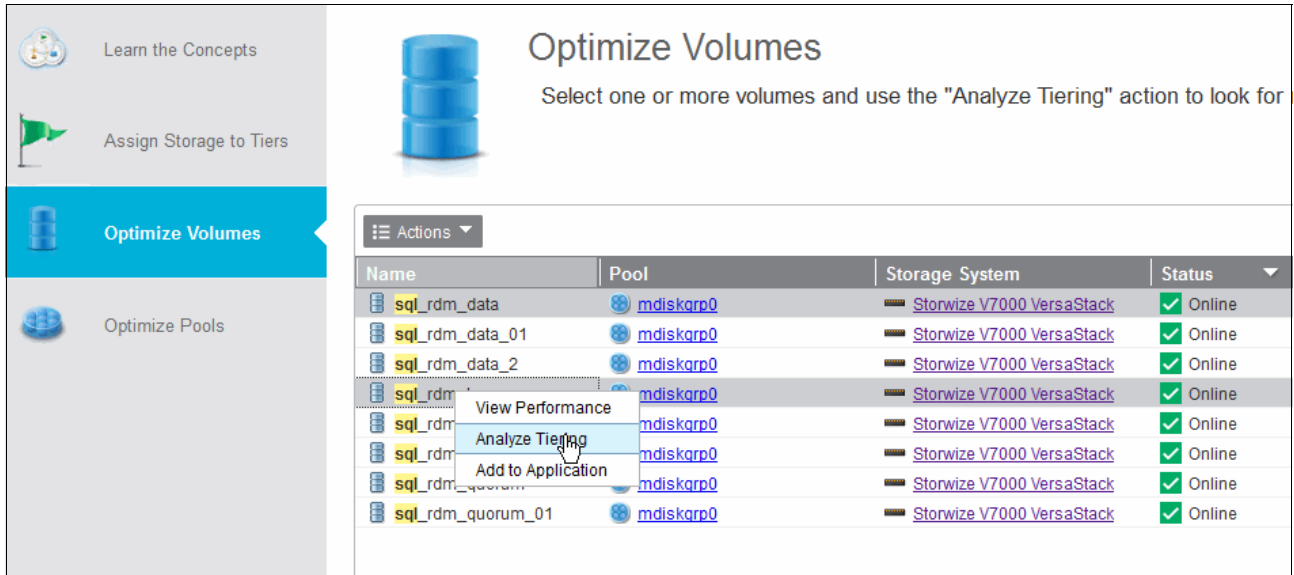


Figure 12-103 VSC Optimize Volumes - Analyze Tiering

4. Select the VersaStack capacity pool and click **Next** (Figure 12-104).

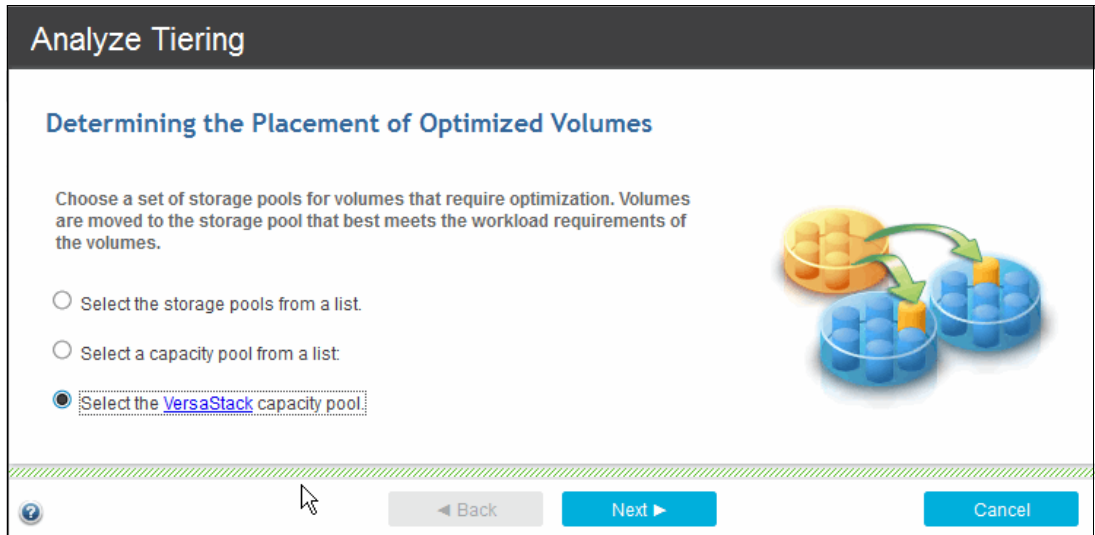


Figure 12-104 VSC Optimize Volumes - Select Capacity Pool

By default, VSC uses the data from the last seven days. You can either set the Volume I/O rate (I/O per second) or Volume I/O Density (I/O per second per GB). You can also define the maximum pool I/O rate for the available tiers to ensure that adding the volume does not cause the total amount of I/O for that pool to be exceeded (Figure 12-105).

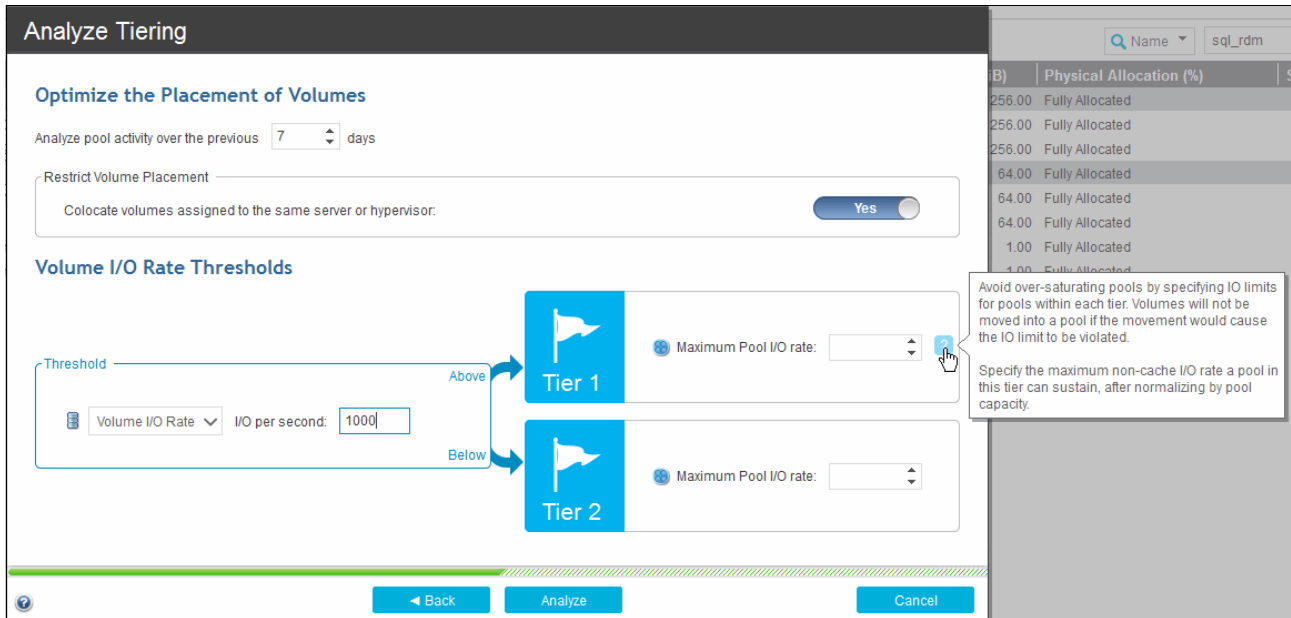


Figure 12-105 VSC Optimize Volumes - Define Thresholds

5. Click **Analyze** to start the process, as shown in Figure 12-106.

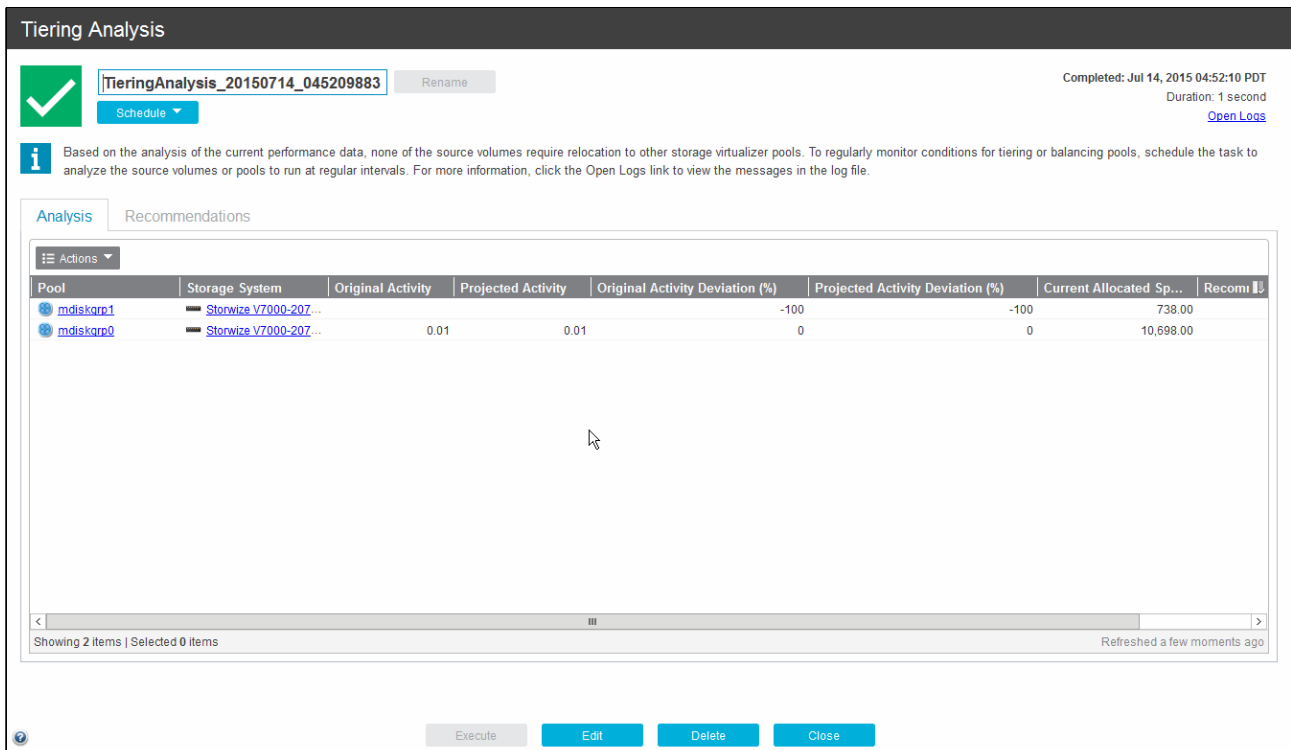


Figure 12-106 Tiering Analysis results

The load over the past seven days on the DB2 cluster in our VersaStack setup does not require the DB2 clustered volumes to be moved into tier 1. We can now schedule this analysis to take place, for example, each week.

The results of the Analyze Tiering is grouped in the Tasks section of the VSC Web GUI. Recommendations for up- or down-tiering can be run after review by the storage administrator or set up for automated running because the tiering migration is transparent and has no impact on the host system.

12.8.3 Integrating servers and virtual machines

In “Integrating the VMware vCenter Hypervisor with Spectrum Control” on page 302, we added the VMware vCenter environment to the Tivoli Productivity Center VSC. As a result, all VMs running on the hypervisor are discovered at the time the scheduled resource probing of these systems takes place.

Within VSC, both physical and virtual systems are grouped in the Servers section.

Figure 12-107 shows the VSC Servers overview filtered by the DB2 VMs.

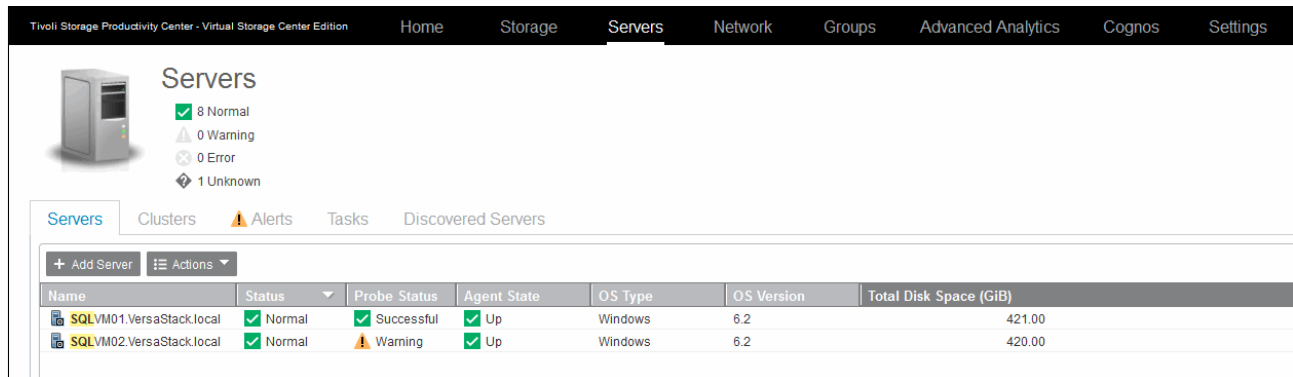


Figure 12-107 VSC Servers overview

Servers can be both physical or virtual systems. VMs are added under the Discovered Servers tab.

On these servers, you can install Storage Resource agents to collect information about storage resources, such as servers, VMs, workstations, HBAs, and fabrics.

You must deploy Storage Resource agents on resources where you want to gather the following information:

- ▶ Asset information
- ▶ File and file system attributes
- ▶ Database application information
- ▶ Network-attached storage (NAS) device information
- ▶ Topology information
- ▶ Information about zoning and the fabrics that are visible to the server

You can also monitor servers without deploying a Storage Resource agent. When you add an agentless server, VSC correlates information about that server with the known host connections on monitored resources. If a match is made between the server and a monitored resource, you can view topology information and the capacity and performance of the storage that is assigned to that server.

We deployed a Storage Resource Agent on both members of the DB2 clusters by completing the following steps:

1. Open the VSC Web GUI.
2. Click **Servers** and click **Add Server**.

Figure 12-108 shows adding a server and deploying a Storage Resource Agent.

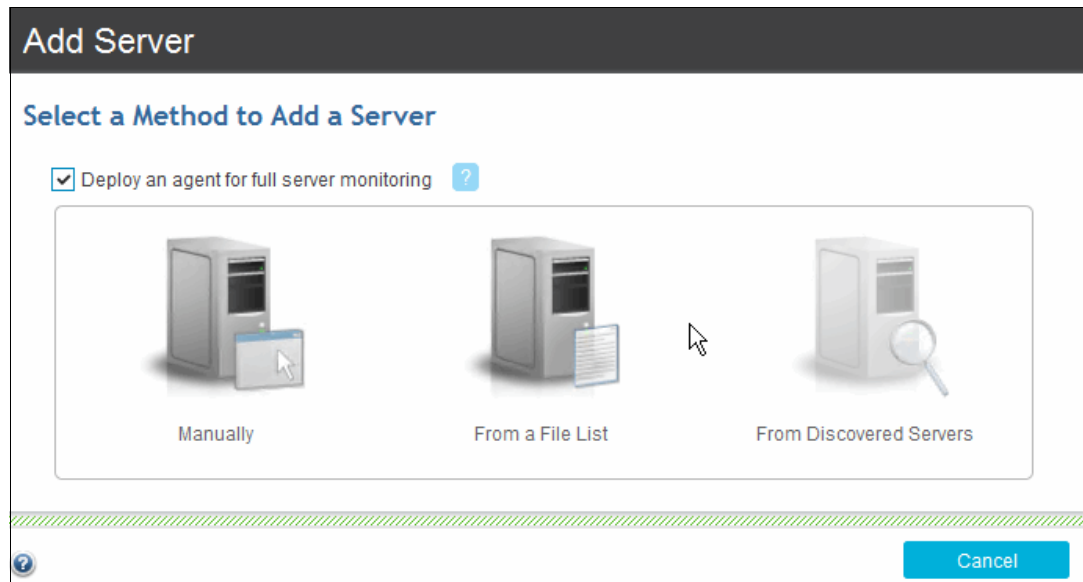
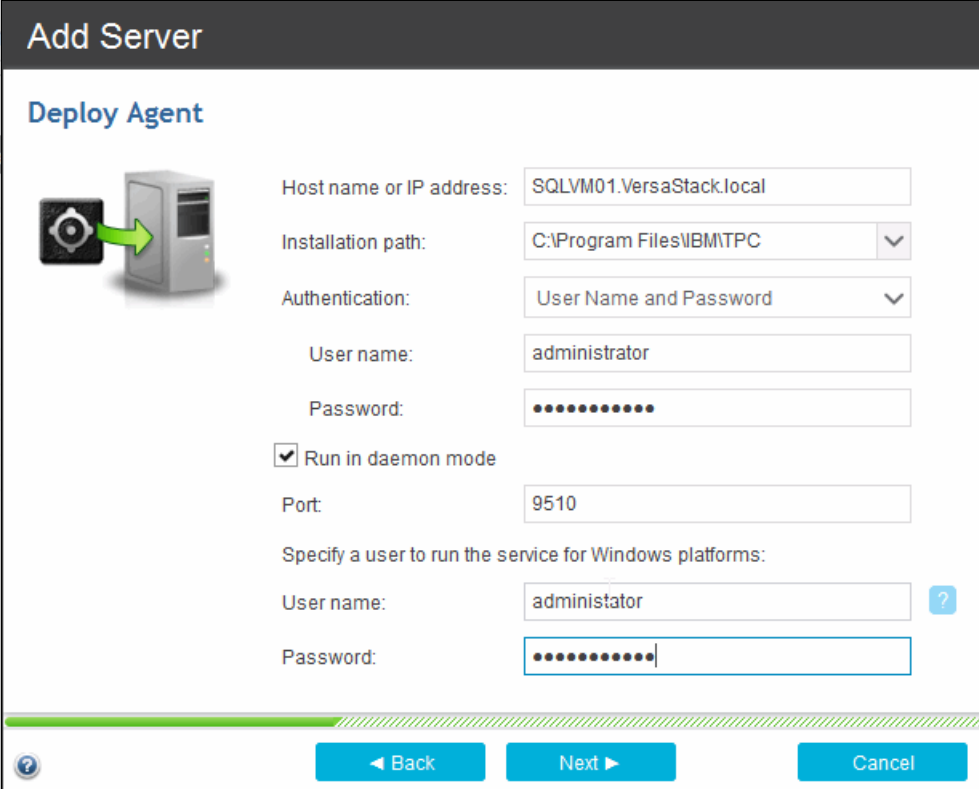


Figure 12-108 VSC Add Server

3. Select the **Deploy an agent for full server monitoring** and click **Manually**. You can also deploy multiple storage resource agents at a time by creating a file list containing the required information (Figure 12-109).



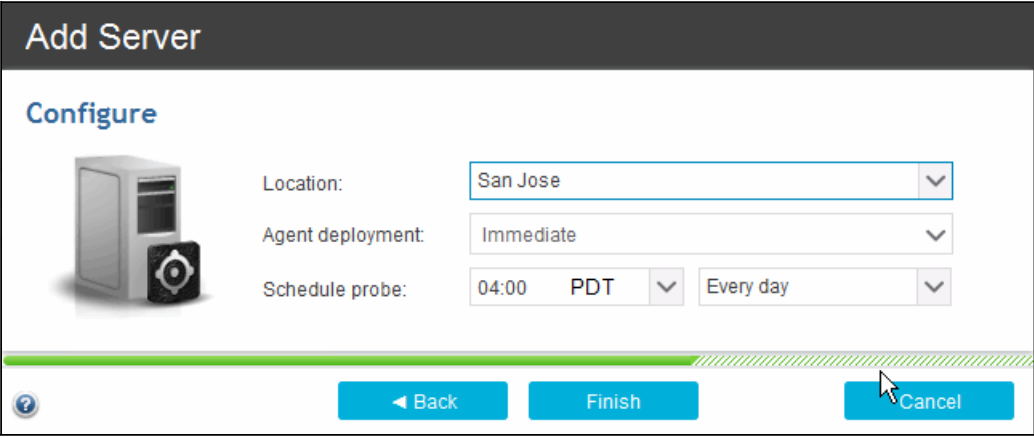
The screenshot shows the 'Add Server' wizard in the 'Deploy Agent' step. On the left, there is an icon of a server with a green arrow pointing to it. The main area contains several input fields and a checkbox:

- Host name or IP address: SQLVM01.VersaStack.local
- Installation path: C:\Program Files\IBM\TPC
- Authentication: User Name and Password
- User name: administrator
- Password: [masked]
- Run in daemon mode
- Port: 9510
- Specify a user to run the service for Windows platforms:
- User name: administrator
- Password: [masked]

At the bottom, there are three buttons: 'Back', 'Next', and 'Cancel'. A progress bar is visible at the bottom of the form area.

Figure 12-109 VSC Add Server - Deploy Agent

4. Provide the required connectivity and login credentials and click **Next** (Figure 12-110).



The screenshot shows the 'Add Server' wizard in the 'Configure' step. On the left, there is an icon of a server. The main area contains several dropdown menus and a checkbox:

- Location: San Jose
- Agent deployment: Immediate
- Schedule probe: 04:00 PDT, Every day

At the bottom, there are three buttons: 'Back', 'Finish', and 'Cancel'. A progress bar is visible at the bottom of the form area.

Figure 12-110 VSC Add Server - Schedule a probe

Similar to other resources, such as the Storwize V7000 storage systems, a schedule is created to probe the newly added server at specified intervals (Figure 12-111).

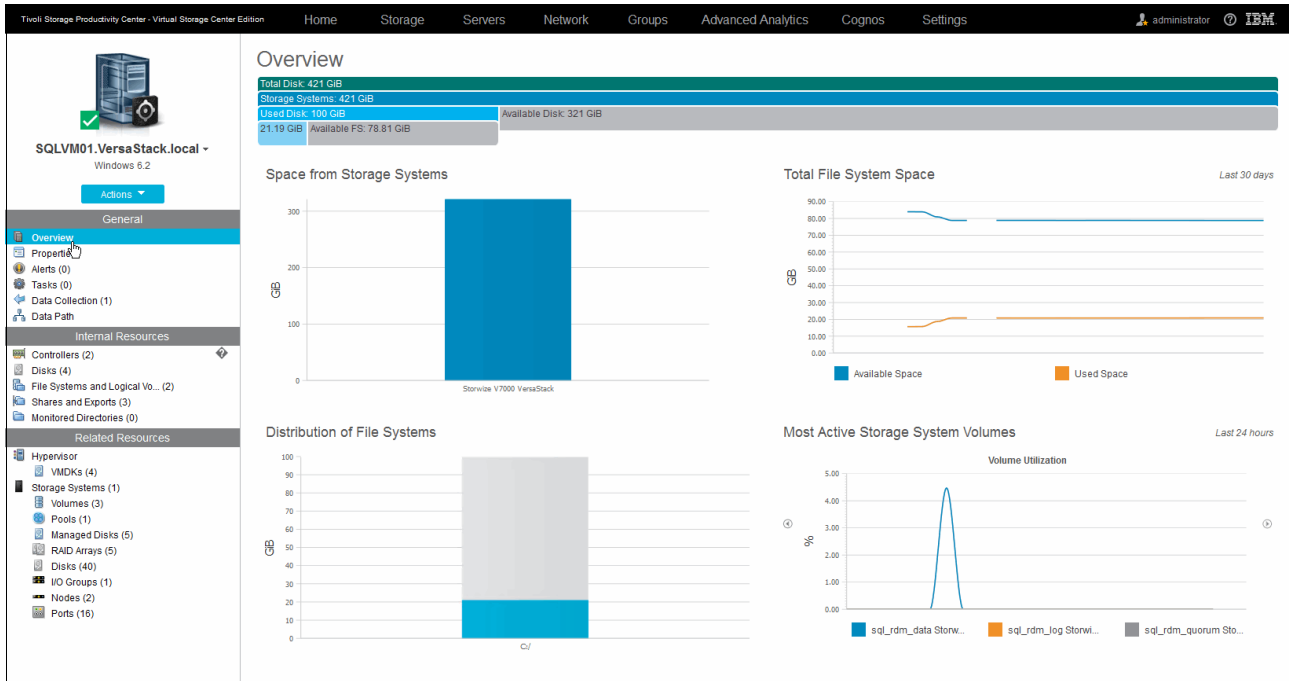


Figure 12-111 VSC Server - detailed overview

Within VSC, we take a uniform approach to group the information for the selected resource in three sections, as you can see in the left pane of Figure 12-111:

- ▶ General:
 - Overview: Gives a graphical representation of the following items:
 - Space from Storage Systems
 - Most Active Switch Ports
 - Most Active Storage Systems Volumes
 - Total File System Space
 - Distribution of File Systems
 - Properties: Displays General, Hardware, Storage, and Agent Related Information.

Figure 12-112 shows the VSC Server General properties.

The screenshot displays the Tivoli Storage Productivity Center - Virtual Storage Center Edition interface. The top navigation bar includes Home, Storage, Servers, Network, Groups, and Advanced Analytics. The main content area is titled 'Properties' and features four tabs: General, Hardware, Storage, and Agent. The 'General' tab is active, showing a table of server properties. The left sidebar contains a navigation menu with sections for Overview, Properties, Alerts, Tasks, Data Collection, Data Path, Internal Resources, and Related Resources.

Properties	
Name	SQLVM01.VersaStack.local
Status	Normal
Acknowledged	No
OS Type	Windows
OS Version	6.2
IP Address	192.168.10.51
Domain Name	versastack.local
Cluster	—
Virtual Machine	Yes
Hypervisor	vm-host-infra-01.versastack.local
Last Start Time	Jul 10, 2015 01:09:16 PDT
Probe Status	Successful
Probe Schedule	Daily. Next run at Jul 15, 2015 12:30:00 PDT
Time Zone	GMT-07:00
Data Source Count	1
Location	San Jose
Custom Tag 1	—
Custom Tag 2	—
Custom Tag 3	—

Figure 12-112 VSC Server Properties

- ▶ Alerts: Shows the alerts that are filtered for the selected server. Within the Definitions tab, you can activate or customize the following alerts:
 - Server Status Change Offline
 - Server Status Change Online
 - HBA Driver Version Changed
 - HBA Firmware Version Changed
 - Probe Failed
 - File System Discovered?
 - File System Low on Free Space?
 - File System Missing
 - Disk Discovered
 - Disk Defect Discovered
 - Disk Failure Predicted
 - Disk Missing
 - Grown Disk Defects Threshold Exceeded

Figure 12-113 shows these alert definitions.

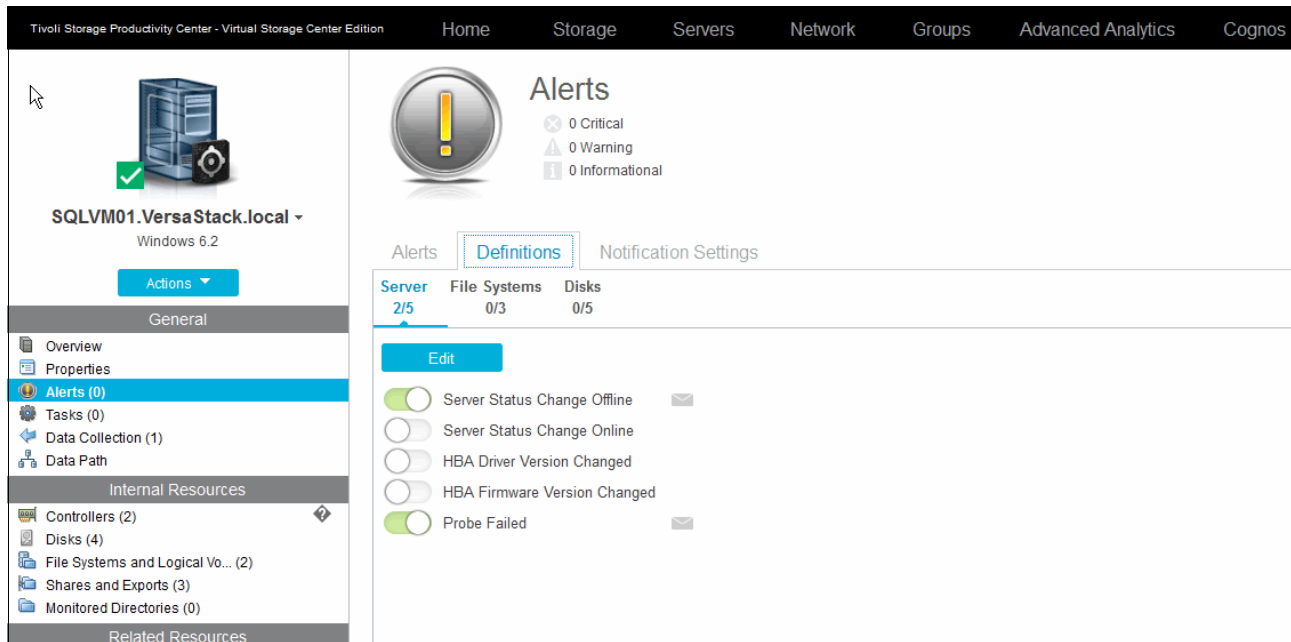


Figure 12-113 VSC Server - alert definitions

For each of these alerts, you can override the default notifications settings and have a script that is run through the Storage Resource Agent to take corrective actions automatically.

- ▶ **Tasks:** Groups tasks, such as provisioning jobs, for the specific server.
- ▶ **Data Collection:** Shows the status result of the latest probe and allows you to perform an Agent Upgrade (Figure 12-114).

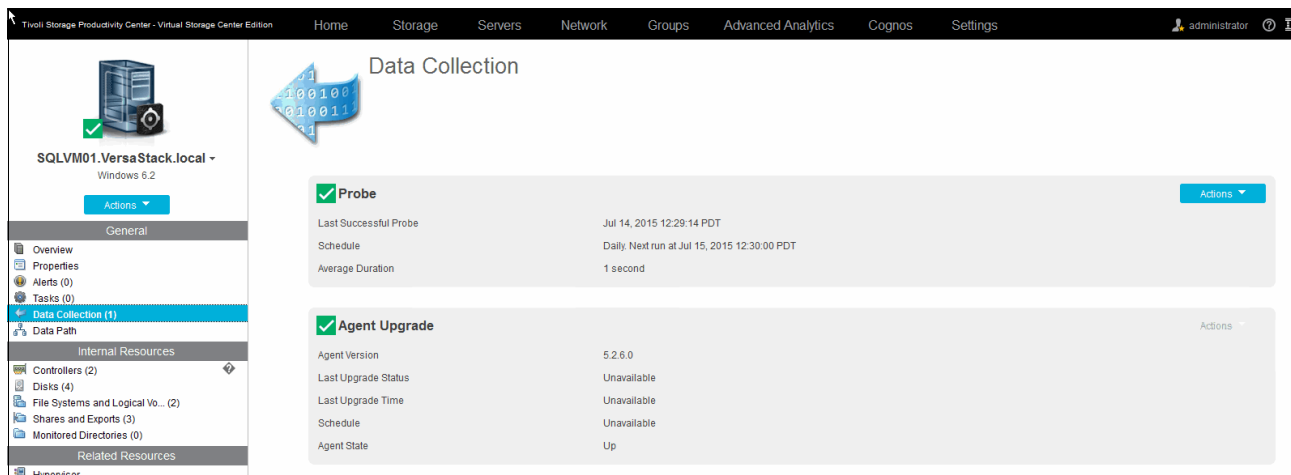


Figure 12-114 VSC Server - Data Collection

- ▶ **Data Path:** Shows the SAN data path.

- ▶ Internal Resources:
 - Controllers: Lists the internal controllers of the server, including information such as type, driver version, firmware, ROM version, HBA WWN, serial number, bus address, bus number, and associated disks.
 - Disks
 - File Systems and Logical Volumes
 - Shares and Exports
 - Monitored Directories
- ▶ Related Resources
 - Hypervisor: Lists the VMDKs that are associated with the server.

Figure 12-115 shows the VSC Server view of the DB2 member and its associated VMDKs.

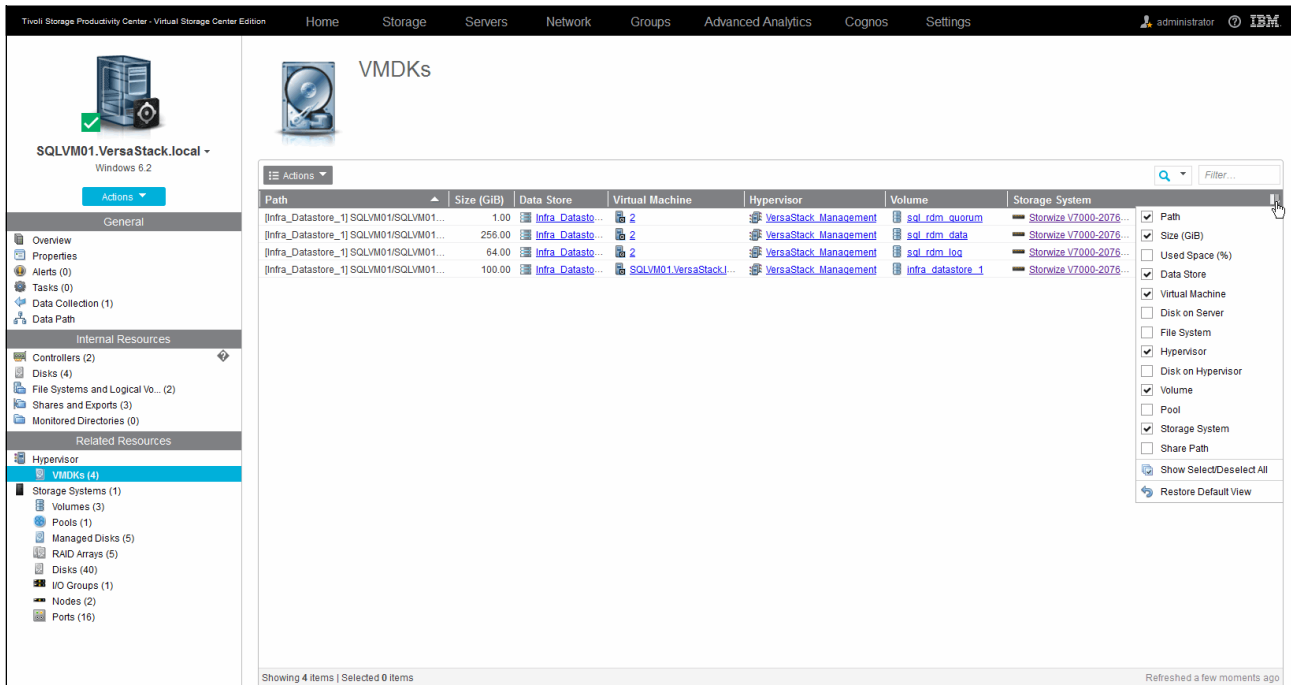


Figure 12-115 VSC Server VMDKs

- Storage Systems: Groups the storage resources that are related to the server.

By deploying Storage Resource Agents on to the DB2 cluster member servers, VSC also can identify the MSCS cluster itself.

Figure 12-116 shows the Servers displaying the cluster that is deployed in the VersaStack environment.

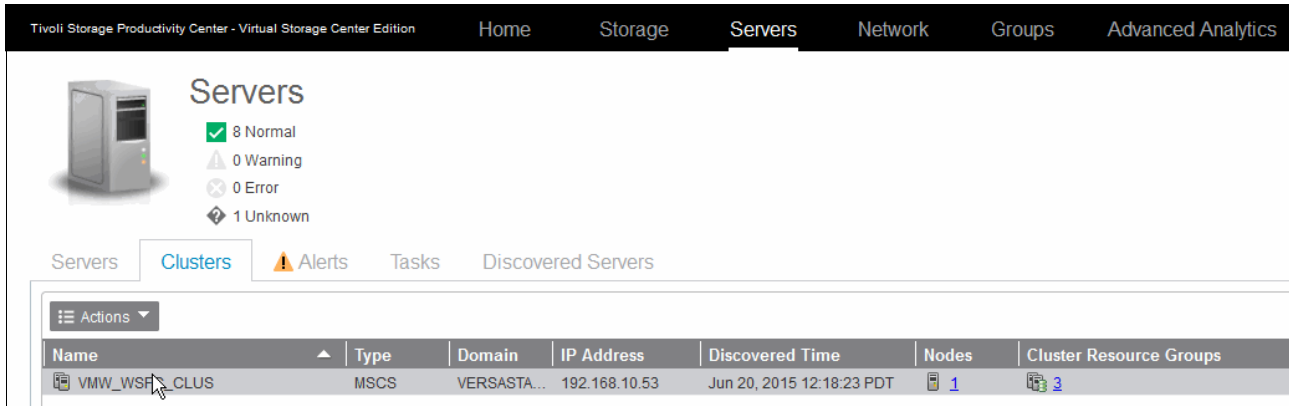


Figure 12-116 VSC Server - Clusters

Within the cluster properties notebook are the following tabs:

- ▶ General: Cluster type, domain, IP address, and discovery time
- ▶ Nodes: Member nodes of the cluster
- ▶ Cluster Resource Groups: Resource groups that are defined on the cluster

Figure 12-117 shows the VSC Cluster Properties notebook showing the Cluster Resource Groups for the MSCS cluster.

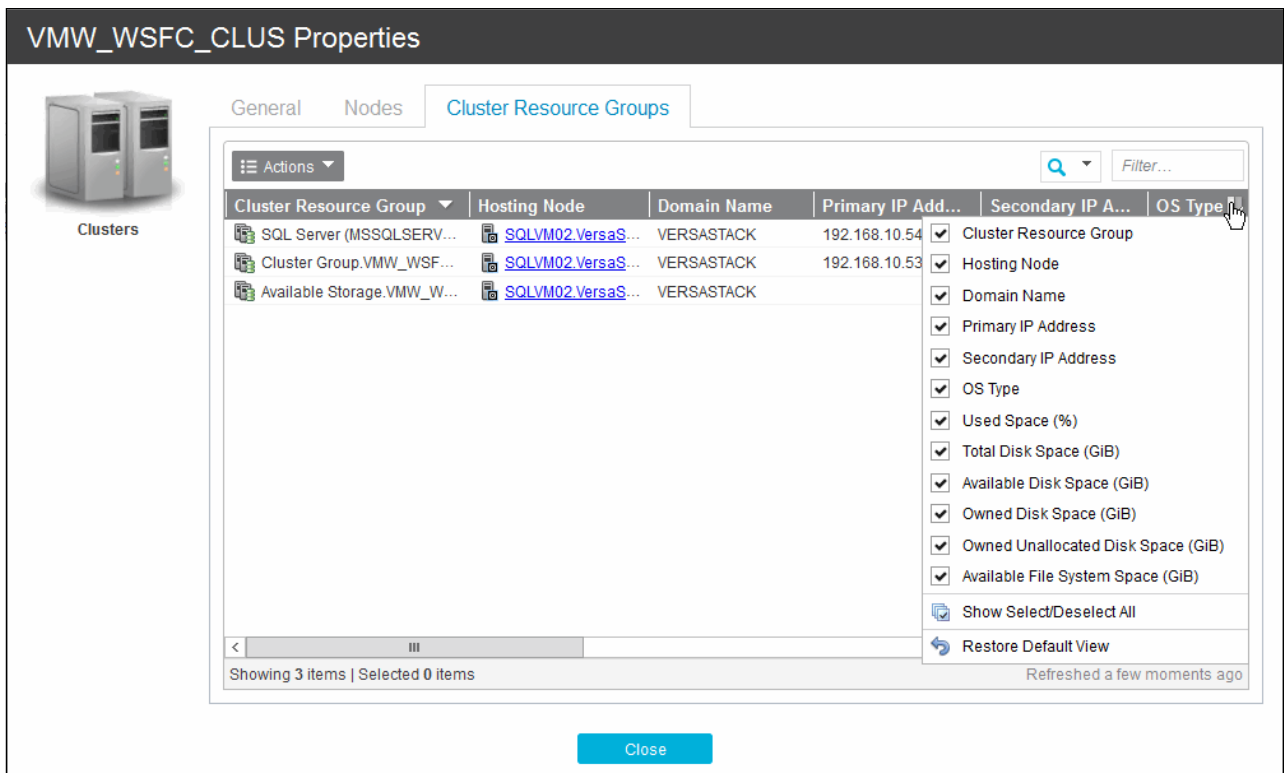


Figure 12-117 VSC Cluster Resource Group

Provisioning

In the initial configuration of the MSCS cluster, we used the Storwize V7000 GUI to provision the LUNs to the cluster member servers. This step can also be performed through the VSC either at the cluster level by provisioning directly to the cluster or by selecting individual servers to which to provision. Section “Provisioning LUNs and volumes to the DB2 cluster” on page 328 outlines how to provision additional data and log volumes to the cluster.

Monitoring

VSC correlates the registered servers (agentless or servers with a SRA deployed) with the Storage Systems and Fabric resources that are used by these systems.

Figure 12-118 shows a one-month performance overview of the cluster volumes.

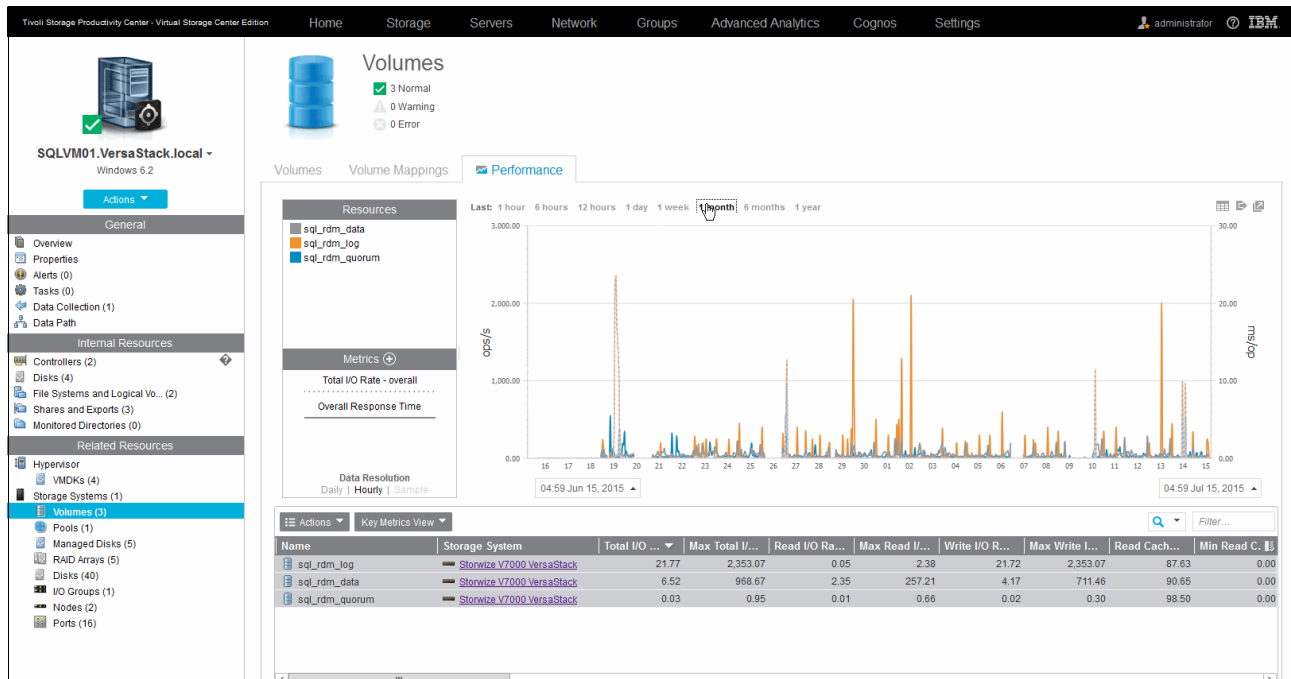


Figure 12-118 VSC Server Volume Performance

For more information about the performance monitoring and alerting capabilities of the Storwize V7000 storage system within the VSC environment, see 12.7.3, “Monitoring and alerting” on page 290.

Protection

For more information about protection, see Chapter 13, “IBM Spectrum Protect integration” on page 355.

12.8.4 Reporting for departments and applications

Environments that offer high levels of automation and flexibility, such as the VersaStack solution with the VMware hypervisor, often pose challenges for the storage administration teams in areas such as troubleshooting, accountability, and chargeback, especially when data is moved dynamically from less expensive SAS to more expensive SSD tiers by using the VSC built-in auto-tiering function.

To accommodate for these issues, you can use the concept of groups. Within the Groups section of the VSC Web GUI, you can define departments and applications.

You can use departments to group all resources for a specific geographical, organizational, or logical unit to have a single pane overview. Figure 12-119 shown an overview of the Cisco department.

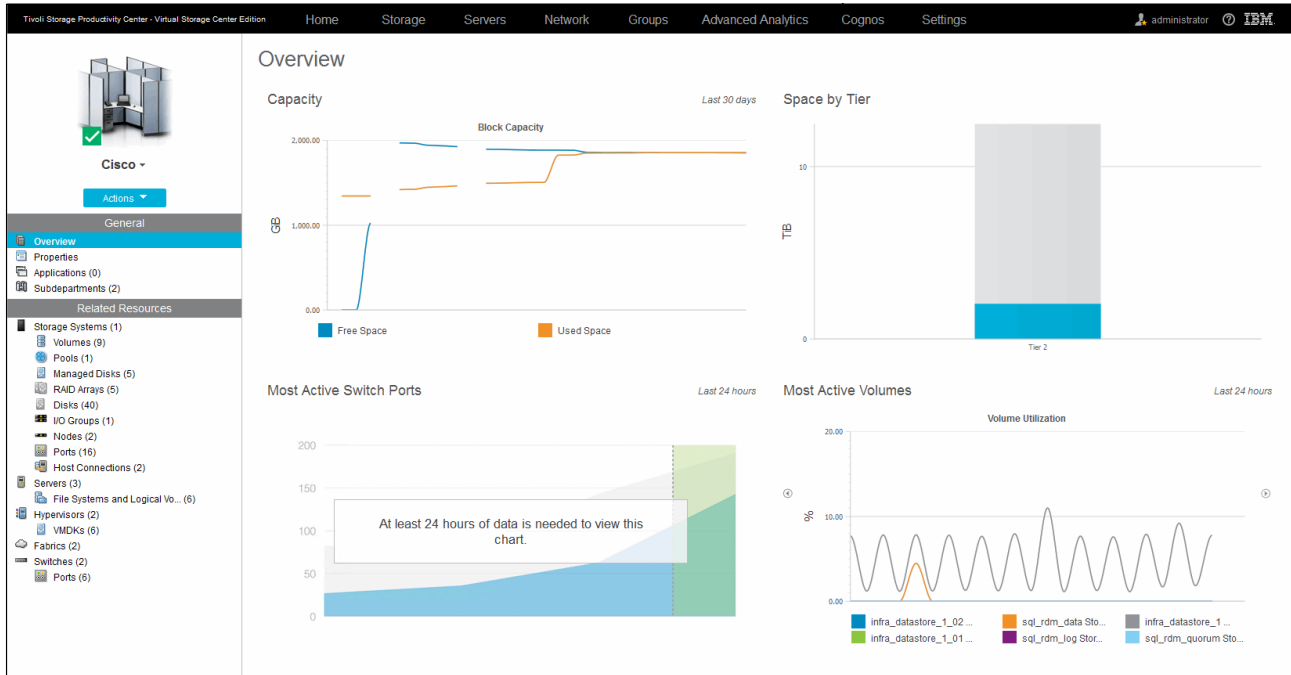


Figure 12-119 VSC Departments Overview

Within a department, you can create subdepartments. Each department or subdepartment can hold one or more applications.

An application is a logical grouping of one or more systems that in turn can be organized into subcomponents.

In the VersaStack setup example, we defined a UCS cluster application and added this application to the Cisco Labs subdepartment.

Figure 12-120 shows the UCS cluster application overview from the VSC Web GUI.

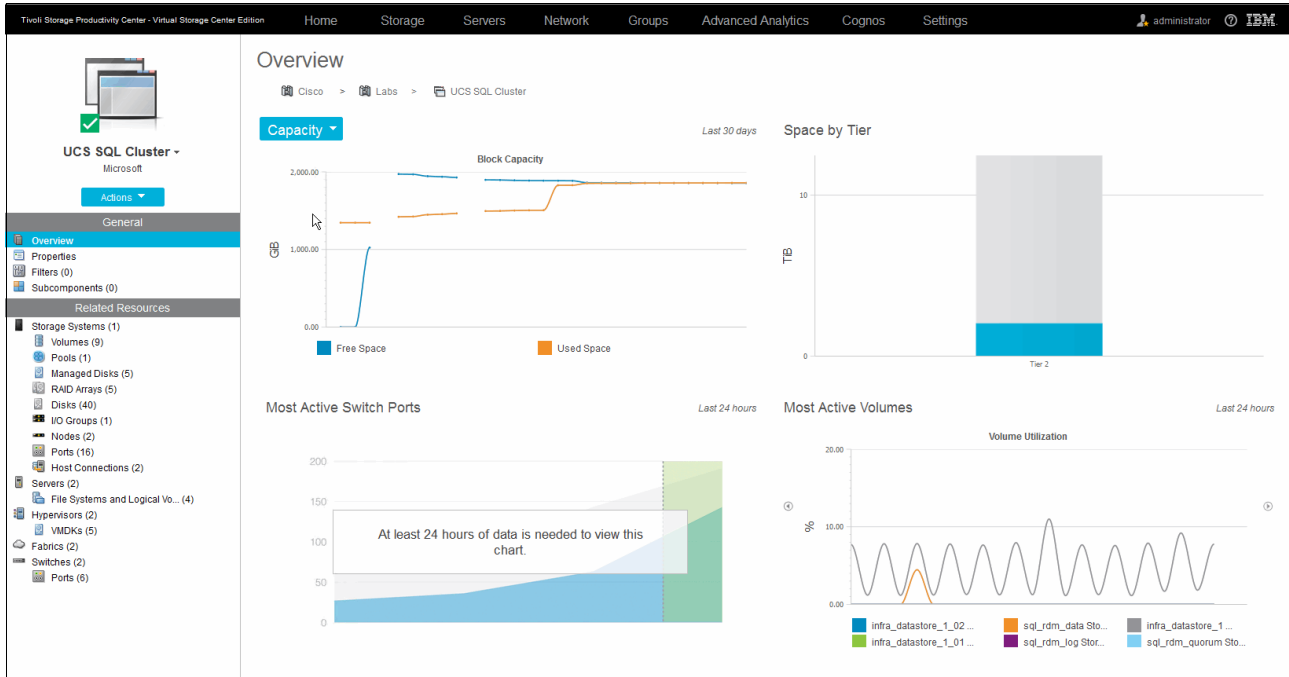


Figure 12-120 VSC Applications

Reporting interfaces

Tivoli Storage Productivity Center VSC provides multiple user interfaces for viewing reports about the storage infrastructure in an enterprise environment.

Cognos Business Intelligence Reporting interface

This interface runs in a web browser. Use this interface to view predefined reports and create custom reports about Tivoli Storage Productivity Center. You access reports from the web-based GUI, and work with the reports in the Cognos BI reporting tool.

To access these reports, select **Cognos** from the VSC Web GUI, as shown in Figure 12-121.




Figure 12-121 VSC Cognos reports

VSC includes several predefined reports and built-in templates to create your own capacity or performance reports with the built-in Report Studio. The following report packages are available when creating custom reports:

- ▶ Capacity and Relationships
- ▶ Historical Capacity
- ▶ Performance
- ▶ Storage Tiering

Both the predefined and your custom created reports can be scheduled to run at specific intervals and stored, exported, printed, or mailed in multiple formats, such as HTML, PDF, and XML.

For the example VersaStack DB2 HA setup, three reports were of specific interest to us and scheduled to be mailed daily:

- ▶ Summarized Performance of Volumes by Server (for the cluster servers)
- ▶ Performance of One Storage System (for the Storwize V7000 storage system)
- ▶ Summarized Performance of Volumes by Hypervisor

For more information about all the available reports, see the following website:

http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.6/com.ibm.tpc_V526.doc/fqz0_c_11_ov_custom_and_predefined_rpts.html

vSphere Web Client extension interface

Use the vSphere Web Client extension for Tivoli Storage Productivity Center to view reports on your virtual environment and storage devices. You can view reports that are customized to use information from Tivoli Storage Productivity Center. The reports include fabric connections, storage mapping information, and performance metrics for storage systems. To view Tivoli Storage Productivity Center storage information in block and file storage reports, you must register Tivoli Storage Productivity Center as a VASA provider.

For more information, see “Spectrum Control hypervisor monitoring and alerting” on page 317.

Stand-alone GUI

You can view detailed information about the storage resources in your environment in the stand-alone GUI. These reports are organized into different types and categories and provide both summary and detailed information, depending on your needs. Many reports are also accessible through the topology viewer, which provides a visual representation of storage topology. To view information about reporting in the stand-alone GUI, go to the product documentation at the following website:

http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.6/com.ibm.tpc_V526.doc/fqz0_c_reporting.html

12.9 Resources

For more information about the topics in this chapter, see the following resources:

- ▶ Tivoli Storage Productivity Center documentation:
http://www.ibm.com/support/knowledgecenter/SSNE44_5.2.6/com.ibm.tpc_V526.doc/tpc_kc_homepage.html
- ▶ VSC wiki:
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20SmartCloud%20Virtual%20Storage%20Center/page/IBM%20SmartCloud%20Virtual%20Storage%20Center%20Wiki>
- ▶ *TPC 5.2.3 Field Level Install Guide for Windows*, found at:
https://www.ibm.com/developerworks/community/wikis/form/anonymous/api/wiki/b6f0fb06-4200-4f2f-9a10-382bddf87c6f/page/f84056cf-76e7-4389-8796-907d9231b2eb/attachment/9d24b843-e00e-4790-b4b5-70e6469fedd0/media/TPC_523_Field_Install_Guide.pdf



IBM Spectrum Protect integration

This chapter describes the implementation of a IBM Spectrum Protect server into the VersaStack environment.

13.1 IBM Spectrum Protect Suite for Unified Recovery overview

This section highlights the features of the IBM Spectrum Protect Suite.

13.1.1 IBM Spectrum Software Defined Storage Suite

New era cloud environments and applications, such as analytics, mobile, and social, are driving huge growth in data volumes, making data the new natural resource. But, cost-effectively optimizing your storage environments while using new opportunities is straining storage budgets. The solution is IBM Spectrum Storage™.

IBM Spectrum Storage unlocks the potential of data and increases your business agility and efficiency in new ways. Spectrum Storage enhances the speed and efficiency of your storage and simplifies migration to new workloads in the following ways:

- ▶ Simplifies and integrates storage management and data protection across traditional and new applications
- ▶ Delivers elastic scalability with high performance for analytics, big data, social, and mobile
- ▶ Unifies silos to deliver data without borders with built-in hybrid cloud support
- ▶ Optimizes data economics with intelligent data tiering from flash storage to tape and cloud
- ▶ Builds on open architectures supporting industry standards, including OpenStack and Hadoop

Figure 13-1 shows IBM Spectrum Protect as part of the IBM Spectrum Software Defined Storage Suite.

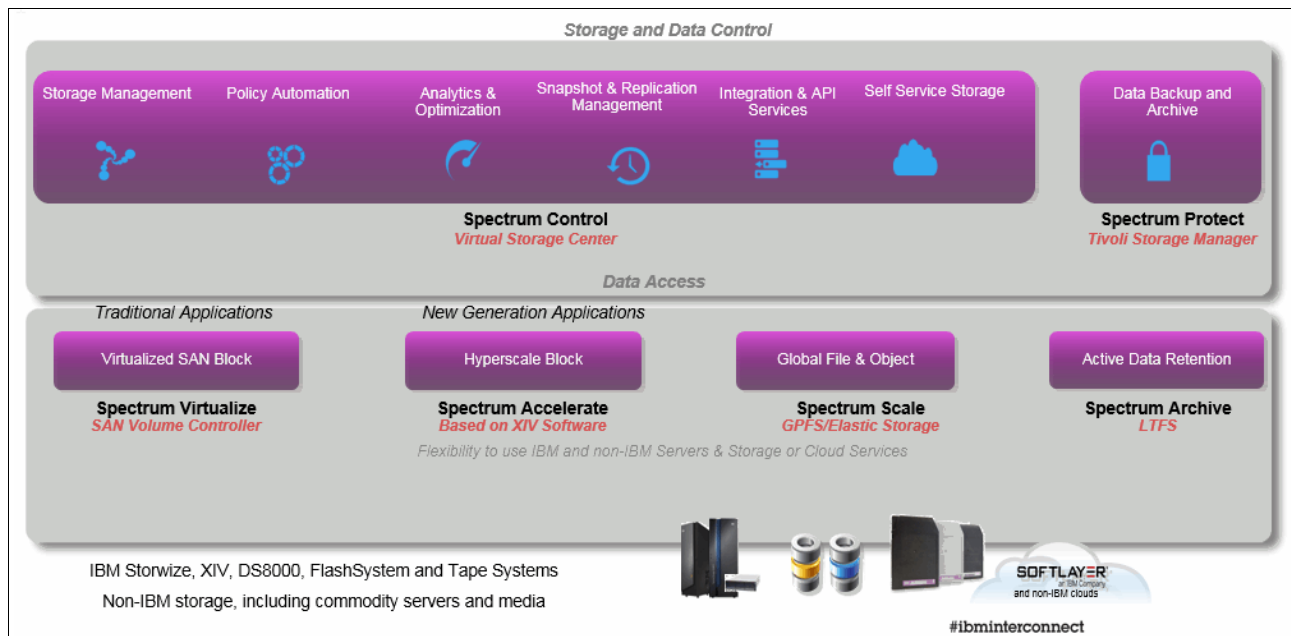


Figure 13-1 IBM Spectrum Software Defined Storage Suite

Within the VersaStack with DB2 solution that is presented in this book, we use the capabilities of three of the IBM Spectrum Software Defined Storage components to complement the functions of the Cisco UCS, IBM Storwize V7000 storage systems, and VMware vCenter. Here are the three components:

- ▶ IBM Spectrum Virtualize™
- ▶ IBM Spectrum Control
- ▶ IBM Spectrum Protect

Here is an overview of all of the components of the Spectrum Software Defined Storage Suite with links to more information:

- ▶ IBM Spectrum Accelerate™
IBM Spectrum Accelerate offers grid-scale block storage with rapid deployment that helps speed delivery of data across an enterprise and adds extreme flexibility to cloud deployments. For more information, see the following website:
<http://www.ibm.com/systems/storage/spectrum/accelerate/index.html>
- ▶ IBM Spectrum Scale™
IBM Spectrum Scale is flash-accelerated, industrial-strength, highly scalable software-defined storage that enables global shared access to data with scalability and agility for cloud and analytics. For more information, see the following website:
<http://www.ibm.com/systems/storage/spectrum/scale/index.html>
- ▶ IBM Spectrum Virtualize
IBM Spectrum Virtualize is at the heart of IBM SAN Volume Controller and IBM Storwize Suite. It enables these systems to deliver industry-leading virtualization that enhances storage to improve resource utilization and productivity, and streamlines deployment for a simpler, more responsive, scalable, and cost-efficient IT infrastructure. For more information, see the following websites:
 - <http://www.ibm.com/systems/storage/software/virtualization/svc>
 - <http://www.ibm.com/systems/storage/storwize>
- ▶ IBM Spectrum Control
IBM Spectrum Control provides efficient infrastructure management for virtualized, cloud, and software-defined storage to simplify and automate storage provisioning, capacity management, availability monitoring, and reporting. For more information, see the following website:
<http://www.ibm.com/software/tivoli/csi/cloud-storage/>
- ▶ IBM Spectrum Protect
IBM Spectrum Protect enables reliable, efficient data protection and resiliency for software-defined, virtual, physical, and cloud environments. For more information, see the following website:
<http://www.ibm.com/software/tivoli/csi/backup-recovery/>
- ▶ IBM Spectrum Archive™
IBM Spectrum Archive enables you to move automatically infrequently accessed data from disk to tape to lower costs while retaining ease of use and without the need for proprietary tape applications. For more information, see the following website:
<http://www.ibm.com/systems/storage/tape/lufs/index.html>

13.1.2 IBM Spectrum Protect Suite for Unified Recovery

IBM Spectrum Protect (formerly known as Tivoli Storage Manager) Suite for Unified Recovery includes the following components:

- ▶ Backup Server
 - IBM Spectrum Protect (Tivoli Storage Manager) Extended Edition. Includes Operations Center for simplified administration, built-in efficiency features, and advanced disaster recovery.
- ▶ Snapshot Management
 - Manages application-aware snapshots on EMC, Hitachi, NetApp, IBM, and Microsoft Volume Shadow Copy Service (VSS)-compatible Windows storage.¹
 - Enables fast and simple recovery of individual files and volumes.
 - Enables “instant” restore for VMware datastores.
- ▶ Advanced Agent for Virtual Environments
 - Incremental “forever” backup for VMware and Hyper-V. Enables flexible, near-instant recovery.
- ▶ Advanced Agents for Core Applications
 - Online, application-aware multi-threaded backups and restores.
 - Mail agents support Microsoft Exchange and IBM Lotus® Domino®.
 - Database agents support Oracle and Microsoft SQL. IBM DB2 and Informix® are supported in the base backup server.
 - Enterprise Resource Planning agent supports SAP and SAP HANA environments.
- ▶ Space Management
 - Policy-based hierarchical space management, for Linux and AIX systems.

Figure 13-2 on page 359 shows data protection with IBM Spectrum Protect Suite for Unified Recovery.

¹ EMC and Hitachi UNIX support requires Device Agents, which are available separately.

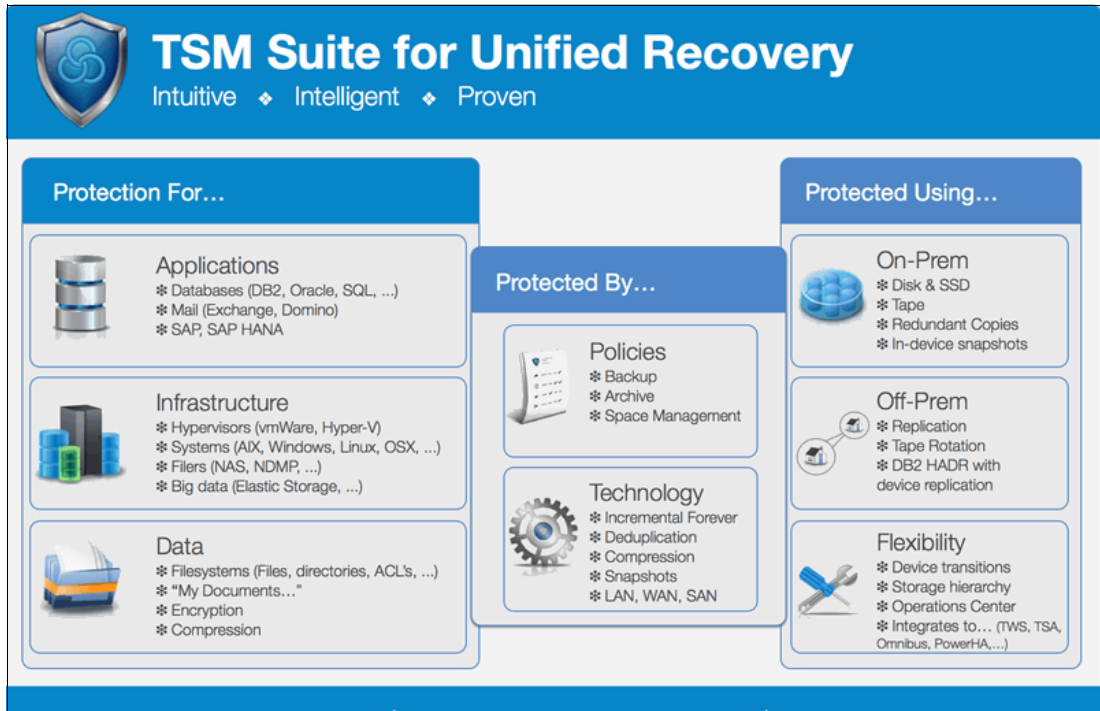


Figure 13-2 IBM Spectrum Protect Suite for Unified Recovery

Within the VersaStack with DB2 environment in this book, we deployed the following components:

- ▶ Tivoli Storage Manager Server
- ▶ Tivoli Storage Manager Operations Center
- ▶ Tivoli Storage Manager/FlashCopy Manager for Virtual Environments
- ▶ Tivoli Storage Manager Backup/Archive Client
- ▶ IBM Tivoli Monitoring for Tivoli Storage Manager

Figure 13-3 shows the Tivoli Storage Manager overview.

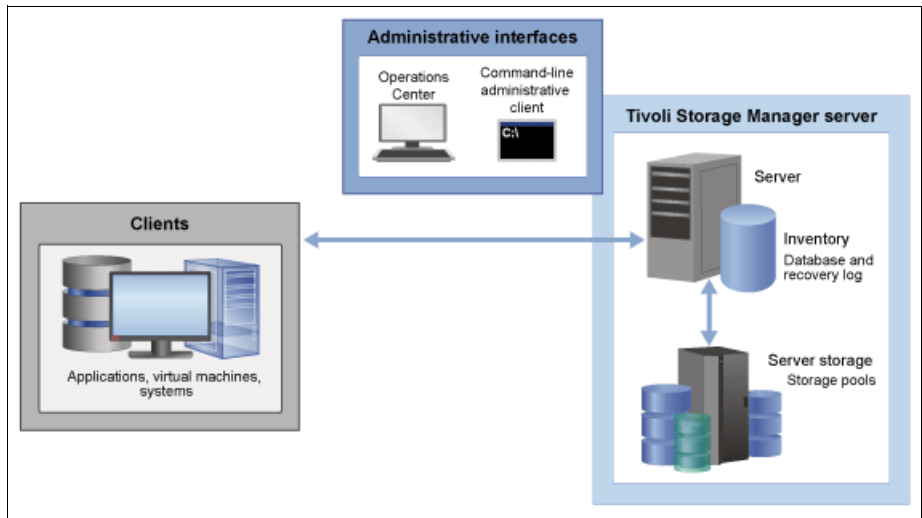


Figure 13-3 Tivoli Storage Manager overview

Server

The Tivoli Storage Manager server stores client data to storage media. The server includes an inventory in which Tivoli Storage Manager stores information about the client data that it is protecting.

Administrative interfaces for the server include a web-based interface that is called the Operations Center and a command-line administrative client. The Tivoli Storage Manager server inventory includes the following components, which can be monitored from the Operations Center.

Database

Tivoli Storage Manager saves information about each file, logical volume, or database that it backs up, archives, or migrates. This inventory data is stored in the server database. The server database also includes information about the policy and schedules for data protection services. Client data is stored in a storage pool.

Recovery log

The recovery log consists of the active and archive logs, and other optional logs. These logs are records of database transactions, which can be used for database recovery. If a failure occurs, such as a power outage or application error, the changes that were made but not committed are rolled back. Then, all committed transactions, which might not yet be written to disk, are redone.

Active log

The active log is a record of the most recent database transactions that are not yet committed.

Archive log

The archive log is a record of the most recent database transactions that are committed but not yet included in a database backup.

Storage

The Tivoli Storage Manager server can write data to hard disk drives (HDDs), disk arrays and subsystems, stand-alone tape drives, tape libraries, and other forms of random-access and sequential-access storage. The media that the server uses are grouped into storage pools.

Storage devices can be connected directly to the server, or connected through a local area network (LAN) or a storage area network (SAN).

Storage pools

Storage pools are a central Tivoli Storage Manager concept. Understanding them is key to effectively managing your Tivoli Storage Manager server environment. Storage pools connect the Tivoli Storage Manager policy hierarchy to the storage devices where client data is stored. A storage pool represents a set of volumes of the same media type, for example, disk or tape volumes.

Tivoli Storage Manager stores all managed data objects in storage pools. You can organize storage pools into one or more hierarchical structures, and each storage hierarchy can span multiple Tivoli Storage Manager server instances.

To obtain the best value from your storage investment, you must store data in the storage pool hierarchy. A disk pool is often first in the hierarchy and can be followed by a tape pool. Tivoli Storage Manager supports many device and media types for sequential access storage.

Figure 13-4 shows how IBM Spectrum Protect automatically places data on the most cost-appropriate tier of storage.

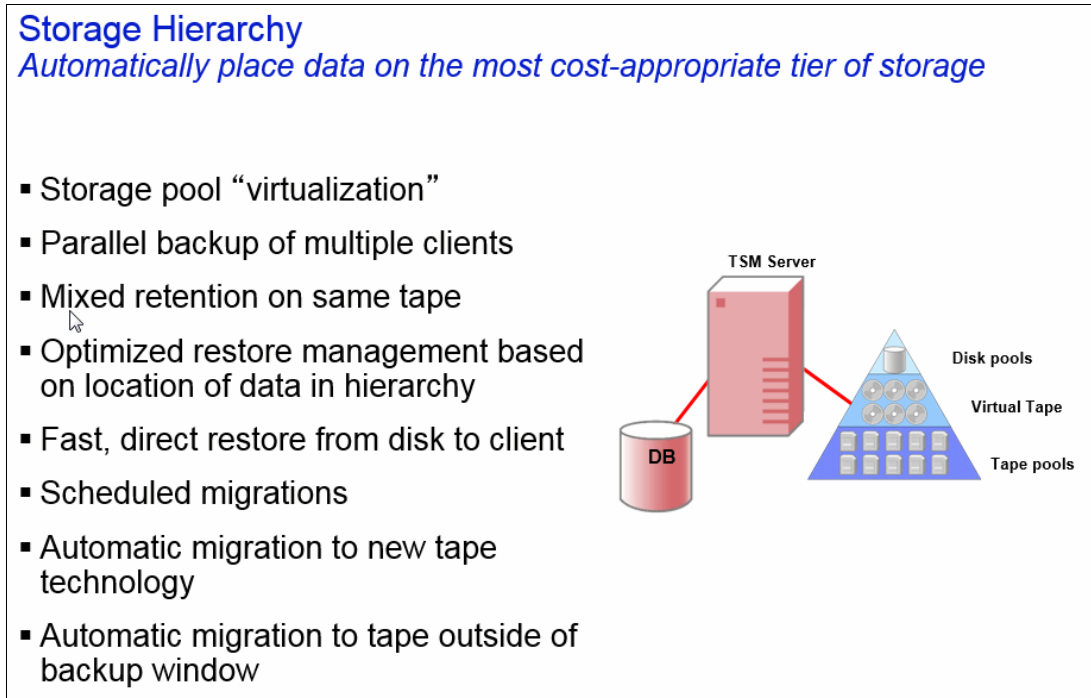


Figure 13-4 IBM Spectrum Protect Storage Hierarchy

Clients

Tivoli Storage Manager clients or client nodes protect data by sending it to a Tivoli Storage Manager server. Client software must be installed on the client system, and the client must be registered with the server.

A client node is equivalent to a computer, such as a backup-archive client that is installed on a workstation for file system backups. A file space is a group of client files that are stored as a logical unit in server storage.

Multiple nodes can be installed on a single computer, as in the case of an IBM DB2 server that contains both an embedded API-based client for DB2 database backups and a backup-archive client for file system backups.

You can define the following clients for use with Tivoli Storage Manager.

Applications

The following clients are application clients. Data that is being protected for these clients is structured data that requires interaction with backup interfaces that are specific to the application:

- ▶ Tivoli Storage Manager for Enterprise Resource Planning
- ▶ Tivoli Storage FlashCopy Manager
- ▶ Tivoli Storage Manager for Databases
- ▶ Tivoli Storage Manager for Mail
- ▶ Tivoli Storage Manager for Virtual Environments

These clients exclude VMware vSphere clients, which are classified as system clients.

A virtual machine (VM) that is backed up by using application client software that is installed on the VM is also classified as an application client.

Virtual machines

A VM is an individual guest that is hosted within a hypervisor. Each VM is represented as a Tivoli Storage Manager file space. Backups for multiple VMs are consolidated together under a common node. Each VM is stored under a separate file space for this common node.

A client is considered a VM when it is protected by either Data Protection for VMware or Data Protection for Microsoft Hyper-V.

Systems

All other clients, for example, backup-archive and API clients, are classified as system clients. These clients back up unstructured data that is contained within files and directories.

System clients also include the following items:

- ▶ A Tivoli Storage Manager source server in a server-to-server virtual volume configuration.
- ▶ A VM that is backed up by using backup-archive client software that is installed on the VM.

Tivoli Storage Manager Operations Center

The Operations Center provides web and mobile access to status information about the Tivoli Storage Manager environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the Tivoli Storage Manager command line.

Figure 13-5 shows the Tivoli Storage Manager Operations Center as deployed in VersaStack and the DB2 environment.

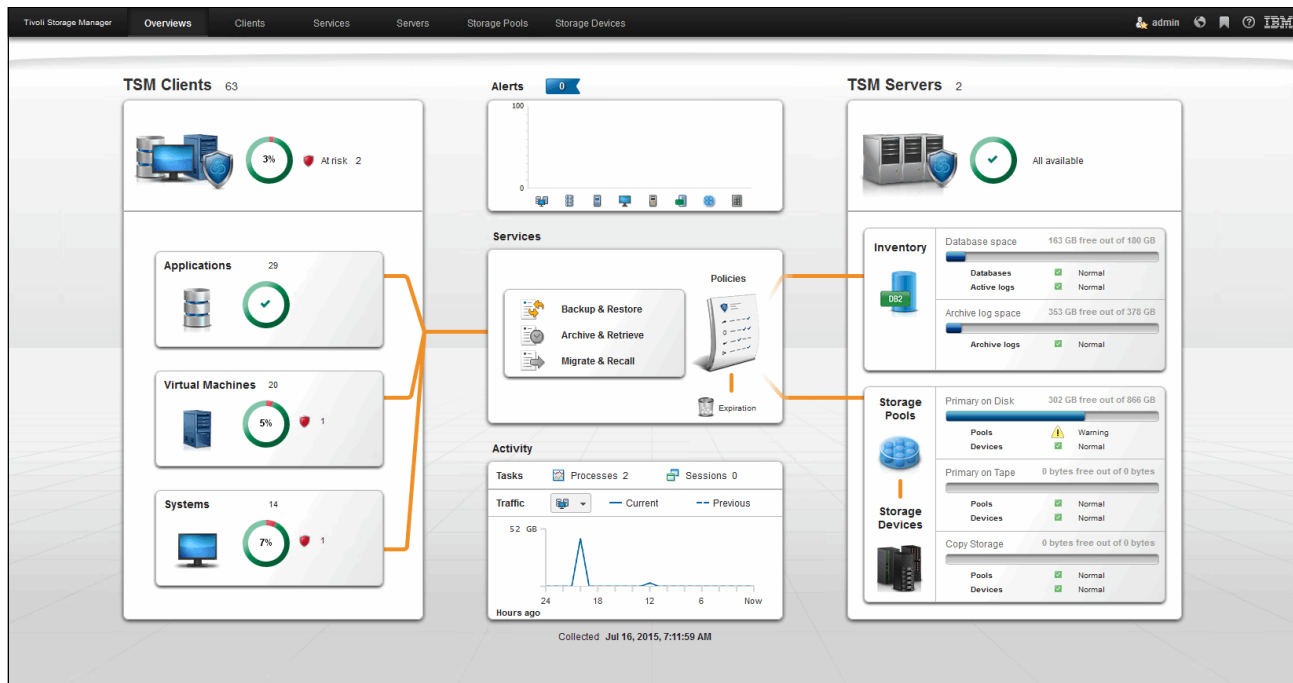


Figure 13-5 Operations Center overview

From the Operations Center, you can complete daily monitoring tasks to ensure that the Tivoli Storage Manager system is functioning correctly.

You can explore the Tivoli Storage Manager Operations Center yourself by exploring the live demonstration environment at the IBM Service Engage website:

<https://demo.tsm.ibm.serviceengage.com:11090/TSMLiveDemo>

For more information, see 13.6, “Monitoring and managing the IBM Spectrum Protect environment” on page 399.

Tivoli Storage Manager for Virtual Environments and FlashCopy Manager for Virtual Environments

IBM Tivoli Storage Manager for Virtual Environments (referred to as Data Protection for VMware) provides a comprehensive solution for protecting VMs.

Data Protection for VMware eliminates the impact of running backups on a VM by offloading backup workloads from a VMware ESX or ESXi-based host to a vStorage Backup server. Data Protection for VMware works with the Tivoli Storage Manager backup-archive client (installed on the vStorage Backup server) to complete full and incremental backups of VMs. The client node that is installed on the vStorage Backup server is called the *data mover node*. This node moves the data to the Tivoli Storage Manager server for storage, and for VM image-level restore later. Instant restore is available at the disk volume level and full VM level. In addition, protection of vApps and organization vDCs in a vCloud Director environment is also available.

Figure 13-6 shows an overview of Tivoli Storage Manager for Virtual Environments and FlashCopy Manager for Virtual Environments.

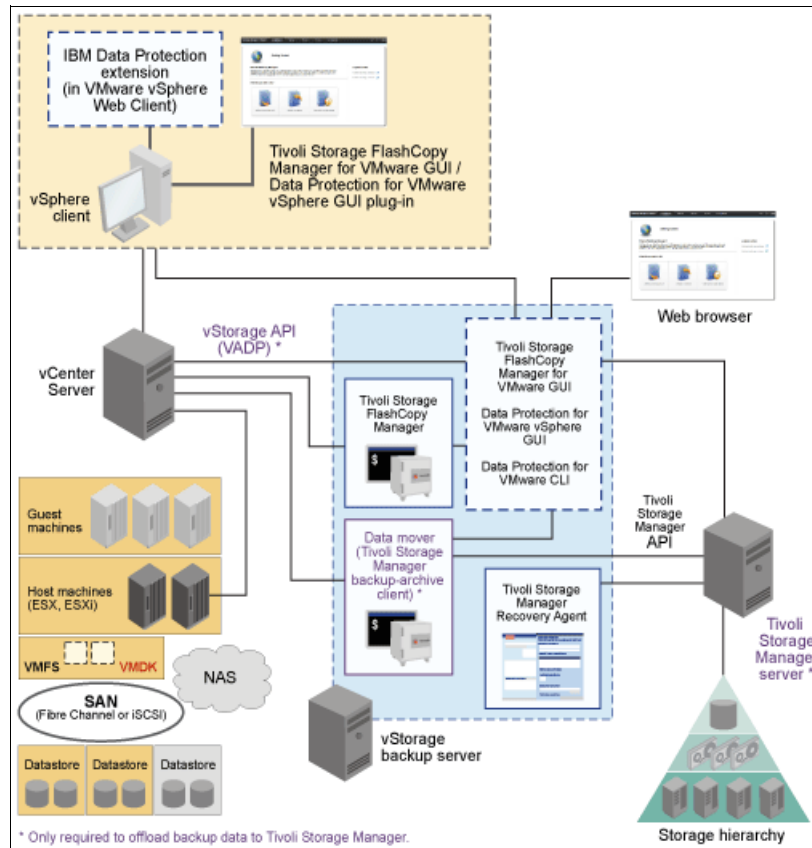


Figure 13-6 Tivoli Storage Manager for Virtual Environments and FlashCopy Manager for Virtual Environments overview

Backup operations in virtualized environments can be separated into in-guest backup, on-host backup, and off-host backup types. Tivoli Storage FlashCopy Manager for VMware uses the off-host backup approach to protect your environment.

Tivoli Storage FlashCopy Manager for VMware supports data protection of virtualized environments by providing off-host storage hardware snapshot backups for VMware VMs. You can install Tivoli Storage FlashCopy Manager for VMware on a physical system or on a VM that has network access to the vCenter Server. The physical or VM where Tivoli Storage FlashCopy Manager for VMware is installed is referred to as the vStorage backup server. Unlike the in-guest backup approach, backup agents are not required to run in each VM. This off-host approach facilitates faster backup operations and is nondisruptive to production applications.

The following list includes the major features when off-host backups are started on a dedicated vStorage backup server or VM:

- ▶ File-level and guest-level image backups can be created and recovered.
- ▶ Centralized management of backup data is provided.
- ▶ Backups are offloaded to free up production server resources.
- ▶ File system consistent backups can be created by using snapshots.
- ▶ Tivoli Storage FlashCopy Manager for VMware and Tivoli Storage Manager for Virtual Environments backups use the VMware vStorage API for Data Protection (VADP).

For more information, see 13.3, “Protecting the VMware infrastructure” on page 379.

Tivoli Storage Manager Backup/Archive Client

You can use the backup/archive client program to back up and archive files from your workstation or file server to storage, and restore and retrieve backup versions and archived copies of files to your local workstation. It includes the following features:

- ▶ An administrative client program that you can access from a web browser or from the command line. A Tivoli Storage Manager administrator can use the program to control and monitor server activities, define storage management policies for backup, archive, and space management services, and set up schedules to perform those services at regular intervals.
- ▶ An application programming interface (API) that you can use to enhance an existing application with storage management services. When an application is registered with a server as a client node, the application can back up, restore, archive, and retrieve objects from storage.
- ▶ A web backup-archive client that enables an authorized administrator, help desk person, or other users to perform backup, restore, archive, and retrieve services by using a web browser on a remote system.

Tivoli Storage Manager uses VSS to back up all system state components as a single object to provide a consistent point-in-time snapshot of the system state. System state consists of all bootable system state and system services components.

Tivoli Storage Manager supports VSS on supported Windows clients.

In the environment in this book, we deployed the Tivoli Storage Manager Backup/Archive Client clients on the SQL server node to back up the operating system component, including the system state. In a VMware environment, you might use the Tivoli Storage Manager for Virtual Environments to back up the VM disk (VMDK) hosting the operating system, but this specific SQL cluster deployment is configured to disallow taking snapshots of the SQL VMs themselves, so there is a need for an in-guest operating system backup.

IBM Tivoli Monitoring for Tivoli Storage Manager

Tivoli Monitoring for Tivoli Storage Manager brings together multiple components to provide real-time monitoring, and historical reporting for your Tivoli Storage Manager servers.

Figure 13-7 shows a schematic flow of IBM Tivoli Monitoring for Tivoli Storage Manager.

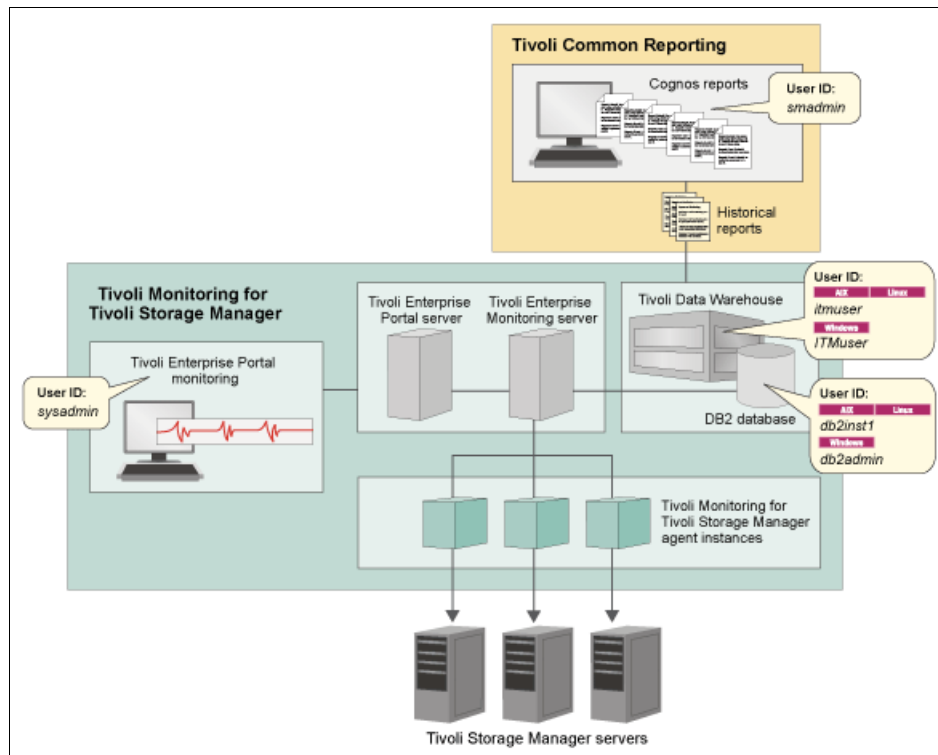


Figure 13-7 IBM Tivoli Monitoring overview

IBM Tivoli Monitoring acts as a monitoring application that provides workspaces for you to monitor real-time information. You can monitor the Tivoli Storage Manager server status, database size, agent status, client node status, scheduled events, server IDs, and so on, by using the monitoring workspaces.

Tivoli Monitoring for Tivoli Storage Manager also provides reports that are based on the historical data that is retrieved. You can use the existing historical reports that are provided, or you can create your own custom reports.

For more information, see 13.6, “Monitoring and managing the IBM Spectrum Protect environment” on page 399.

13.1.3 Licensing metrics

IBM Spectrum Protect Suite for Unified Recovery provides a comprehensive set of data protection capabilities with simplified licensing on a tiered per-terabyte (TB) metric.

The suite features capacity-based licensing, advanced agents for virtual environments and core applications, snapshot management, and hierarchical space management. It lets you get started quickly to gain the benefits of IBM data protection software.

You can use IBM Spectrum Protect Suite for Unified Recovery to more easily modernize data protection. It enables the following capabilities with simple pay as you grow licensing:

- ▶ Protect data with confidence with policy-based management, visual administration, and a scalable IBM platform.
- ▶ Reduce costs for backup infrastructure and administration so that you can invest more in innovation.
- ▶ Add advanced capability that your organization needs to deliver maximum data availability and mitigate the risk of data loss. Application-aware backup agents and hardware-assisted snapshot management capabilities are included.

Figure 13-8 shows the IBM Spectrum Protect license bundles per server or per capacity.

Tivoli Storage Manager Solution bundles at-a-glance			
	Per component per server	Capacity	Capacity
	IBM Tivoli Storage Manager Entry	IBM TSM Suite for Unified Recovery - Front End	IBM Tivoli Storage Manager Suite for Unified Recovery
Available Components			
IBM Tivoli Storage Manager	Standard Edition	Extended Edition	Extended Edition
IBM Tivoli Storage Manager for Virtual Environments	✓	✓	✓
IBM Tivoli Storage Manager for Mail	✓	✓	✓
IBM Tivoli Storage Manager for Databases	✓	✓	✓
IBM Tivoli Storage Manager for Enterprise Resource Planning	✓	✓	✓
IBM Tivoli Storage Manager backup-archive client for file systems	✓	✓	✓
IBM Tivoli Storage Manager for Storage Area Networks	✓	✓	✓
IBM Tivoli Storage Manager for Space Management		✓	✓
IBM Tivoli Storage Manager FastBack®		✓	✓
IBM Tivoli Storage Manager Fastback for Microsoft Exchange		✓	✓
IBM Tivoli Storage Manager Fastback for Bare Machine Recovery		✓	✓

Figure 13-8 IBM Spectrum Protect solution bundles

You can choose from flexible licensing options to get the most favorable plan for your organization:

- ▶ Front End: Capacity is licensed the same way users see it, which simplifies show-back and charge-back auditing.
 - Tiered per-terabyte (TB) license metric with built-in discounts as data grows.
 - Entry: Save up to 55%. Entry versions are limited to 100 TB of managed backup data, and two Tivoli Storage Manager servers per enterprise.

- ▶ Back End: Capacity is measured at the backup servers after efficiency features are used.
 - Tiered per-terabyte (TB) license metric with built-in discounts as data grows.
 - Entry: Save up to 55%. Entry versions are limited to 100 TB of managed backup data, and two Tivoli Storage Manager servers per enterprise.
 - Archive: Save up to 80%. Archive option applies to data imported through an archive operation and backed up to Tivoli Storage Manager VTL or tape archive pools. Data that is backed up to other storage pools is fully supported and charged at the standard Suite for Unified Recovery capacity rate.
 - IBM ProtecTIER® Option: Save up to 75%. ProtecTIER Option measures capacity after IBM ProtecTIER data deduplication is used. Assuming 4:1 data deduplication, capacity-based licensing for data that is stored with ProtecTIER is 75% less than the top tier rate.

For more information, see the following website:

<http://www.ibm.com/software/products/en/tsm-suite-for-unified-recovery>

13.2 IBM Spectrum Protect implementation

This section details the components that are implemented in the DB2 on VersaStack environment in this book.

13.2.1 Architectural overview

This section outlines which IBM Spectrum Protect server components are deployed in the DB2 on VersaStack environment.

IBM Spectrum Protect components

Here are the IBM Spectrum Protect core components:

- ▶ IBM Spectrum Protect backup server
- ▶ Spectrum Operations Center
- ▶ IBM Tivoli Monitoring for IBM Spectrum Protect Reporting and Monitoring Server

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 supported operating systems, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21243309#Server%20Table>

Within the DB2 on VersaStack setup, we deploy IBM Spectrum Protect V7.1.1.300 on SUSE Linux Enterprise Server 11 SP3 running in a VM on the second hypervisor of the DB2 on the VersaStack setup in `vm-host-infra-02.versastack.local`.

The minimum requirements 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

As an alternative, Microsoft Windows or another supported guest OS running on the VMware ESXi hypervisor might have been chosen for deployment, as described in 13.2.2, “Guest support for virtual machines and virtualization” on page 370.

Given the limited size of DB2 in our VersaStack lab setup and the expected payload, we assigned the following resources to the IBM Spectrum Protect VM:

- ▶ VM Version: 8
- ▶ CPU: Four vCPUs
- ▶ Memory: 32 GB
- ▶ VNIC0: VM-Production 1 GbE for the management interfaces
- ▶ VNIC1: VM-Backup 1 GbE for the backup data transport
- ▶ VM Virtual Disks:
 - Hard Disk 1: 64 GB, operating system
 - Hard Disk 2: 32 GB, DB2 database
 - Hard Disk 3: 64 GB, DB2 log files
 - Hard Disk 4: 192 GB, DB2 archive log files
 - Hard Disk 5: 512 GB, data deduplication enabled storage pool

All disks are thick-provisioned and lazy-zeroed and hosted on a dedicated datastore on the Storwize V7000 that is named Protect_Datastore_1.

Note: In this lab setup, both the primary production and the secondary backup environment are hosted on the same VersaStack physical infrastructure. In a real-world scenario, use dedicated storage for the backup environment in combination with the IBM Spectrum Protect Node-Replication towards a secondary backup or server, or invest in dedicated backup hardware.

A secondary IBM Spectrum Protect server that acts as the replication target server is deployed with the same specifications as the primary server but hosted on a separate Cisco UCS blade (vm-host-infra-03.versastack.local in) on the Protect_Datastore_2 in the VersaStack_DC_2 datacenter. This is a logical separation because the same underlying hardware is used in our lab setup.

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.

The Operations Center hardware and software requirements can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21653418>

The Operations Center follows a hub-spoke model with the first IBM Spectrum Protect server connected to it acting as the hub server. This system connects to the spoke servers to query information and run commands.

There is a co-relation between the version of the Operations Center and the version of the hub server as new functions are introduced over time, which requires updates on the IBM Spectrum Protect Servers themselves. For more information, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21640917>

Figure 13-9 shows the Operations Center hub-spoke model running on a hub server or a separate computer.




Figure 13-9 Operations Center hub-spoke model

In the lab setup, we deployed the Linux version of the Operations Center on the primary *spectrumprotect* server itself, which acts as the hub server with the *spectrumreplica* server being the monitored spoke.

For larger environments with multiple IBM Spectrum Protect servers, use dedicated IBM Spectrum Protect server instances running in a virtual environment as the hub server so that you can upgrade the Operations Center and the hub server to new code levels and plan for upgrades to the production spoke servers later.

IBM Spectrum Protect Reporting and Monitoring Server

In the DB2 on VersaStack lab setup, we deploy the IBM Tivoli Monitoring Server for IBM Spectrum Protect on a Windows Server 2008R2 server running in a VM with the following specifications:

- ▶ VM Version: 8
- ▶ CPU: Four vCPUs
- ▶ Memory: 16 GB
- ▶ VNIC0: VM-Production 1 GbE for the management interfaces
- ▶ VM Virtual Disks: Hard Disk 1: 80 GB, operating system and application, thin-provisioned

You can also deploy the Reporting and Monitoring Server on AIX and Linux operating systems. The list of hardware and software requirements can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21678084>

On this system, we configured two IBM Spectrum Protect Monitoring Agents that perform an hourly agentless query towards the IBM Spectrum Protect servers.

13.2.2 Guest support for virtual machines and virtualization

VM and virtualization guest support for IBM Spectrum Protect products is subject to the following supported configurations and limitations for all virtualization technologies:

- ▶ The guest must be running an operating system that is supported by the Tivoli Storage Manager product.
- ▶ Tivoli Storage Manager products and components that rely on other IBM and third-party products are supported if the IBM and third-party components are supported by the virtualized environment. Examples of these dependencies are listed but not limited to the following ones:
 - For the Tivoli Storage Manager Server product, DB2 must also support running within the virtualized environment.
 - For the Data Protection products, the application being protected must also support running on that operating system inside a guest.
- ▶ The performance of Tivoli Storage Manager applications, especially the Tivoli Storage Manager server, ultimately depends on the resources that are available to the application, whether deployed in a physical or virtual environment. For guidance about resource considerations for the Tivoli Storage Manager server, see 13.2.3, “Blueprints” on page 371.

VMware ESX and ESXi Guest

The support position for the following Tivoli Storage Manager products and components is for backup and recovery within the VMware ESX virtual guest, which includes all versions of ESX and ESXi supported by VMware.

Figure 13-10 shows an overview of IBM Spectrum Protect Components that run as a guest on ESX.

TSM PRODUCT/COMPONENT	SUPPORT?	ADDITIONAL SUPPORT INFORMATION
TSM Server TSM Operations Center TSM Reporting and Monitoring	Yes	<ul style="list-style-type: none"> • IBM can make no guarantees with respect to the performance and scalability in a virtualized environment • No support for attached tape drivers or tape libraries, either virtual or physically attached • No support for LAN-Free data movement to tape or disk
TSM Backup-Archive and API clients	Yes	<ul style="list-style-type: none"> • No support for LAN-Free data movement • No support for backupset restore from tape
TSM Storage Agent	No	<ul style="list-style-type: none"> • No support
TSM UNIX HSM (Space Management) clients	Yes	<ul style="list-style-type: none"> • No support for LAN-Free data movement
TSM HSM for Windows clients	Yes	<ul style="list-style-type: none"> • No known restrictions
TSM for Mail (DP for Domino, DP for Exchange)	Yes	<ul style="list-style-type: none"> • No support for LAN-Free data movement
TSM for Databases (DP for Oracle, DP for SQL)	Yes	<ul style="list-style-type: none"> • No support for LAN-Free data movement
TSM for Enterprise Resource Planning	Yes	<ul style="list-style-type: none"> • No support for LAN-Free data movement
TSM FastBack for Workstations / CDP for Files	Yes	<ul style="list-style-type: none"> • No known restrictions
TSM FastBack	Yes	<ul style="list-style-type: none"> • No known restrictions
TSM for SysBack	Not applicable	<ul style="list-style-type: none"> • Not applicable

Figure 13-10 IBM Spectrum Protect supported components on ESX

Almost all IBM Spectrum Protect components are supported in a virtual environment except for the IBM Spectrum Protect Storage Agent, and LAN-free and Tape Library support. A complete list of all virtual environments and the supported IBM Spectrum Protect components can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21239546>

13.2.3 Blueprints

In our DB2 on VersaStack lab setup, we deployed the IBM Spectrum Protect servers manually within the Linux VMs by using the following outline:

1. Deploy SUSE Linux Enterprise Server 11 in the VM and configure the core networking through YaST.
2. Create a user that is named spadmin to host the IBM Spectrum Protect server instance.
3. Created a group that is named tivoli for the spadmin user.
4. Format the VM hard disks for the database, log, archive log and data by using YaST and mount them under the following directories:
 - /tsmdb
 - /tsmlog
 - /tsmarchlog
 - /tsmdedupe
5. Create the /tsminst1 directory to hold the IBM Spectrum Protect instance configuration files and assign spadmin:tivoli ownership to all the directories that are listed in step 4.
6. Copy the TSM_7111_LIN86_AGT_ML.bin file into the VM, extract it, and start the IBM Installation Manager by running install.sh. Select the IBM Spectrum Protect Extended Edition, License, and Operations Center components for deployment.
7. Run the IBM Spectrum Protect Instance configuration wizard (/opt/tivoli/tsm/server/bin/dsmicfgx) and use the directories and user settings.
8. Start the IBM Spectrum Protect Operations Center from <https://spectrumprotect.versastack.local:11090/oc> and complete the initial configuration wizard.
9. Use the built-in CLI from the Operations Center and delete the three default storage pools (backuppool, archivepool, and spacemgpool).
10. Use the built-in CLI from the Operations Center to create the deduplication-enabled storage pool and define a domain for the backup data that uses this pool.

Example 13-1 shows sample commands to define a deduplication-enabled storage pool and VersaStack logical domain and assign the storage pool to the default domain.

Example 13-1 IBM Spectrum Protect commands

```
define devc dedup devtype=file mountl=100 maxcap=10G dir=/tsmdedupe
define stgpool spectrumpedupe dedup maxscr=51 deduplicate=yes identifyprocess=0

def domain VersaStack descript="VersaStack Domain"
def policyset VersaStack PS_VersaStack
def mgmt VersaStack PS_VersaStack MC_VersaStack
def copyg VersaStack PS_VersaStack MC_VersaStack dest=spectrumdedupe
assign defmgmt VersaStack PS_VersaStack MC_VersaStack
validate policyset VersaStack PS_VersaStack
```

```
activate policyset VersaStack PS_VersaStack
```

```
update copygroup STANDARD STANDARD STANDARD STANDARD destination=spectrumdedupe  
update copygroup STANDARD STANDARD STANDARD STANDARD type=archive  
destination=spectrumdedupe  
validate policyset STANDARD STANDARD  
activate policyset STANDARD STANDARD
```

11. Disable deduperquiresbackup and set registration to open through the server properties in the Operations Center.

For more information, see the following website:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.install.doc/t_srv_install_luw-linux.html

IBM Spectrum Protect Blueprints

IBM released installation instructions and an automation tool to perform the tasks that are described in 13.2.3, “Blueprints” on page 371. These instructions are called IBM Spectrum Protect Blueprints.

The Tivoli Storage Manager Blueprint improves time-to-value for Tivoli Storage Manager deployments by providing a set of hardware blueprints for small, medium, and large Tivoli Storage Manager Server architectures. These reference architectures are based on hardware running AIX, Linux, or Windows, and are optimized as disk-only storage by using IBM Storwize or IBM Elastic Storage™ Server (based on IBM Spectrum Scale technology) and Tivoli Storage Manager data deduplication. The architectures are tested to determine the optimal workloads and limits for each size: Small, Medium, or Large. The value proposition is to speed up the sales cycle by matching customer workload requirements to one of the three predefined sizes.

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 does a configuration check to verify 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 by using preferred practices to accomplish the following 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 a Tivoli Storage Manager database backup to disk.
- ▶ Creates a dsmserv.opt file with preferred 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 and provides performance measurements that can be used to compare as a reference against those measured on the blueprint configuration.

When deploying IBM Spectrum Protect in your VersaStack environment, follow the instructions that are outlined for the Small configuration when running IBM Spectrum Protect in a virtual environment on one of the Cisco UCS blades.

Figure 13-11 shows a small Storwize System IBM Spectrum Protect Blueprint configuration overview.

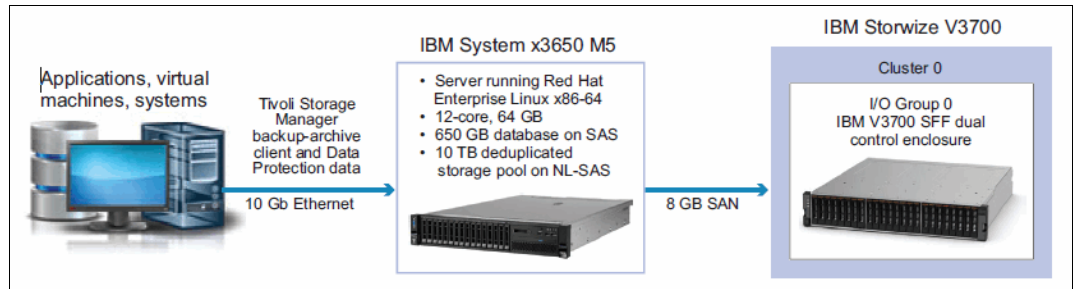


Figure 13-11 IBM Spectrum Protect Blueprint

For proof of concept (PoC) purposes, IBM developed a IBM Spectrum Protect Virtual Appliance that hosts the following components:

- ▶ IBM Spectrum Protect Server
- ▶ IBM Spectrum Protect Operations Center
- ▶ IBM Spectrum Protect for Virtual Environments

This PoC IBM Spectrum Protect VM is based on the Small system version in the IBM Spectrum Protect Blueprints that are published at the following website:

<https://ibm.biz/TivoliStorageManagerBlueprints>

Plans to release IBM Spectrum Protect as Virtual Appliance are being investigated. For more information about the IBM Spectrum Protect PoC Appliance, contact your IBM representative or IBM Business Partner.

13.2.4 Multi-site setup

Deploying a IBM Spectrum Protect server as VM by using shared resources on your primary environment gives you the benefit of advanced data and application protection technologies while maximizing your investment in the VersaStack environment.

Figure 13-12 shows IBM Spectrum Protect running (virtualized) in the primary environment.

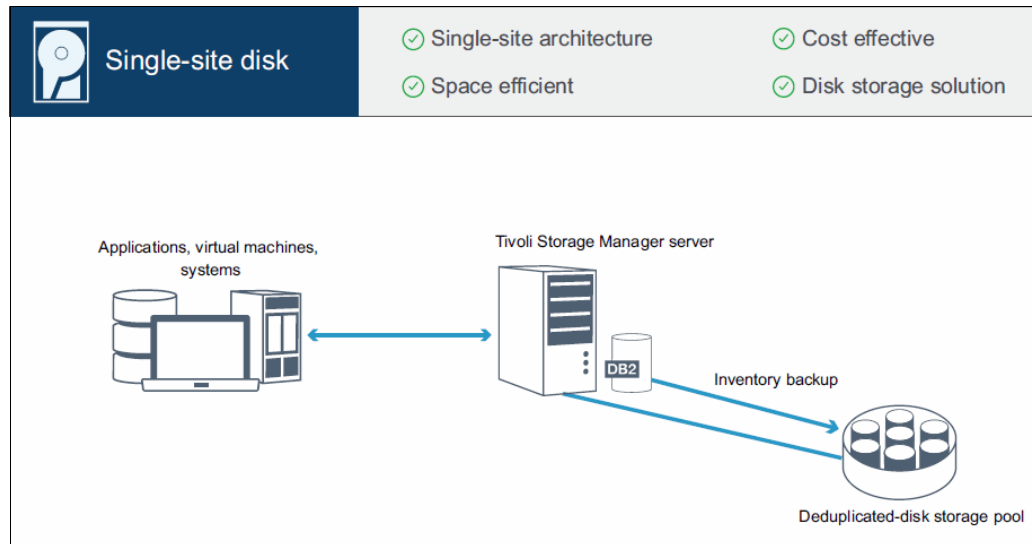


Figure 13-12 IBM Spectrum Protect single-site solution

However, having both your primary data and backup data on the same environment without a secondary copy is not preferred (this is just a lab setup and not a production environment).

IBM Spectrum Protect has multiple high availability and disaster recovery solutions built-in, depending on the storage hardware that is used and the specific requirements:

- ▶ Backup of the IBM Spectrum Protect configuration files, database, and a secondary copy of your primary storagepool to a cospool on tape that is externalized through tape-vaulting
- ▶ Cross-site backup with primary data from site 1 being backed up to site 2 and a copy being sent back to site 1 outside the backup window with cross-site server configuration backup
- ▶ DB2-HA mirroring the IBM Spectrum Protect database and instance in combination with storage mirroring and cross-site cospools to have automated failover between two sites
- ▶ Per client (node) replication between two (cross-site) or multiple (many-to-one) IBM Spectrum Protect servers

Figure 13-13 on page 375 shows IBM Spectrum Protect servers on primary and secondary sites that use node replication.

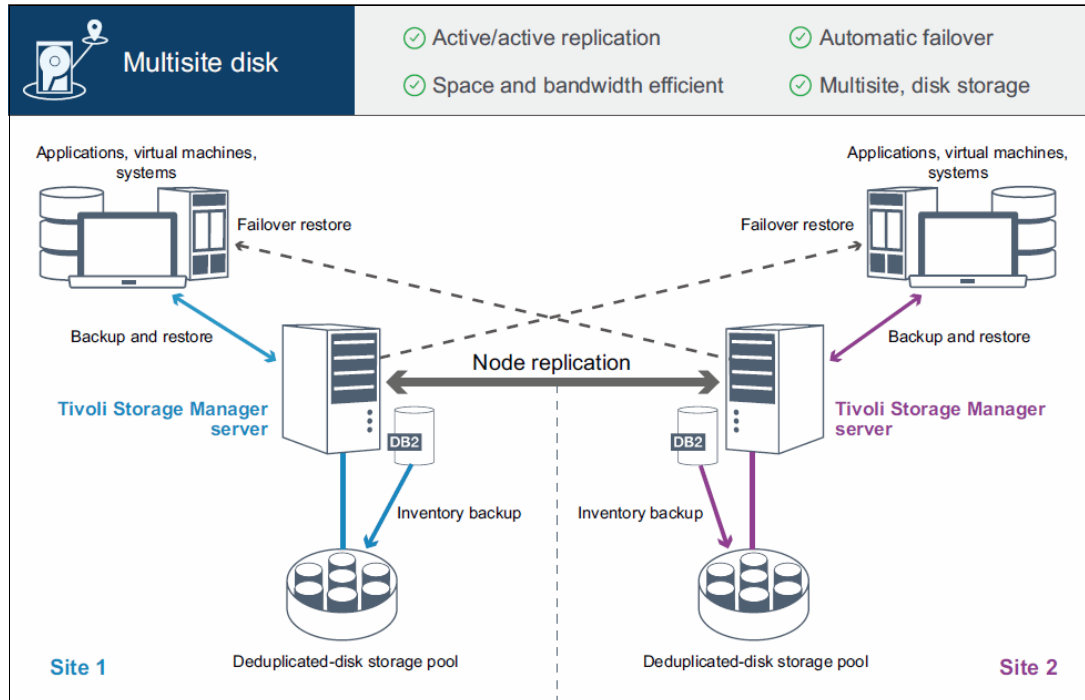


Figure 13-13 IBM Spectrum Protect Multisite

In the DB2 on VersaStack lab setup, we deploy a secondary IBM Spectrum Protect server that is similar to the primary server deployment that is outlined in 13.2.3, “Blueprints” on page 371. After the initial configuration, complete the following steps:

1. Start the IBM Spectrum Protect Operations Center and use the built-in CLI to define a server-to-server connection from the primary to the secondary server over the backup VLAN. Example 13-2 shows how to define server to server communication.

Example 13-2 Define server-to-server communication

```
define server spectrumprotectreplica serverpassword=Object00
hladdress=192.168.60.11 lladdress=1500
```

ANR1660I Server SPECTRUMPROTECTREPLICA defined successfully.

2. Use the Monitor Spoke wizard to register the secondary server as a spoke server and define the primary server on the secondary server.

Figure 13-14 shows the IBM Spectrum Protect servers configured in the Operations Center.

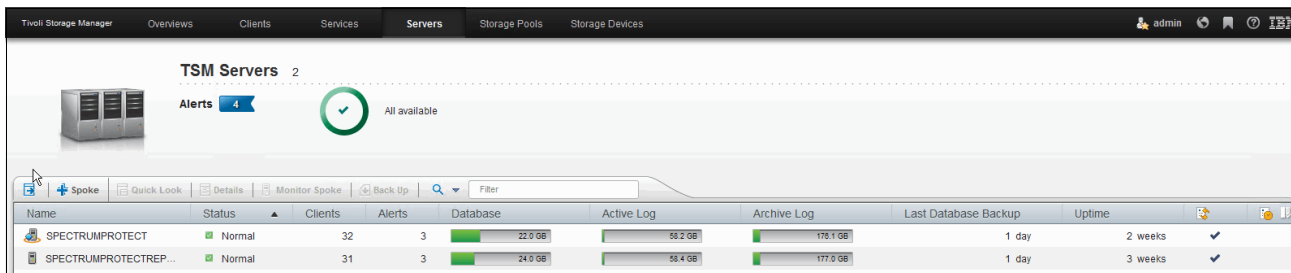


Figure 13-14 IBM Spectrum Protect Operations Center servers overview

3. Within the Operations Center, select the primary server spectrumprotect and click **Details**.
4. Scroll down to Replication, set outbound replication to **Enabled**, and choose spectrumprotectreplica as the peer replication server.

Figure 13-15 shows setting the replication target server through the IBM Spectrum Protect Operations Center.

The screenshot displays the 'SPECTRUMPROTECT' server details in the IBM Spectrum Protect Operations Center. The interface includes a navigation bar at the top with tabs for 'Tivoli Storage Manager', 'Overviews', 'Clients', 'Services', 'Servers', 'Storage Pools', and 'Storage Devices'. The main content area shows the server name 'SPECTRUMPROTECT' with a status indicator 'Normal'. A left-hand sidebar contains navigation options: 'Summary', 'Properties', 'Alerts' (3), 'Active Tasks' (13), and 'Completed Tasks' (12). Below these are 'Related resources' including 'Clients', 'Maintenance', 'Schedules', 'Services', and 'Storage Pools'. The main configuration table is as follows:

Property	Value
Central scheduling	ACTIVE
Poll by client or prompt by server	ANY
Duration for one-time actions	5 days <input type="checkbox"/> Indefinite
Schedule randomization	25 % <input type="checkbox"/> All schedules run at beginning of startup
Client-polling interval	hours <input checked="" type="checkbox"/> Determined by client
Retain schedule events	14 days <input type="checkbox"/> Today only
Inbound sessions disabled	—
Outbound sessions disabled	—
History and Logs	
Retain activity log data	30 days <input type="checkbox"/> Do not retain
Retain activity summary data	30 days <input type="checkbox"/> Do not retain
Create accounting records	OFF
Replication	
Outbound replication	ENABLED
Peer server	SPECTRUMPROTECTREPLICA
Default archive rule	ALL_DATA
Default backup rule	ALL_DATA
Default space-management rule	ALL_DATA
Retain replication history	30 days <input type="checkbox"/> Do not retain

Figure 13-15 IBM Spectrum Protect Operations Center server details

5. Perform the same action on the secondary server spectrumprotectreplica to enable cross-site replication. Both servers are now enabled for node replication.

Node replication

Node replication is the process of incrementally copying, or replicating, data that belongs to backup-archive client nodes. Data is replicated from one Tivoli Storage Manager server to another Tivoli Storage Manager server.

The server from which client node data is replicated is called a *source replication server*. The server to which client node data is replicated is called a *target replication server*. A replication server can function as either a source server, a target server, or both.

Use replication processing to maintain the same level of files on the source and the target servers. When client node data is replicated, only the data that is not on the target server is copied. As part of replication processing, client node data that was deleted from the source server is also deleted from the target server. Client node data is marked for deletion during replication processing, but it is not deleted until expiration processing runs on the target server.

You can maintain different versions of files on the source and target servers or you can maintain files for more or less time on the target server than they are being maintained on the source server. To do this task, you must configure the source and target servers to allow the target server to manage replicated files by using the target server policies.

If a disaster occurs and the source server is temporarily unavailable, client nodes can recover their data from the target server. If the source server cannot be recovered, you can convert client nodes to store data on the target server. When there is an outage, the source server can automatically fail over to a target server for data recovery.

You can use replication processing to recover damaged files. You must replicate the node to the target server before the file damage occurs. Subsequent replication processes detect damaged files on the source server and replace them with undamaged files from the target server.

Figure 13-16 shows automated recovery from the replication server if there are damaged volumes or files on the source server.

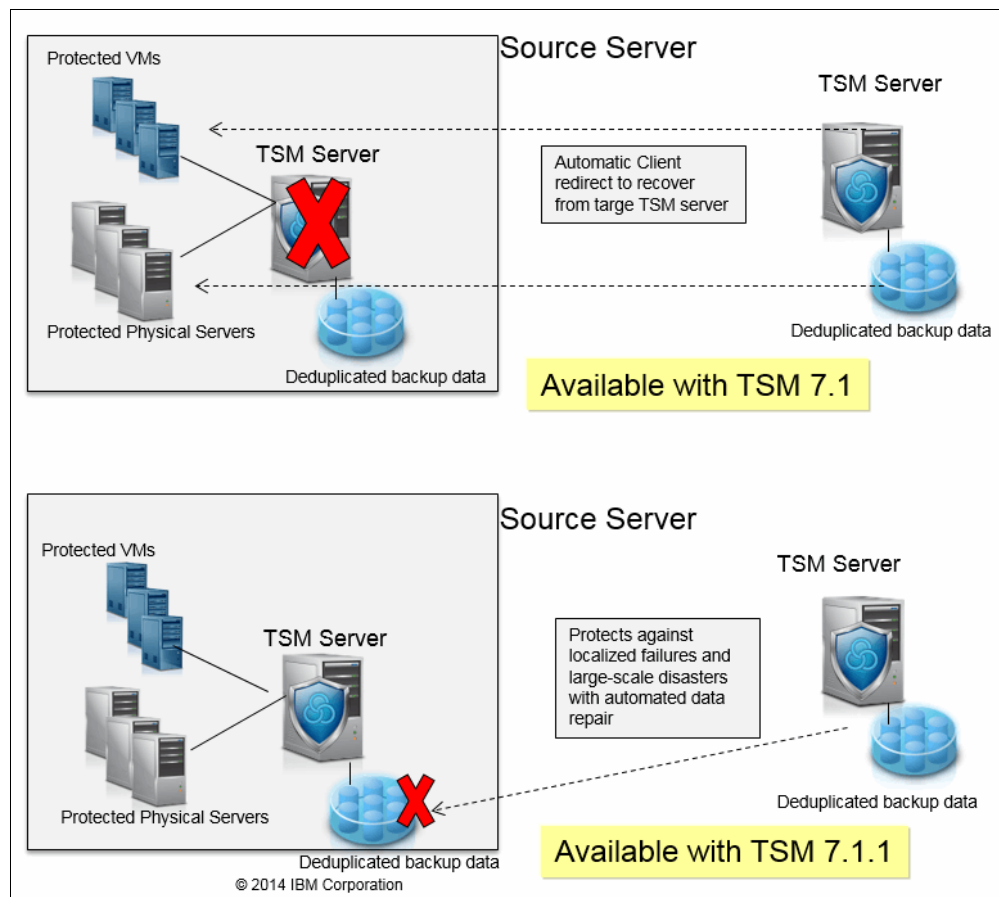


Figure 13-16 IBM Spectrum Protect Node Replication 7.1.1 enhancements

You can replicate the following types of client node data:

- ▶ Active and inactive backup data together, or only active backup data
- ▶ Archive data
- ▶ Data that was migrated to a source server by Tivoli Storage Manager for Space Management clients

Use node replication for data recovery at a disaster recovery site and to maintain the same level of files on the source and target servers. Node replication is used for the following objectives:

- ▶ Controlling network throughput by scheduling node replication at specific times
- ▶ Recovering data from large-scale site loss
- ▶ Recovering damaged files on the source server

Automatic failover for data recovery overview

Automatic failover for data recovery occurs if the source replication server is unavailable because of a disaster or a system outage.

During normal operations, when the Tivoli Storage Manager Version 7.1 client logs in to a source replication server, it receives connection information for the target failover server. The client node stores the failover connection information in the client options file. During client restore operations, the Tivoli Storage Manager server automatically changes clients to the target replication server and back again. Only one failover server can exist per node at any time. The server information is stored in the options file. The failover server can be modified only if the default replication server is modified and another replication is completed for the node.

If the client cannot connect to the source replication server, it uses the failover connection information and attempts to log on to the target failover server. The client logs on to the target replication server and is allowed only to recover data. The client cannot store data during failover processing.

When a new client operation is started, the client attempts to connect to the source replication server. The client resumes operations on the source server if the source replication server is available.

13.2.5 Summary

This section covered the following topics:

- ▶ IBM Spectrum Protect components
- ▶ IBM Spectrum Protect Core Architecture
- ▶ IBM Spectrum Protect deployment by using Blueprints
- ▶ IBM Spectrum Protect high availability setup by using node replication

The base overview of the IBM Spectrum Protect server is complete. The next section covers the backup of the VMware environment on the DB2 on VersaStack setup by deploying the IBM Spectrum Protect for Virtual Environments application module.

13.3 Protecting the VMware infrastructure

This section describes how we used IBM Spectrum Protect in our example VMware environment.

13.3.1 Deploying IBM Spectrum Protect for Virtual Environments

IBM Spectrum Protect for Virtual Environments is an add-on that runs on a separate system called the vStorage backup server, as shown in Figure 13-6 on page 363.

vStorage backup server

This vStorage backup server can either be virtual or physical (when SAN-based data movement towards physical or virtual tape library is a requirement) and hosts the following components:

- ▶ IBM Spectrum Protect/FlashCopy Manager for Virtual Environments stand-alone GUI
- ▶ IBM Spectrum Protect/FlashCopy Manager for Virtual Environments vSphere GUI plug-in
- ▶ IBM Spectrum Protect/FlashCopy Manager for Virtual Environments Command-Line Interface
- ▶ IBM Spectrum Protect for Virtual Environments Datamover
- ▶ IBM Spectrum Protect for Virtual Environments Recovery Agent

Both Windows x64 and Linux x86_64 are supported operating systems for IBM Spectrum Protect for Virtual Environments. However, IBM Spectrum Protect FlashCopy Manager for Virtual Environments requires a Linux x86_64 operating system. Therefore, we deploy two vStorage backup servers in the DB2 on VersaStack setup.

Figure 13-17 shows the opening window of the IBM Spectrum Protect/FlashCopy Manager for VMware vSphere GUI.

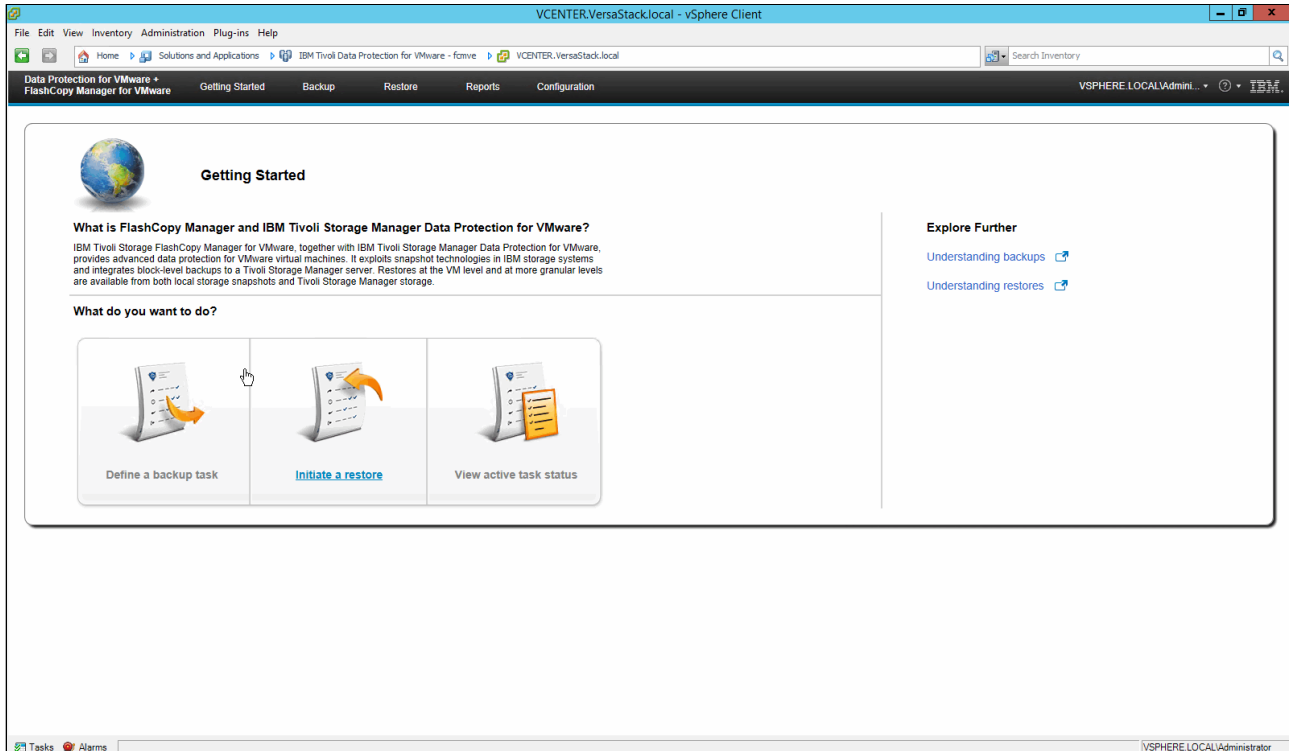


Figure 13-17 IBM Spectrum Protect/FlashCopy Manager for Virtual Environments vSphere GUI

SpectrumVStorage

This VM is a SUSE Linux Enterprise Server 11 SP3 VM with the following specifications:

- ▶ VM Version: 8
- ▶ CPU: Two vCPUs
- ▶ Memory: 4 GB
- ▶ VNIC0: VM-Production 1 GbE for the management interfaces
- ▶ VNIC1: VM-Backup 1 GbE for the backup data transport
- ▶ VM Virtual Disks: Hard Disk 1: 64 GB, operating system, IBM Spectrum Protect for Virtual Environments GUI, vSphere GUI, CLI, Datamover, and Recovery Agent

SpectrumDm

This VM is a Windows Server 2012 VM with the following specifications:

- ▶ VM Version: 8
- ▶ CPU: Two vCPUs
- ▶ Memory: 4 GB
- ▶ VNIC0: VM-Production 1 GbE for the management interfaces
- ▶ VNIC1: VM-Backup 1 GbE for the backup data transport
- ▶ VM Virtual Disks: Hard Disk 1: 64 GB, operating system, IBM Spectrum Protect for Virtual Environments Datamover, and Recovery Agent

The hardware and software requirements for IBM Spectrum Protect for Virtual Environments can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21697958>

The hardware and software requirements for IBM Spectrum Protect FlashCopy Manager for Virtual Environments can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21701160>

The installation process has the following steps:

1. Deploy SUSE Linux Enterprise server on the SpectrumVStorage VM.
2. Deploy Windows Server 2012 on the SpectrumDm VM.
3. Deploy IBM Spectrum Protect for Virtual Environments on the SpectrumVStorage VM, selecting all components.
4. Deploy IBM Spectrum Protect for Virtual Environments on the SpectrumDm VM, selecting the Datamover and Recovery Agent components.
5. Start the IBM Spectrum Protect for Virtual Environments GUI and complete the initial configuration wizard to register the application on the IBM Spectrum Protect Server.
6. Deploy IBM Spectrum Protect FlashCopy Manager for Virtual Environments on the SpectrumVStorage VM.
7. Prepare the target FlashCopy volumes on the Storwize V7000 storage system.
8. Start the IBM Spectrum Protect FlashCopy Manager for Virtual Environments GUI and complete the initial configuration wizard.

This process is outlined in more detail in the *FlashCopy Manager V4.1.1 for VMware and Tivoli Storage Manager for Virtual Environments (Data Protection for VMware 7.1.1) Integrated Installation Cookbook*, which can be found at the following websites:

- ▶ https://www.ibm.com/developerworks/community/groups/service/html/communityview?communityUid=869bac74-5fc2-4b94-81a2-6153890e029a#fullpageWidgetId=W1420ccd1a64d_45f8_8f76_fdbd1fa5cb3e&file=e4f9e51d-32cf-4942-8e00-1f51fa1f5476
- ▶ <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Data%20Protection%20for%20VMware>

In the DB2 on VersaStack setup, we deploy Versions 4.1.2 and 7.1.2, but the same installation instructions apply.

13.3.2 Storwize V7000 FlashCopy mapping

IBM Spectrum Protect FlashCopy Manager for VMware V4.1.2 requires the target volumes to be created and mapped on the Storwize V7000 storage system. For the example environment in this book, complete the following steps

1. Following the instructions that are outlined in *FlashCopy Manager V4.1.1 for VMware and Tivoli Storage Manager for Virtual Environments (Data Protection for VMware 7.1.1) Integrated Installation Cookbook*, create two thin-provisioned target volumes for the `infra_datastore_1` and `infra_datastore_2` volumes that host the VMware datastores.
2. Use the Create FlashCopy Manager from within the FlashCopy Mappings section of the Storwize V7000 GUI to create a FlashCopy consistency group and map the source and target volumes.

Figure 13-18 on page 382 shows using the Storwize V7000 Create FlashCopy Mapping wizard to link the source and target volumes.



Figure 13-18 Storwize V7000 FlashCopy Mappings

In the lab setup, we defined two target volumes. The number of target volumes determines the number of FlashCopy based backups that you can make. If you surpass this number when running a FlashCopy based backup, then the oldest target volume is overwritten. So, with two target volumes, only two restore points are available. Adjust the number of target volumes to your requirements.

Figure 13-19 shows the status of the FlashCopy mapping on the Storwize V7000 GUI.

 A screenshot of the IBM Storwize V7000 GUI showing the 'FlashCopy Mappings' overview. The table has columns for Mapping Name, Status, Source Volume, Target Volume, Progress, Group, and Flash Time.

Mapping Name	Status	Source Vol...	Target Volume	Progress	Group	Flash Time
fcmap4	Copying	infra_datastore_1	infra_datastore_1_02	10%	fccstgrp1	Jun 19, 2015, 7:31:30 PM
fcmap1	Copying	infra_datastore_1	infra_datastore_1_01	8%	fccstgrp0	Jun 23, 2015, 11:12:18 AM
fcmap0	Copying	infra_datastore_2	infra_datastore_2_01	0%	fccstgrp0	Jun 23, 2015, 11:12:18 AM
fcmap2	Copying	infra_datastore_2	infra_datastore_2_02	39%	fccstgrp1	Jun 19, 2015, 7:31:30 PM

Figure 13-19 Storwize V7000 FlashCopy Mappings overview

As you can see, the status of the mappings is Copying because we selected a thin-provisioned volume as the FlashCopy Manager targets to reduce the space that is required for the FlashCopy copies.

You also must specify the **NOCOPY** configuration parameter, as described in *FlashCopy Manager V4.1.1 for VMware and Tivoli Storage Manager for Virtual Environments (Data Protection for VMware 7.1.1) Integrated Installation Cookbook* or by editing the `/opt/tivoli/tsm/tdpvmware/common/scripts/vmcliprofile` file directly. Example 13-3 shows the `vmcliprofile` file.

Example 13-3 IBM Spectrum Protect FlashCopy Manager for Virtual Environments vmcliprofile file

```
>>> GLOBAL
ACS_DIR /home/tdpvmware/tdpvmware/config
ACSD fcmve 57328
# ENFORCE_TLS12 NO
# TRACE NO
<<<
```



```

>>> ACSD
ACS_REPOSITORY /home/tdpvmware/tdpvmware/config/repo
# REPOSITORY_LABEL TSM
# SYNCHRONOUS_RECONCILE RESTORE_AND_DELETE
<<<

>>> VMWARE
VCENTER_SERVER vcenter
AUXILIARY_ESX_HOST vm-host-infra-02.versastack.local
# VCENTER_SERVER_VM_NAME
VCENTER_SERVER_USER administrator@vsphere.local
# FCM_VM_NAME
# VM_BACKUP_MODE SNAPSHOT_EXCL_MEM
# NUMBER_CONCURRENT_VM_TASKS 1
MAX_VERSIONS ADAPTIVE
# HOST_NAME_MAPPING
# TIMEOUT_PARTITION 3600
# TIMEOUT_PREPARE 3600
# TIMEOUT_FLASH 300
# TIMEOUT_VERIFY 3600
# TIMEOUT_CLOSE 3600
# TIMEOUT_FLASHRESTORE 3600
# TIMEOUT_COMPLETERESTORE 3600
<<<

>>> VMCLI
VE_TSM_SERVER_NAME      spectrumprotect
VE_TSM_SERVER_PORT      1500
VE_TSMCLI_NODE_NAME     fcmtsmve_vmcli
VE_VCENTER_NODE_NAME    fcmtsmve_vccenter
DERBY_HOME /home/tdpvmware/tdpvmware
VE_DATACENTER_NAME VersaStack_DC_1::FCMTSMVE_VERSASTACK_DC_1
VMCLI_TRACE NO
VMCLI_SCHEDULER_INTERVAL 60
VMCLI_TASK_EXPIRATION_TIME 864000
VMCLI_RESTORE_TASK_EXPIRATION_TIME 2592000
VMCLI_GRACE_PERIOD 2592000
VMCLI_RECON_INTERVAL_FCM 600
VMCLI_RECON_INTERVAL_TSM 1200
VMCLI_DB_BACKUP AT 00:00
VMCLI_DB_BACKUP_VERSIONS 3
VMCLI_LOG_DIR logs
VMCLI_DB_HOST localhost
VMCLI_DB_PORT 1527
VMCLI_CACHE_EXPIRATION_TIME 600
VMCLI_DB_NAME VMCLIDB
VE_DATACENTER_NAME      VersaStack_DC_2::VERSASTACK_DC_2
<<<

>>> DEVICE_CLASS V7000
COPYSERVICES_HARDWARE_TYPE SVC
COPYSERVICES_PRIMARY_SERVERNAME v7000
COPYSERVICES_USERNAME superuser
# SVC_COPY_RATE 80
# SVC_CLEAN_RATE 50

```

```

# SVC_GRAIN_SIZE 256
COPYSERVICES_REMOTE NO
# COPYSERVICES_COMMPROTOCOL HTTPS
# COPYSERVICES_CERTIFICATEFILE NO_CERTIFICATE
# COPYSERVICES_SERVERPORT 5989
FLASHCOPY_TYPE NOCOPY
# COPYSERVICES_TIMEOUT 6
# RECON_INTERVAL 12
TARGET_SETS 1 2
TARGET_NAMING %SOURCE_0%TARGETSET
<<<

```

13.3.3 Protecting VMware data

With both IBM Spectrum Protect for Virtual Environments and IBM Spectrum Protect FlashCopy Manager for Virtual Environments, you can perform hardware-assisted FlashCopy snapshot-based backups, which have the following features:

- ▶ Near instantaneous backups by using hardware snapshots.
- ▶ The backups are on the Storwize V7000 storage system and require capacity on the primary storage system.
- ▶ Low RTO and RPO possibilities (less than 1 hour) for all the VMs that are hosted on the VMware environment.
- ▶ The primary (storage) environment must be operational in case of restore.
- ▶ Recovery can be either at the VM or at datastore level.
- ▶ Recovery granularity is at the VM level (same or alternative location) or file level by attaching the VMDK to the source or an alternative target VM.

Figure 13-20 shows selecting a FlashCopy based restore point to attach a backed-up VMDK to a VM.

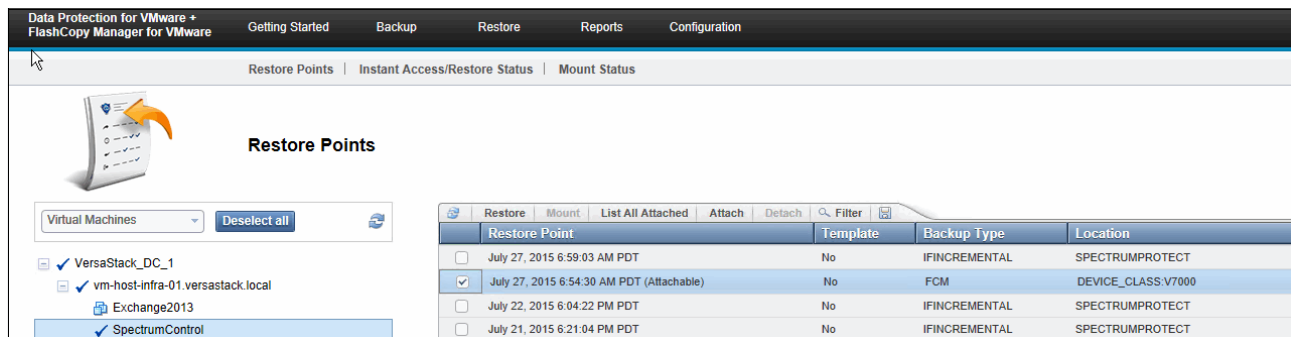


Figure 13-20 IBM Spectrum Protect FlashCopy Manager for Virtual Environments Restore Points

Software-based backups to the IBM Spectrum Protect server

Software-based backups to the IBM Spectrum Protect server have the following features:

- ▶ These backups can use data reduction technologies such as client-side data deduplication and compression and incremental forever backups to perform bandwidth and storage usage optimized backups.
- ▶ Independent backup and long-term copies that are stored on the IBM Spectrum Protect server.

- ▶ RPO in general is 24 hours with RTOs depending on the chosen back-end infrastructure. The Instant-Restore function can provide low RTOs for individual VMs or individual VMs disks.
- ▶ Recovery is at the VM level.
- ▶ Recovery granularity is at the VM (same or alternative location), VMDK (full VMDK), or instant VMDK restore or file level by either attaching the backup copy as a virtual mount point within the source or alternative VM or exposing this virtual mount point as a network share to the user.

Figure 13-21 shows selecting a Tivoli Storage Manager based restore point to expose a backed up VMDK over the network.

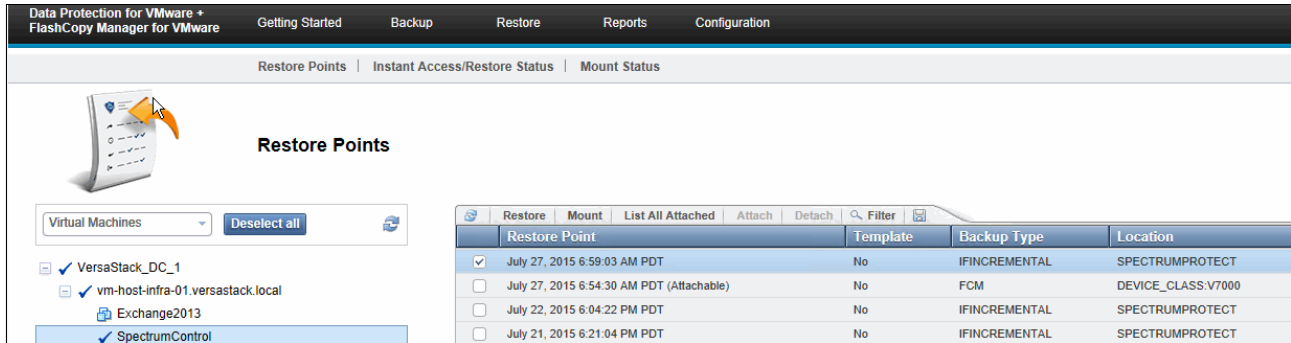


Figure 13-21 IBM Spectrum Protect for Virtual Environments Restore Points

Both the FlashCopy and IBM Spectrum Protect based backups and restores are managed from the same GUI.

IBM Spectrum Protect for Virtual Environments GUI

This GUI can be accessed directly through a web URL or through a vSphere client plug-in, as shown in Figure 13-17 on page 380.

The interface is divided into five sections for easy access to the main functions:

- ▶ Getting Started: Provides information about the available backup and restore functions and links to accomplish the following tasks:
 - Define a backup task.
 - Initiate a restore.
 - View the active task status.

▶ Backup

Figure 13-22 is an overview of IBM Spectrum Protect for Virtual Environment backup schedules.

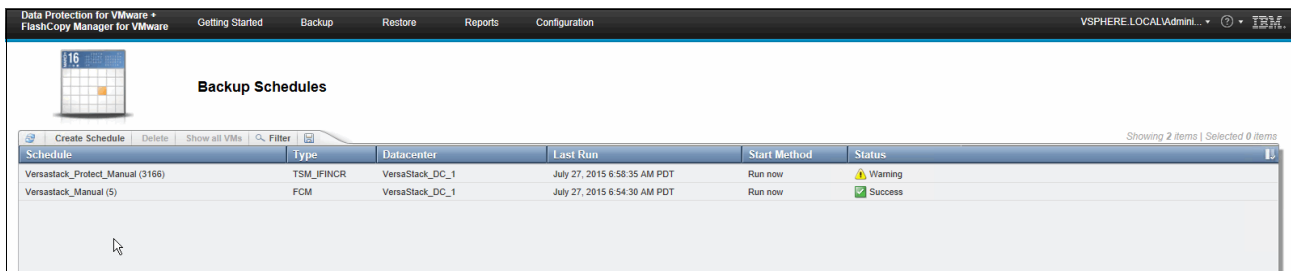


Figure 13-22 IBM Spectrum Protect for Virtual Environments GUI backup overview

From within this section, you can use the Create Schedule wizard to define a backup schedule that performs FlashCopy based backups, IBM Spectrum Protect based backups, or combined backups. You can also define manual *ad hoc* based backups to be run immediately.

A single Datamover instance in the vStorage backup server can back up multiple VMs in parallel (up to 50). This function, when combined with the incremental forever backup technology, reduces the scheduling complexity and the number of schedules that are required.

VMs can be selected at the cluster, host, folder, or VM level with the option to include automatically newly created VMs. VMs or VMDKs can be excluded from backup by using wildcards with IBM Spectrum Protect for Virtual Environments.

For FlashCopy Manager for Virtual Environments clusters, hosts or VMs can be selected for backup by either selecting the cluster, host, or datastore (to have newly created VMs automatically be incorporated in the backup) or individual VMs within specific datastores.

► Restore

Figure 13-23 shows an overview of IBM Spectrum Protect for Virtual Environments datastore restore points.

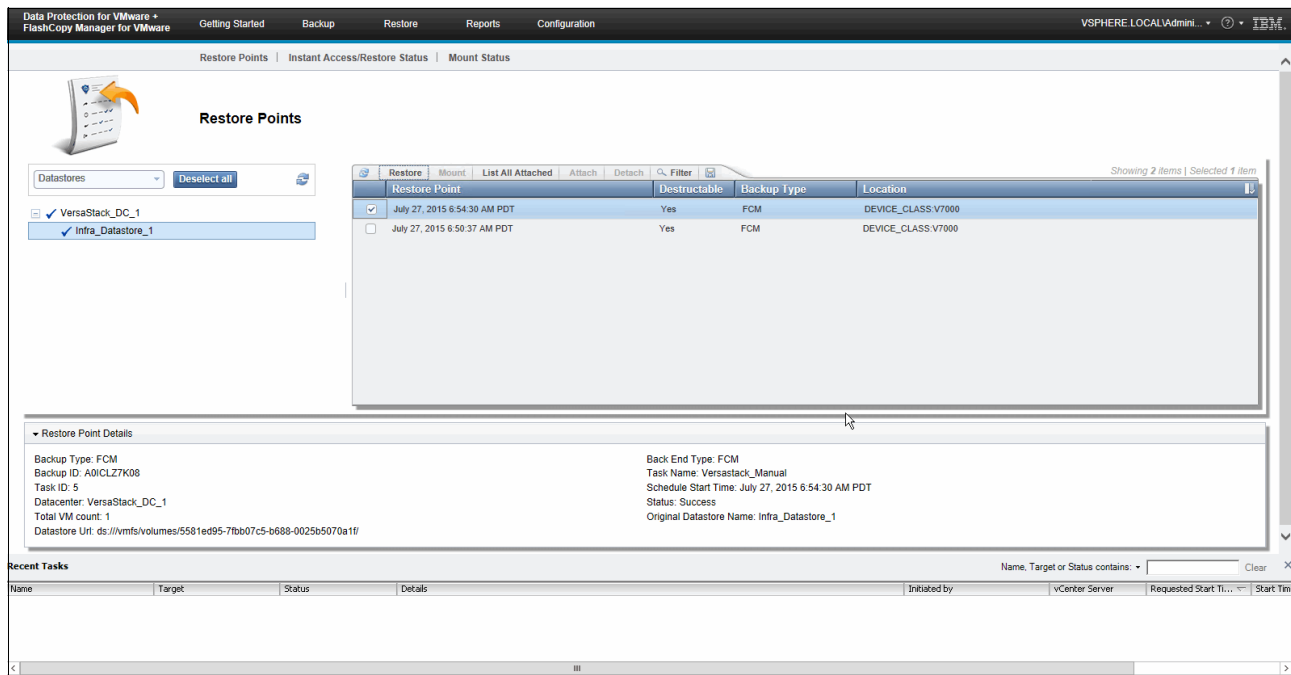


Figure 13-23 IBM Spectrum Protect for Virtual Environments GUI Restore overview

From this pane you switch between datastore- or VM-based restore points:

- VMs restore points:
 - You can select (multiple) VMs to be restored to their original or alternative location.
 - You can select a VMDK of a VM to be mounted onto the vStorage backup server or be exposed through a network (CIFS/NFS) share for specific users.

- You can perform an instant restore of VM or instant access to a VM where the IBM Spectrum Protect server is used as a temporary datastore in the VMware environment with the backup VM started from this datastore for restore consistency verification (instant access) or started from and moved onto the production datastores with VMotion.
 - You can attach a VMDK from a FlashCopy backup to the source or an alternative VM.
- Datastore Restore Points:
- You can select the datastores and VMs of those datastores when you make the FlashCopy backup to perform an instant restore at datastore level and have the selected VMs registered in the VMware environment.

Note: All VMs in the selected datastores are overwritten by the instant restore process. Do not use the instant restore process if you have VMs that are not backed up in the selected restore point.

Next to the Restore Points section, you also have two overview panes that show you the Instant Access/Restore Status and the Mount Status operations that are in progress.

► Reports

Figure 13-24 shows the Backup Status in the IBM Spectrum Protect for Virtual Environments Reports overview.

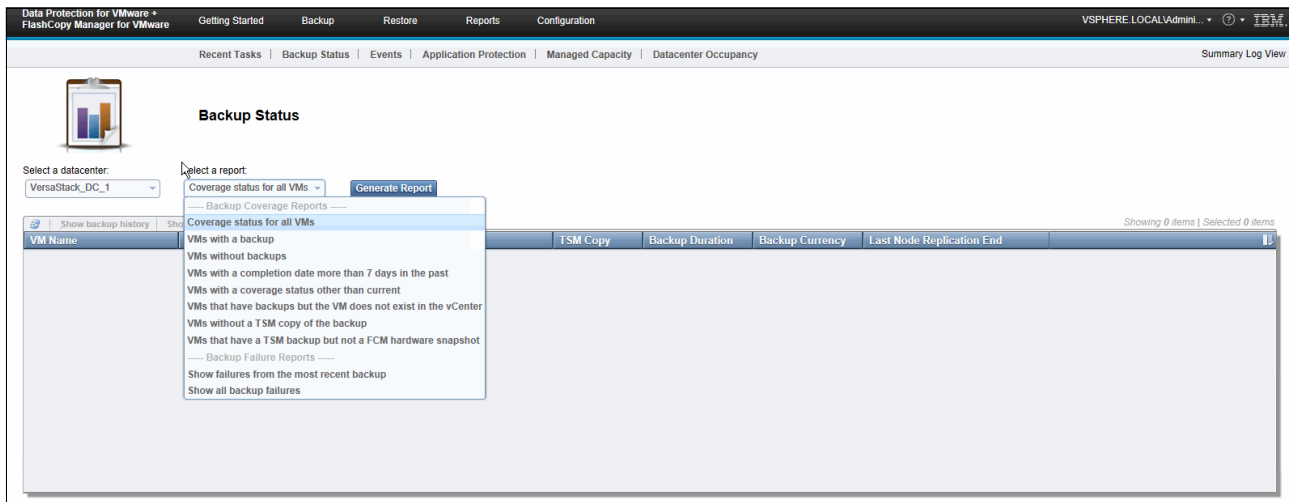


Figure 13-24 IBM Spectrum Protect for Virtual Environments GUI Reports overview

This pane hosts the following subsections:

- Recent Tasks: Gives an overview of the tasks recently run from within the GUI.
- Backup Status: Determines which VMs have a backup (either on IBM Spectrum Protect or on FlashCopy), most recent backup failures, VM coverage status, and more, as shown in Figure 13-24.
- Events: Shows an overview of all events, and completed events with the option to see all or failed VMs that are related to the corresponding event.

- Application Protection: IBM Spectrum Protect for Virtual Environment can scan Windows based VMs to determine which applications that run in these VMs are supported by its agentless application protection capability. Here you can see the following information:
 - Application Configuration Status: Shows which supported applications run in the VM and what kind of IBM Spectrum Protect client or application is deployed in the VM.
 - Unified Component Backup Status: Shows the backup status for both the agentless VM backup as the in-guest IBM Spectrum Protect client or application backup status.
 - Backup Activity Status: Shows the consolidated view of all backups (agentless and in-guest) for the VMs in the selected virtual datacenter.
 - Managed Capacity: The capacity of the datastores that are protected through FlashCopy Manager for Virtual Environments
 - Datacenter Occupancy: The number of VMs in the protected virtual datacenters with the number of VMs being backed up and the occupancy on the IBM Spectrum Protect Server
- Configuration

Figure 13-25 shows IBM Spectrum Protect nodes relationship for the Virtual Environment.

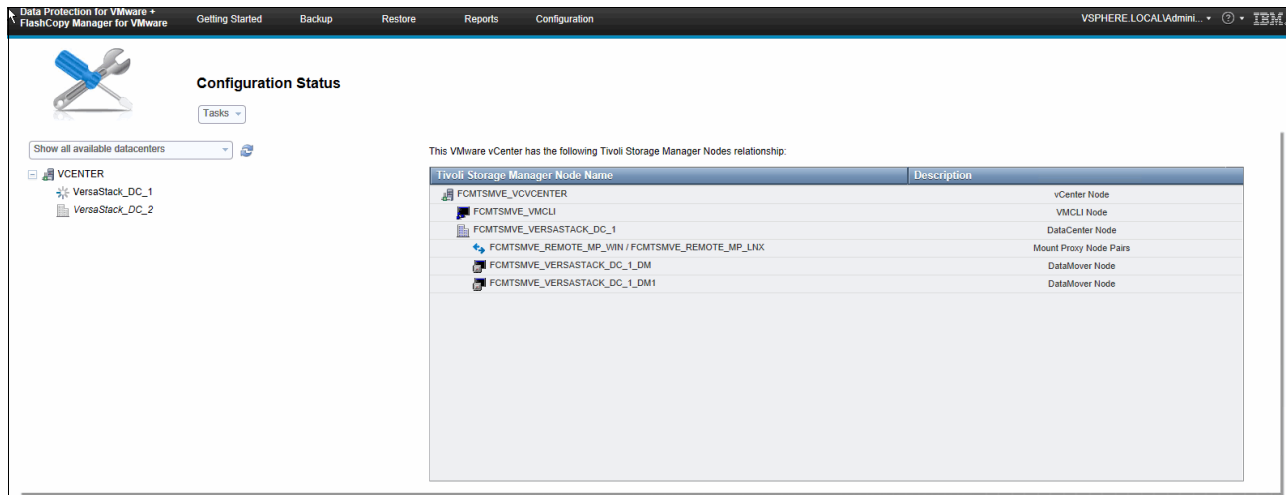


Figure 13-25 IBM Spectrum Protect for Virtual Environments GUI Configuration Overview

The vCenter and its virtual datacenters are mapped to specific virtual nodes on the IBM Spectrum Protect server. VMs are stored on a common virtual datacenter node so that they can be backed up and restored by any datamover. An overview of the IBM Spectrum Protect to VMware node relationship can be seen in this pane. You can also query the connectivity towards the datamovers, run the IBM Spectrum Protect or FlashCopy Configuration wizards, or edit the Tivoli Storage Manager configuration from the Tasks drop-down menu.

IBM Spectrum Protect for Virtual Environments vSphere Web Client Extension

Within the IBM Spectrum Protect Suite, there is a dual approach to backup management. Central backup administrators can manage all backups and restores centrally either through the IBM Spectrum Protect Administration Center, the Operations Center, or through the administrative client by using the command-line interfaces (CLIs).

Also, the backup consumer can run the restore in their familiar working environment by using the IBM Spectrum Protect/FlashCopy Manager for Virtual Environments vSphere Web Client extension.

Figure 13-26 shows the vSphere Web Client Extension Tasks overview.

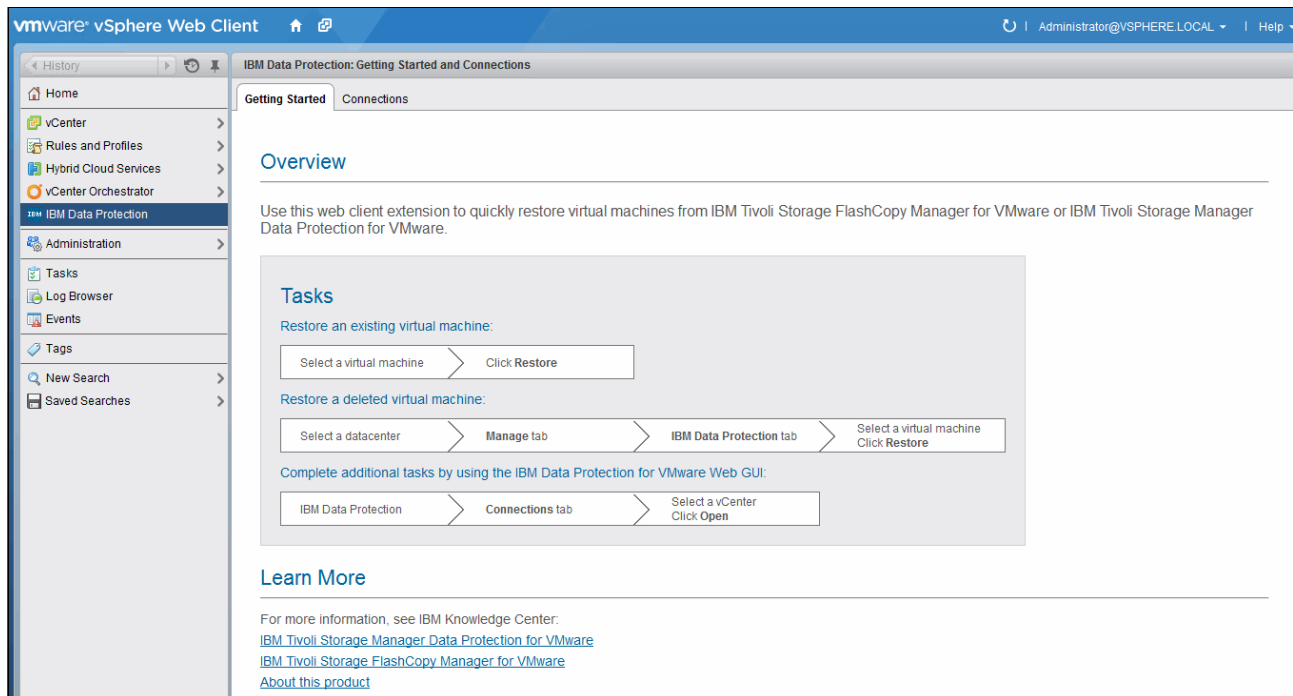


Figure 13-26 IBM Spectrum Protect for Virtual Environments vSphere Web Client Extension

The extension has the following features:

- ▶ You can restore a full VM to its original or an alternative location (IBM Spectrum Protect/FlashCopy Manager for Virtual Environments).
- ▶ Restore a VMDK to its original location (FlashCopy Manager for Virtual Environments).

13.3.4 Summary

This section reviewed the deployment of IBM Spectrum Protect for Virtual Environments and FlashCopy Manager for Virtual Environments and described which functions these products deliver to protect the VMware environment on which the DB2 on VersaStack systems are running.

You used these features to back up the Spectrum Control and other auxiliary VMs in the test setup, such as the Exchange 2013 mail server to which the Spectrum Control and IBM Spectrum Protect automated reports and alert emails were sent.

13.4 Protecting the DB2 HADR virtual machines

Chapter 11, “IBM DB2 High Availability server and failover cluster implementation” on page 171 outlined two alternatives to deploy a DB2 HA setup:

- ▶ Deploying a clustered DB2 instance on the Microsoft Failover Cluster by using physical Raw Device Mapping (pRDM) volumes as data volumes.
- ▶ Deploying two DB2 instances in non-clustered VMs with regular VMDKs for the data volumes and configuring them for use with the built-in HADR feature of DB2.

The first scenario prohibits the use of IBM Spectrum Protect for Virtual Environments to make an agentless backup because no VMware snapshots can be taken from the pRDM by using VADP. In such a scenario, an in-guest backup agent or application must be deployed, as described in 13.4.1, “Using IBM Spectrum Protect to back up DB2 data” on page 390.

In the case of the second scenario, you can use IBM Spectrum Protect for Virtual Environments and IBM Spectrum Protect FlashCopy Manager for Virtual Environments (see 13.3, “Protecting the VMware infrastructure” on page 379) in combination with pre-freeze and post-thaw scripts to quiesce the DB2 databases. Sample commands can be found in 13.4.2, “DB2 quiescing commands” on page 391.

13.4.1 Using IBM Spectrum Protect to back up DB2 data

IBM DB2 has built-in support for IBM Spectrum Protect as a backup target. Here are the configuration steps that are required for DB2 to work with IBM Spectrum Protect on Windows:

1. Install Tivoli Storage Manager API software, which is included with the IBM Spectrum Protect Backup-Archive Client installation package.
2. Create a plain text file to be used as the options file for the DB2 node in `c:\progra~1\common~1\Tivoli\TSM\api`.

The default name is `dsm.opt`, but a different file name can be used, for example, `db2dsm.opt`. The options file must contain at least the following lines:

```
TCPSERVERADDRESS spectrumprotect.versastack.local
NODENAME SQLVM01
PASSWORDACCESS GENERATE
```

3. Register the node `SQLVM01` on the IBM Spectrum Protect server. Use separate node names for your DB2 node and your file-system node.

Be sure to give the node the authority to delete its backups (`BACKDEL=YES`) because DB2 controls how long to keep the backup data.

The DB2 backup objects are all uniquely named and are in an `ACTIVE` state on the IBM Spectrum Protect Server until they are removed by DB2. When DB2 performs the deletions of the data, this action marks the object for immediate expiration on the IBM Spectrum Protect server. If `BACKDEL` is set to `N0` (the default value), then DB2 cannot delete the object. In this case, DB2 makes a second deletion attempt and marks the object as `INACTIVE`. The data then is processed according to the retention values for the management class.

A DB2 restore can only access data that has an ACTIVE state, so set up the management class with a backup copygroup that has the following retention settings:

```
VEREXIST=1  
VERDELETED=0  
RETEXTRA=0  
RETONLY=0
```

Do not set this management class as the default within the policy domain unless you are registering your DB2 nodes in their own policy domain. The above retentions keep only one active copy of a backup, which is not normally the wanted retentions for other file system data.

4. On the system where DB2 and the IBM Spectrum Protect client is installed, create an environment variable called `DSMI_CONFIG` (in uppercase) within the `SYSTEM` environment variable settings (search the Control Panel for “Environment Variables”) that contains the path and file name to the DB2 node's options file. For example:

```
c:\progra~1\common~1\Tivoli\TSM\api\dsm.opt
```

Do not use the `USER` variable except in specific circumstances, such as when there are multiple instances on the same machine that must be backed up through unique node names or options files.

5. Restart the system or restart the DB2 services so that the `DSMI_CONFIG` variable is loaded into the DB2 runtime engine.

Note: If you are binding the DB2 data to a non-default management class, update the DB cfg parameter `TSM_MGMTCLASS` before restarting DB2. For example:

```
db2 update db cfg for db_name using tsm_mgmtclass my_mgmtclass
```

6. Sign on as the administrator and change to the `...\sqllib\adsm` directory.
7. Run `dsmapipw` and follow the prompts. You are prompted for the current password, then for a new password, and confirmation of the new password. The utility should report whether the command was successful. To keep the existing password, enter the same password at all three prompts.
8. It is possible to run the DB2 command `db2adut1 query` to confirm that the password was set correctly. `Db2adut1.exe` is typically in the `c:\program files\sqllib\bin` directory.
9. You can now perform a backup by using the DB2 command line or the IBM DB2 Connect™ GUI, for example:

```
db2 backup db GSDB use tsm
```

13.4.2 DB2 quiescing commands

When you use IBM Spectrum Protect for Virtual Environments or IBM Spectrum Protect Snapshot for Virtual Environments, you can run custom quiescing scripts inside Windows VMs, as described in the VMware Knowledge Base at the following website:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1006671

With these scripts, you take the DB2 database offline, take an agentless backup of the whole VM, and start the DB2 database by putting the following commands into the pre-freeze-script.bat file:

```
CONNECT TO database-alias
QUIESCE DATABASE IMMEDIATE FORCE CONNECTIONS;
UNQUIESCE DATABASE;
TERMINATE;
DEACTIVATE DATABASE database-alias
In the post-thaw-script.bat you would put:
ACTIVATE DATABASE database-alias
```

Likewise, you can deploy a IBM Spectrum Protect Backup/Archive client and the IBM Spectrum Protect API within the VM and have an online backup to IBM Spectrum Protect triggered when the VM that hosts the DB2 is being backed up with IBM Spectrum Protect for Virtual Environments.

13.5 Using IBM Spectrum Protect Advanced Protection and Recovery Technologies

This section briefly covers some of the key IBM Spectrum Protect advanced protection and recovery technologies:


- ▶ Progressive incremental backups
- ▶ Client- and server-side data deduplication
- ▶ IBM Spectrum Protect server high availability

13.5.1 Progressive incremental backups

Figure 13-27 on page 393 shows how much capacity that you can potentially save.

Data Reduction: Progressive Incremental Backup

How Much Can You Save?



**Backup Capacity
Needed for 1 Month:**
Vendor A: 26TB
Vendor B: 14TB
IBM TSM: 7TB

Assumes: Full backup completed, 2TB data to start, 26% annual growth rate, 10% new/changed data per day

Figure 13-27 IBM Spectrum Protect progressive incremental backups

One of the core technologies for file-based backups within IBM Spectrum Protect is *Progressive Incremental Backup*. After the first full backup, only incremental backups are made, which provides the following advantages:

- ▶ Lowers the backup window by eliminating regular full backups.
- ▶ Reduces the back-end storage that is required to hold the backup data.
- ▶ Eliminates and reduces the restore complexity because a single pass restore versus full+differential or full+incremental restores is run.
- ▶ True progressive incremental at backup time with no resource consuming post-backup synthetic full back-up reconsolidation

With the shift towards virtualized server environments, this technology is incorporated into the IBM Spectrum Protect for Virtual Environments application as Incremental Forever, as shown in Figure 13-28.

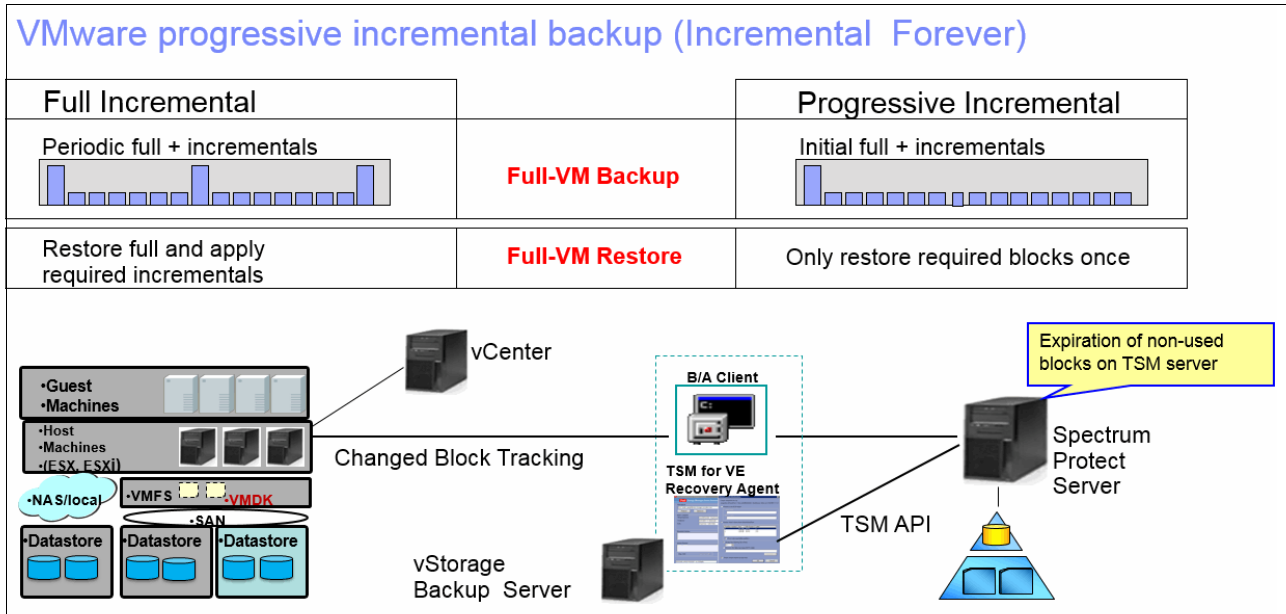


Figure 13-28 IBM Spectrum Protect for Virtual Environments Incremental Forever

Data that is backed up by using progressive incremental backups can be stored on any supported storage medium and does not suffer from tape scattering over time because IBM Spectrum Protect has a built-in collocation mechanism that stores and groups backup data sets (at the file-system or VM level) by using the least amount of tapes, ensuring adequate restore times and eliminating time-consuming redundant tape mounts.

Here is a link to an ESG Lab Review: Tivoli Storage Manager for Virtual Environments that describes in detail the potential savings in network and backup infrastructure resources by using both the progressive incremental and data deduplication technologies to achieve 95% data reduction over just 11 days of backups:

<https://ibm.biz/BdXuvj>

13.5.2 Data deduplication

Data deduplication is a method of reducing storage needs by eliminating redundant data.

Two types of data deduplication are available on IBM Spectrum Protect: client-side data deduplication and server-side data deduplication.

Client-side data deduplication is a data deduplication technique that is used on the backup-archive client to remove redundant data during backup and archive processing before the data is transferred to the Tivoli Storage Manager server. Using client-side data deduplication can reduce the amount of data that is sent over a local area network.

Figure 13-29 on page 395 shows the process of Tivoli Storage Manager client-side data deduplication.

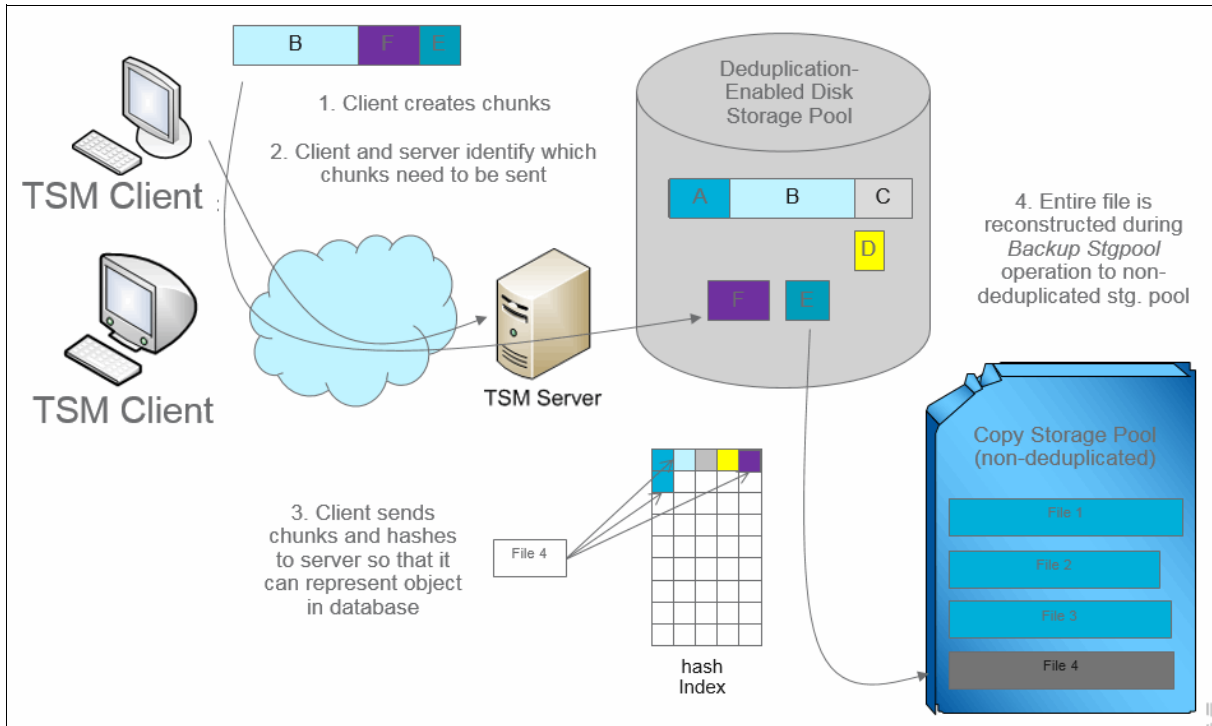


Figure 13-29 IBM Spectrum Protect client-side data deduplication

Server-side data deduplication is a data deduplication technique that is done by the server.

Figure 13-30 shows the process of Tivoli Storage Manager server-side data deduplication.

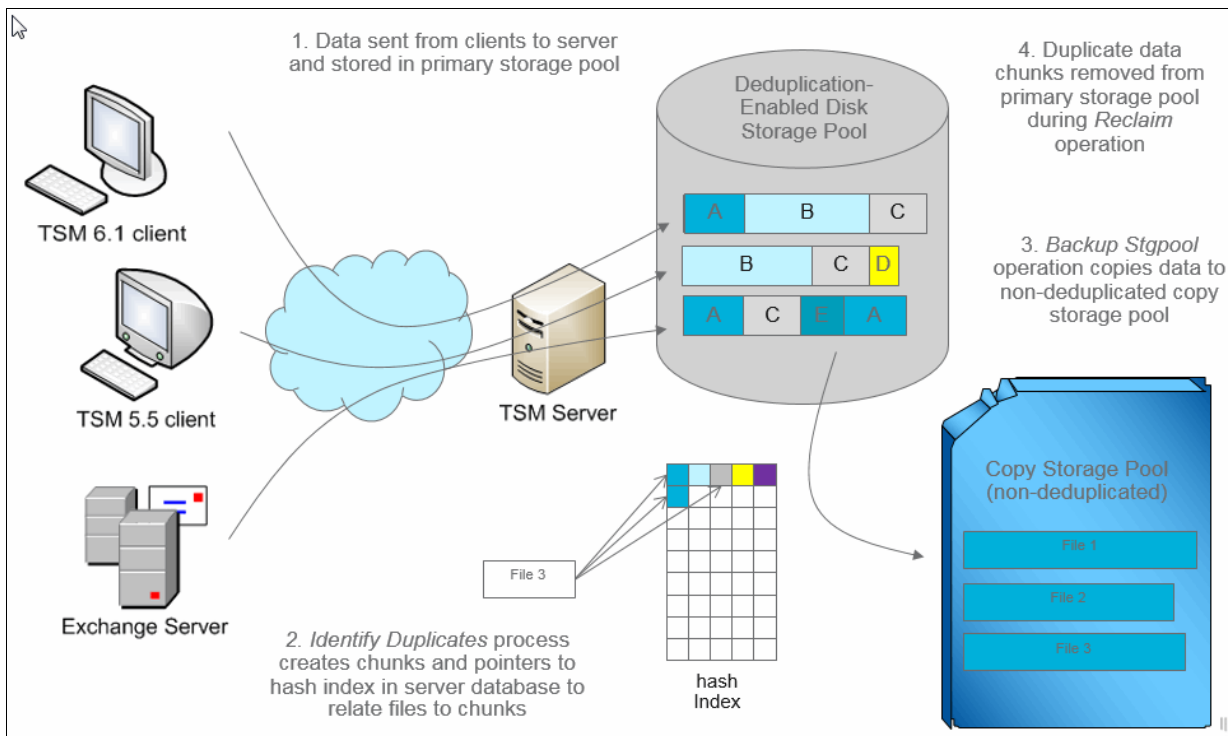


Figure 13-30 IBM Spectrum Protect server-side data deduplication

The Tivoli Storage Manager administrator can specify the data deduplication location (client or server) to use with the **DEDUP** parameter on the **REGISTER NODE** or **UPDATE NODE** server command.

With client-side data deduplication, you can accomplish the following tasks:

- ▶ Exclude specific files on a client from data deduplication.
- ▶ Enable a data deduplication cache that reduces network traffic between the client and the server. The cache contains extents that were sent to the server in previous incremental backup operations. Instead of querying the server for the existence of an extent, the client queries its cache.
- ▶ Enable both client-side data deduplication and compression to reduce the amount of data that is stored by the server. Each extent is compressed before it is sent to the server. The trade-off is between storage savings and the processing power that is required to compress client data. In general, if you compress and de-duplicate data on the client system, you are using approximately twice as much processing power as data deduplication alone.

Client-side data deduplication uses the following process:

1. The client creates extents. Extents are parts of files that are compared with other file extents to identify duplicates.
2. The client and server work together to identify duplicate extents. The client sends non-duplicate extents to the server.
3. Subsequent client data deduplication operations create extents. Some or all of those extents might match the extents that were created in previous data deduplication operations and sent to the server. Matching extents are not sent to the server again.

Client-side data deduplication provides several advantages:

- ▶ It can reduce the amount of data that is sent over the local area network (LAN).
- ▶ The processing power that is required to identify duplicate data is offloaded from the server to client nodes. Server-side data deduplication is always enabled for deduplication-enabled storage pools. However, files that are in the deduplication-enabled storage pools and that were deduplicated by the client do not require additional processing.
- ▶ The processing power that is required to remove duplicate data on the server is eliminated, allowing space savings on the server to occur immediately.

For further data reduction, you can enable client-side data deduplication and compression together. Each extent is compressed before it is sent to the server. Compression saves space, but it might increase the processing time on the client workstation.

With client-side data deduplication, the server does not have whole copies of client files until you back up the primary storage pools that contain client extents to a non-deduplicated copy storage pool (extents are parts of a file that are created during the data deduplication process). During storage pool backup to a non-deduplicated storage pool, client extents are reassembled into contiguous files.

By default, primary sequential-access storage pools that are set up for data deduplication must be backed up to non-deduplicated copy storage pools before they can be reclaimed and before duplicate data can be removed. The default ensures that the server has copies of whole files at all times, in either a primary storage pool or a copy storage pool.

For more information about IBM data deduplication solutions, see *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888.

In the DB2 on VersaStack deployment, we used IBM Spectrum Protect Node-Replication rather than working with a copy storage pool.

13.5.3 Node replication with automated failover

Node replication is the process of incrementally copying or replicating client node data from one IBM Spectrum Protect server to another IBM Spectrum Protect server for the purpose of disaster recovery.

The server from which client node data is replicated is called a source replication server. The server to which client node data is replicated is called a target replication server.

Node replication avoids the logistics and security exposure of physically moving tape media to a remote location. If a disaster occurs and the source replication server is unavailable, backup-archive clients of Tivoli Storage Manager can recover their data from the target replication server. If you cannot recover the source replication server, you can convert client nodes to non-replicating nodes for store operations on the target replication server.

Figure 13-31 shows the benefits of data replication for recovery.

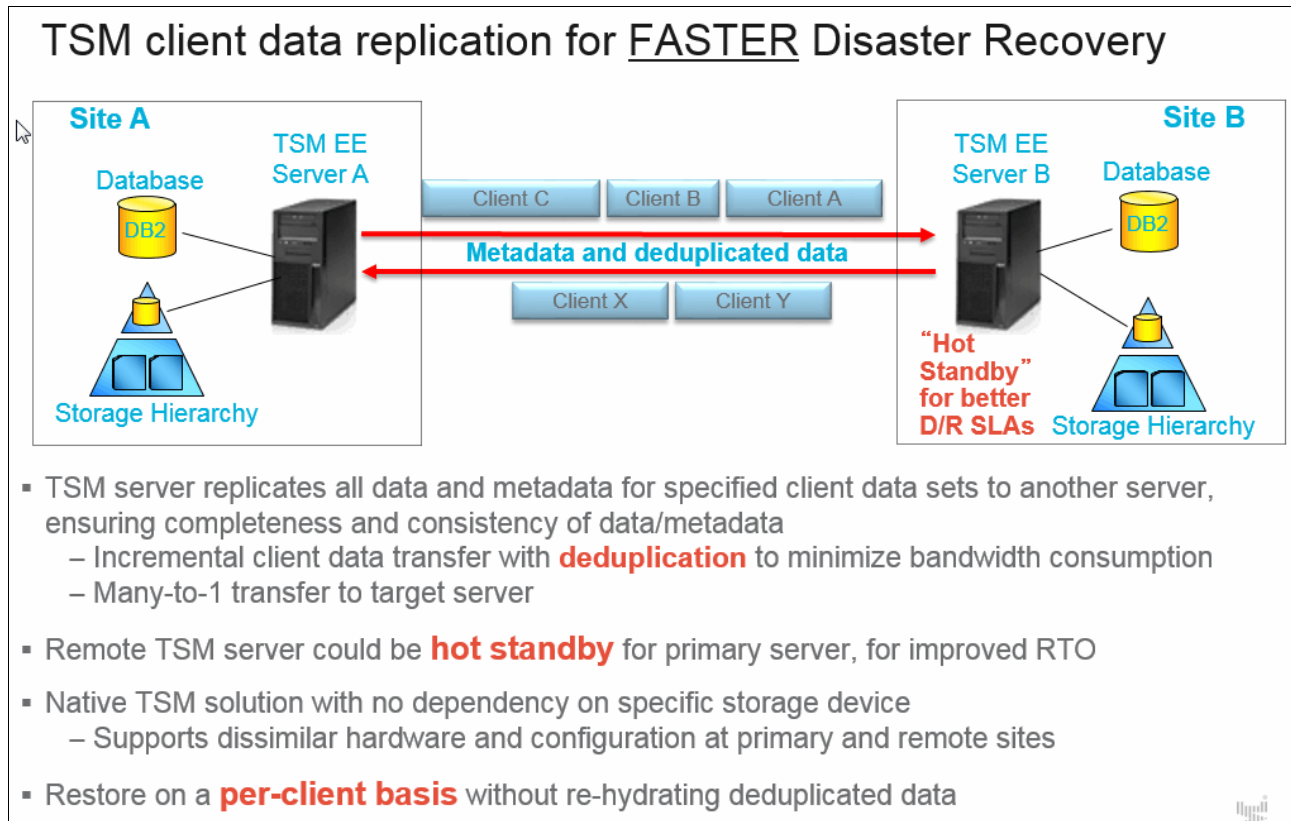


Figure 13-31 IBM Spectrum Protect Node-Replication

As of Version 7.1, IBM Spectrum Protect Node-Replication is enhanced to offer automated failover. When a backup/archive client or a data protection application starts, it attempts to open a session to its primary backup server. If this action fails, a connection to the secondary replication server is established, allowing for restores without requiring backup operator intervention.

Figure 13-32 shows the process of replication with automated failover.

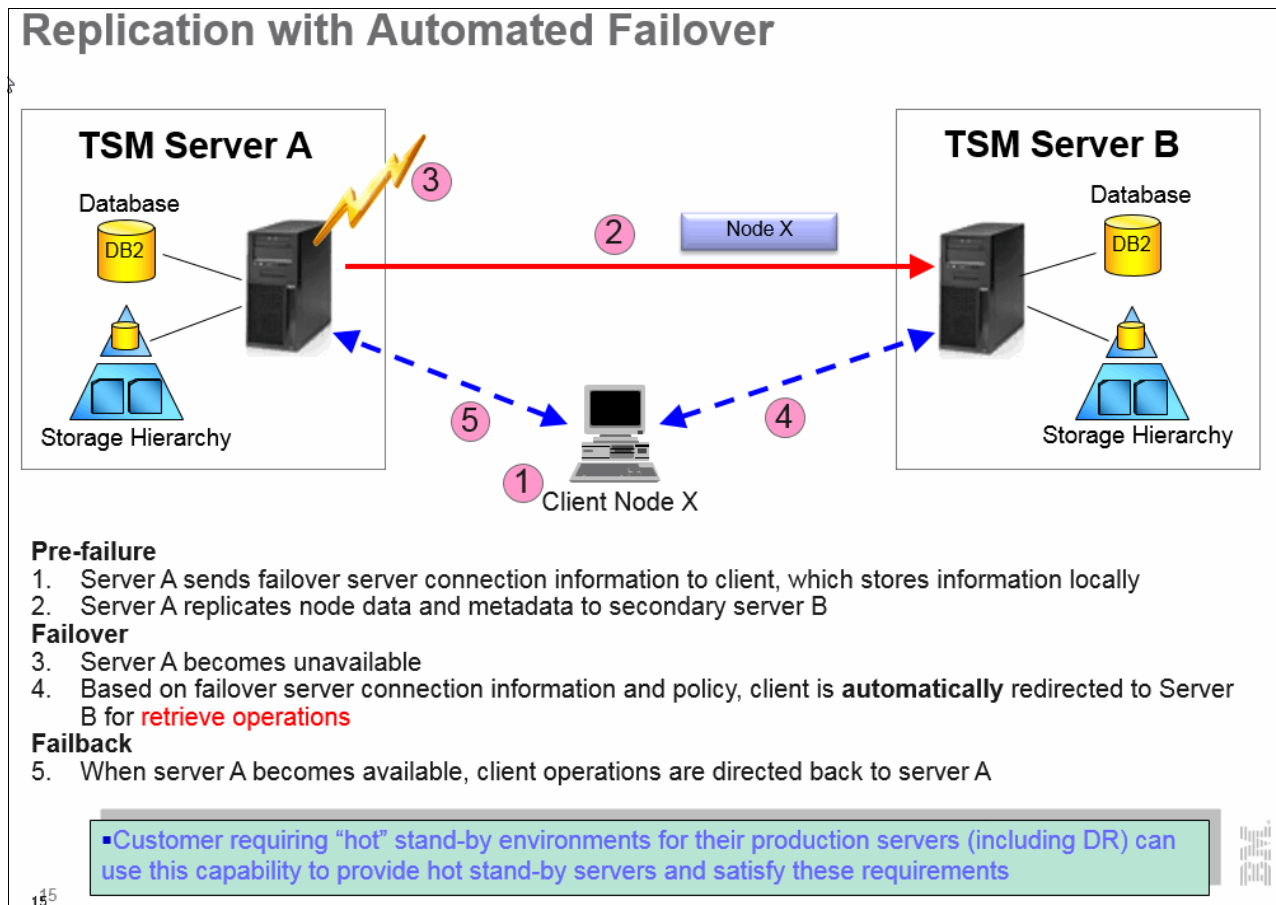


Figure 13-32 IBM Spectrum Protect Node-Replication automated failover

With IBM Spectrum Protect V7.1.1, Policy Driver Remote Replication is introduced. You can use this feature to have dissimilar versions or retention time on the source and target servers. A typical use case is to have a limited number of versions in a branch office for fast local restore with more versions in the central datacenter, or more versions on the primary production server and a limited longer term subset on the secondary server.

Figure 13-33 on page 399 shows Policy Driven Remote Replication.

Policy Driven Remote Replication – Dissimilar Policies

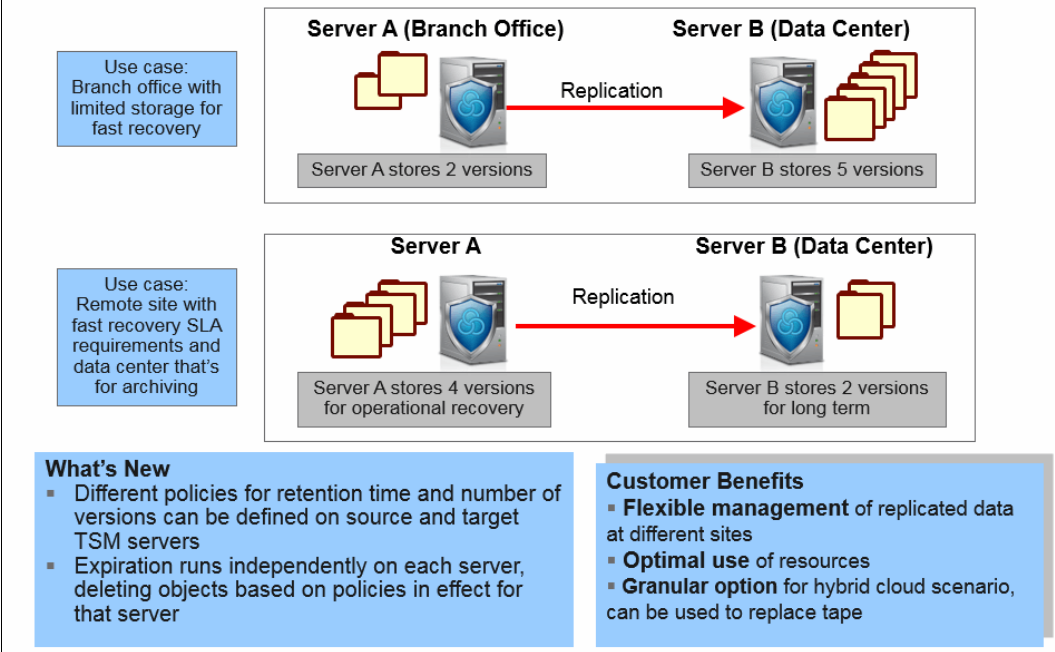


Figure 13-33 IBM Spectrum Protect Node-Replication Policy Driven Remote Replication - Dissimilar Policies

When using Node-Replication, you can also configure the IBM Spectrum Protect servers to automatically recover damaged data on the primary server by retrieving it from the secondary replication server, as shown in Figure 13-16 on page 377.

13.6 Monitoring and managing the IBM Spectrum Protect environment

Within the IBM Spectrum Protect product range, there is a dual approach to monitoring and managing the backup environment. With one approach, the backups can be centrally managed, monitored, and reported on, and with the other approach, the GUIs, backup and restore processes, and local reporting are integrated closely with the native environment that the user uses for the application by using, for example, the IBM Spectrum Protect Data Protection for SQL or Data Protection for VMware application modules.

This section describes the following topics:

- ▶ Data Protection for VMware monitoring and reporting through the GUI.
- ▶ How you can use the Operations Center as a central operational management dashboard.
- ▶ How you can have long-term statistical data and reports that are generated and automatically distributed through the Reporting and Monitoring Component.

13.6.1 Data Protection for VMware

The VMware administrator can use the Data Protection for VMware vSphere plug-in or web GUI to monitor the backup activities that are related to their VMware environment.

Figure 13-34 shows the Data Protection for VMware Recent Tasks window.

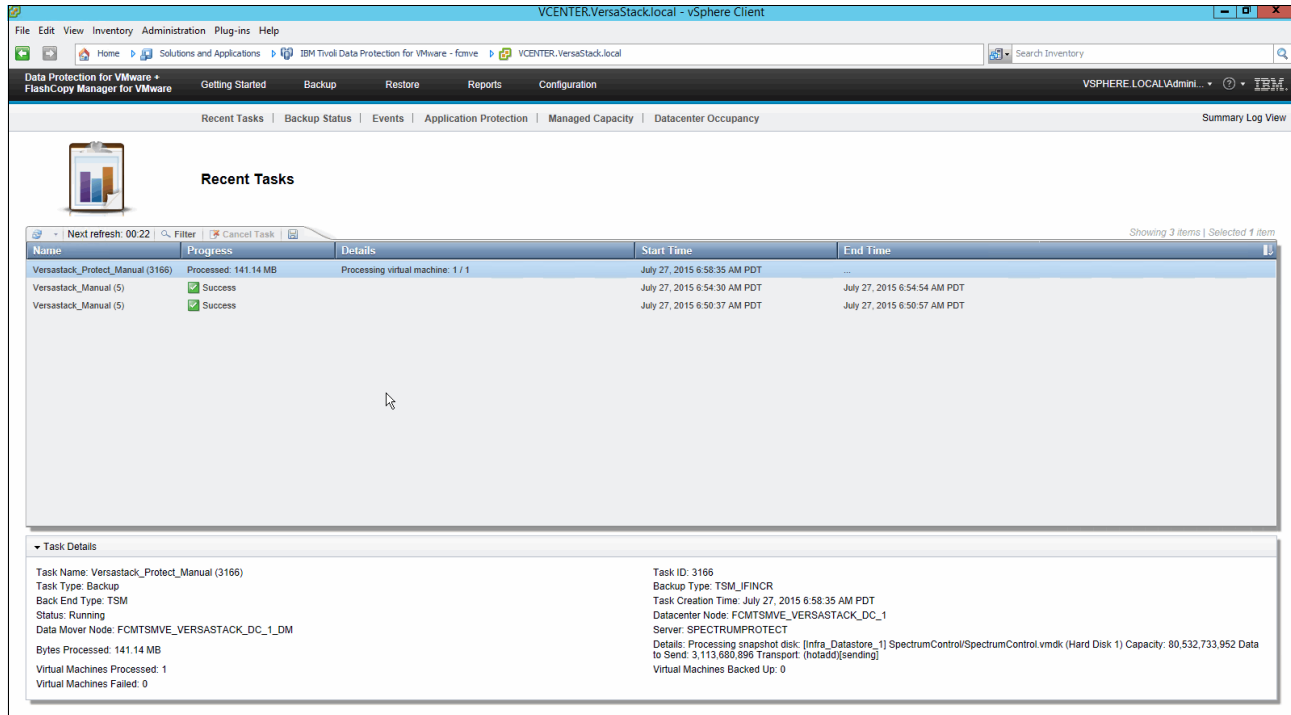


Figure 13-34 Data Protection for VMware Recent Tasks window

For some applications, in-guest backup clients and applications must be deployed as they are deployed in the DB2 on VersaStack environment. The plug-in offers a holistic view for the VMware administrator so that he can see the status of the agentless VM backups and retrieve information about the backups that are run for the specific VMs through his familiar VMware interface.

13.6.2 IBM Spectrum Protect Operations Center

The Operations Center provides web and mobile access to status information about the IBM Spectrum Protect environment. You can use the Operations Center to monitor multiple servers and to complete some administrative tasks. The Operations Center also provides web access to the IBM Spectrum Protect command line.

From the Operations Center, you can complete daily monitoring tasks to ensure that the IBM Spectrum Protect system is functioning correctly.

Figure 13-35 on page 401 shows the overview window of the Operations Center.

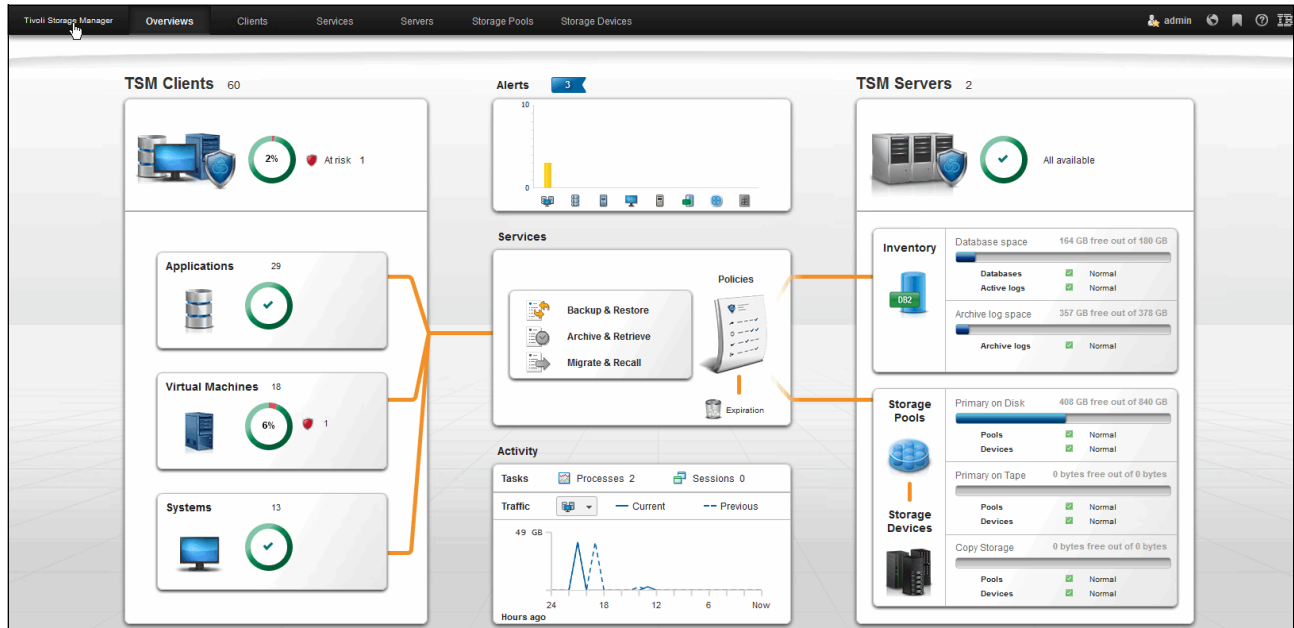


Figure 13-35 Operations Center - overview window

Tivoli Storage Manager Clients window

You can determine whether any clients are at risk of being unprotected because of failed or missed backups. You can open the Tivoli Storage Manager Clients window to view details for the following items:

- ▶ Applications: Groups applications that are protected by the IBM Spectrum Protect Data Protection modules
- ▶ VMs: Shows the VMs that are protected through the Data Protection for VMware module and the results of the latest backup run for these VMs
- ▶ Systems: Lists the physical or VMs that have a IBM Spectrum Protect Backup/Archive client installed.

Figure 13-36 shows an overview of the Tivoli Storage Manager clients from within the Operations Center.

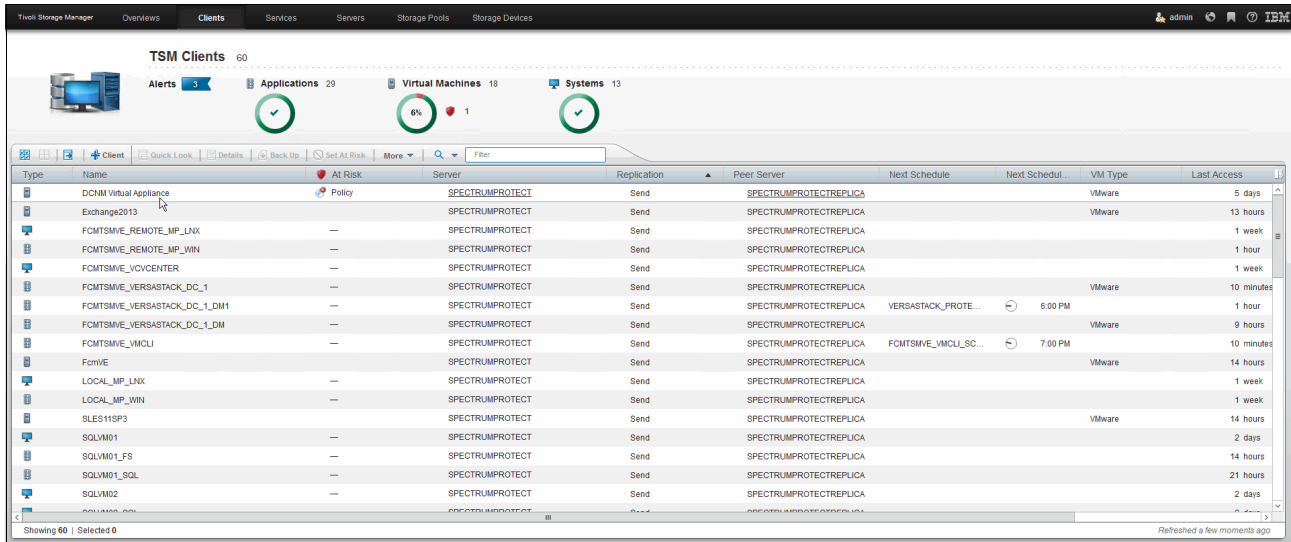


Figure 13-36 Operations Center - Tivoli Storage Manager Clients window

The Tivoli Storage Manager Clients window shows all the registered clients for both the source and replication IBM Spectrum Protect servers. Within each pane, the information that is displayed can be toggled and you can set advanced filters to display only the systems that you want to review, as shown in Figure 13-37.

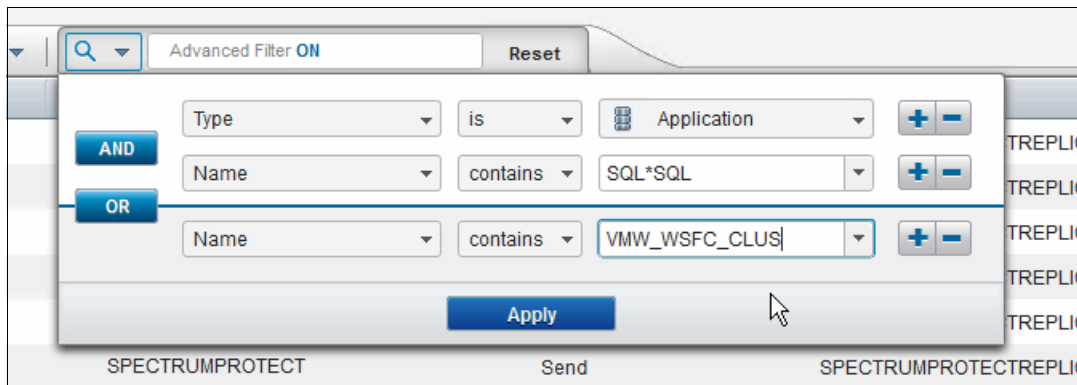


Figure 13-37 Operations Center - Advanced Filter

In this book, we define a filter to show both the VMs and the logical cluster nodes on the IBM Spectrum Protect Servers.

Backup schedules for physical systems or VMs that have a backup/archive client that is deployed can be triggered from within this section of the Operations Center.

Figure 13-38 on page 403 shows the SQLVM02 Operations Center client summary showing the activity for the last two weeks.




Figure 13-38 Operations Center Client Summary

Double-clicking a Tivoli Storage Manager client or selecting the **Client Details** shows the activities that are related to that client for the last two weeks.

From within this window, you can also change the client properties on the IBM Spectrum Protect server, start the remote web client through the Client Access function, and retrieve diagnostic information for the backups that are related to that client.

Alerts

Determine whether any client-related or server-related errors require attention.

Click the Alerts area to view details. Activity log messages are available on the Alerts window.

Activity

Verify that the amount of data that was recently sent to and from the Tivoli Storage Manager servers is within the expected range.

Tivoli Storage Manager Servers

Verify that the Tivoli Storage Manager servers that are managed by the Operations Center are available to provide data protection services to clients.

Click the Tivoli Storage Manager Servers area to view details and to go to more details for a specific server, for example:

- ▶ On the Servers window, select a server, and click **Details**.
- ▶ See the Summary, Active Tasks, and Completed Tasks tabs.

In the Active Tasks view, you can view or cancel the sessions that are in progress. You can also view activity log messages.

In the Completed Tasks view, you can view the sessions and processes that succeeded or failed. You can also view activity log messages.

From the Summary and Completed Tasks tabs, you can view information about the following processes:

- ▶ Database backups
- ▶ Scheduled server maintenance processes, such as reclamation, storage pool backups, and storage pool migrations

Inventory

If problems are indicated for the server database and associated logs, click **Inventory** to view details, for example:

- ▶ Check the amount of used and free space for the database, the active log, and the archive log.
- ▶ Verify that database backups are running as expected.

Storage Pools

If problems are indicated for primary or copy storage pools, click **Storage Pools** to view details.

For example, verify that the storage pools have enough free space.

Tip: If data deduplication is enabled, see the Completed Tasks view for the respective server to ensure that processes are completing successfully.

Storage Devices

If problems are indicated for devices, click **Storage Devices** to view details. Check for the following problems that can affect the status:

- ▶ For DISK device classes, volumes might be offline or have a read-only access state.
- ▶ For tape or shared FILE device classes, libraries, paths, or drives might be offline.
- ▶ For FILE device classes that are not shared, directories might be offline. Also, adequate free space might not be available for allocating scratch volumes.

Command Line

From the Operations Center command line, you can issue commands to manage Tivoli Storage Manager servers that are configured as hub or spoke servers.

The IBM Spectrum Protect Operations Center provides you with a management interface and a dashboard that holds up to 14 days of data that is related to the IBM Spectrum Protect environment. This short-term data is stored on the IBM Spectrum Protect Server that acts as the Operations Center Hub server.

Longer-term data is collected separately and stored in a Tivoli Monitoring for IBM Spectrum Protect data warehouse outside of the IBM Spectrum Protect servers' databases. This data warehouse is queried by the supplied Cognos Business Intelligence report creation tool to offer automated reports, historical trending, audit logs, and so on.

13.6.3 Reporting and monitoring for IBM Spectrum Protect

This section describes the reporting and monitoring components of IBM Spectrum Protect.

Tivoli Monitoring for IBM Spectrum Protect

Tivoli Monitoring for IBM Spectrum Protect brings together multiple components to provide Tivoli Storage Manager data collection, real-time monitoring of that data, and historical reports.

IBM Tivoli Monitoring acts as a monitoring application that provides workspaces for you to monitor real-time information. You can monitor the Tivoli Storage Manager server status, database size, agent status, client node status, scheduled events, server IDs, and so on, using the monitoring workspaces.

Tivoli Monitoring for IBM Spectrum Protect also provides reports that are based on the historical data that is retrieved. You can use the existing historical reports that are provided, or you can create your own custom reports.

Tivoli Monitoring for IBM Spectrum Protect consists of the following components:

- ▶ IBM DB2: Stores historical data that is obtained from Tivoli Storage Manager servers that are monitored by IBM Tivoli Monitoring.
- ▶ IBM Tivoli Monitoring: Consists of a number of components that accumulate and monitor historical data for reporting:
 - Tivoli Enterprise Portal Server
 - Tivoli Data Warehouse
 - Tivoli Enterprise Monitoring Server
 - Summarization Pruning agent
 - Warehouse Proxy agent
 - Tivoli Monitoring for IBM Spectrum Protect agent

The Tivoli Monitoring for IBM Spectrum Protect agent queries and formats data that is presented to you in the following ways:

- ▶ As workspaces from the Tivoli Enterprise Portal
- ▶ As reports that use the Tivoli Data Warehouse and the reporting portion of Tivoli Monitoring for IBM Spectrum Protect

Figure 13-39 on page 406 shows the IBM Spectrum Protect Servers overview.

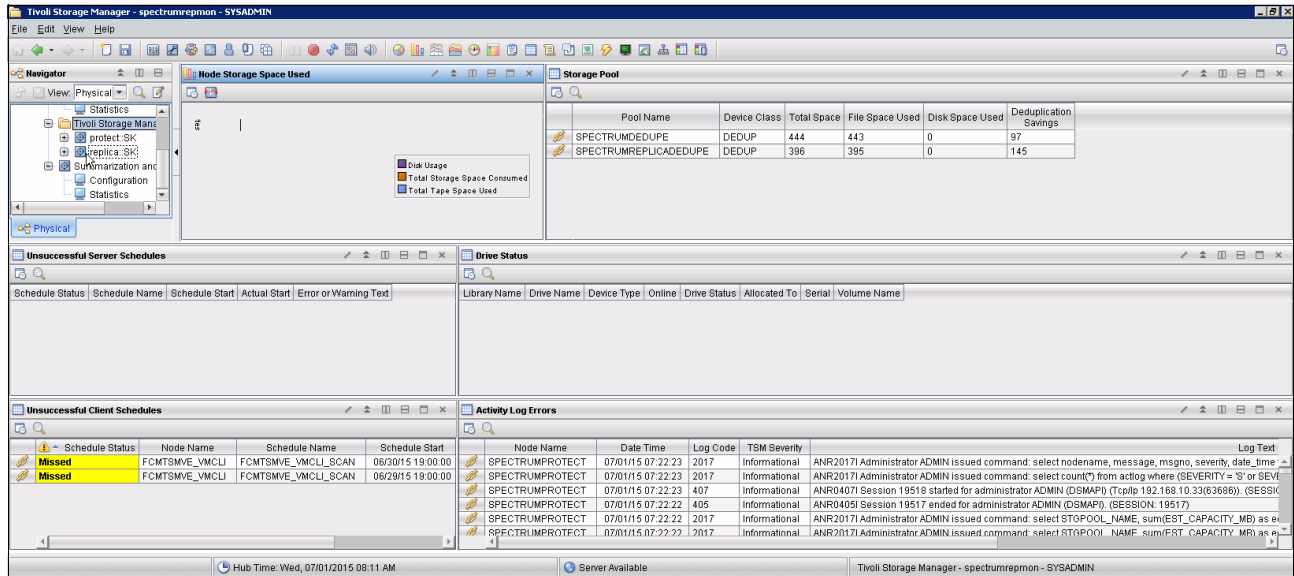


Figure 13-39 Tivoli Enterprise Portal - IBM Spectrum Protect Servers Overview

Tivoli Enterprise Portal for IBM Spectrum Protect

You can monitor your Tivoli Storage Manager server in real time by using the workspaces that are provided in the Tivoli Enterprise Portal. Client and server activities are monitored by the monitoring agent, and are displayed in workspace views.

When you open the Tivoli Enterprise Portal and go to the Tivoli Storage Manager view, a dashboard workspace displays commonly viewed information in a single location.

The dashboard workspace can be customized to suit your monitoring needs, but the default settings display the following information:

- ▶ Storage space that is used for each node that is defined on the server
- ▶ Storage pool summary details
- ▶ Unsuccessful client and server schedules, including all missed or failed schedules
- ▶ Client node activity for all nodes on the server
- ▶ Activity log errors, including all severe error messages

These workspaces are provided as part of the Tivoli Enterprise Portal:

- ▶ **Activity log:** This workspace provides information about activity log messages that are based on the parameters selected. The data can be used to generate aggregated reports that are grouped by server, and subgrouped by client. By default, only error messages are displayed. To display warning and informational messages, you can update the agent environment file to update the KSK_QUERYWARN and KSK_QUERYINF environment variables.
- ▶ **Activity summary:** This workspace provides summarized activity log information about virtual environments.
- ▶ **Agent log:** This workspace provides trace file information that is produced by the agent without having to enable tracing. It provides messages information, such as login successes and failures, and agent processes.
- ▶ **Availability:** This workspace provides the status and the performance of the agent that is running for each of the different workspaces that are listed under the Tivoli Storage Manager agent. It can help to identify problems with the gathering of historical data.

- ▶ **Client node storage:** The main workspace displays information about client node storage, disk, and tape usage data. This data can help you identify the clients that are using the most resources on the server. Disk and tape usage information is displayed in graph format.

The subworkspaces display data in a tabular format and a graph format. To display the subworkspaces, select the Client Node Storage workspace, right-click, select **Workspace**, and click the subworkspace that you want to view. Additional subworkspaces include:

- File space usage
 - Tape usage
 - Total storage space used
 - Storage pool media
- ▶ **Client missed files:** This workspace provides the status of missing files that are reported during client backups. It displays the client node name, the name of the server, the missing file name, and the full path to the missing file. This workspace can help to identify clients with many missing files.
 - ▶ **Client node status:** This workspace provides the date of the last successful backup, successful backup dates with warnings, and dates of any failed backups, for a client node name. You can click the chain-link icon for more details about each node. Click the green back arrow to return to the main workspace view.
 - ▶ **Database:** This workspace provides information about the status of database backups, including the last full backup and the last incremental backup. This information can be used to determine when all of the allocated database space is used. If all the allocated space is used, expansion operations must be taken to ensure that the database continues to operate. Because a Tivoli Storage Manager server processes client requests for backup-archive operations, the Tivoli Storage Manager database is updated with current and historical types of data. Total capacity and total space used data is displayed in a bar chart format, and database details such as percentage of space that is used and total space that is used is displayed in a tabular format.
 - ▶ **Drives:** This workspace provides status information about the drives, including drive name, library name, device type, drive status, such as loaded or empty, the volume name, and whether the drive is online. An additional subworkspace is Drives drill down.
 - ▶ **Libraries:** This workspace provides status information about libraries, such as the library name, type, if it is shared or not, LAN-free, auto label, number of available scratch volumes, whether the path is online, and the serial number.
 - ▶ **Node activity:** This workspace provides activity metrics for a specific node over a 24-hour period, for example, activity metrics include the amount of data that is backed up, the number of objects that are inspected, and the number of processed objects.

The subworkspaces display data in a tabular format and a graph format. To display the subworkspaces, select the Node Activity workspace, right-click and select **Workspace**, and click the subworkspace that you want to view. Additional subworkspaces include:

- Client activity backup
- Client activity restore
- Client activity archive
- Client activity retrieve
- NAS activity
- Server activity DB backup
- Server activity file expiration

- ▶ **Occupancy:** This workspace provides tabular and graphical information about where backup and archive data is stored on the server and how much data is stored, for example, number of files, physical MB, logical MB, and by node name. Click the chain-link icon to display more details. Bar graph details depict the space that is used, in MB, by the storage pool and the number of files that are used by the storage pool.

The subworkspace displays data in a tabular format and a graph format. To display the subworkspaces, select the Occupancy workspace, right-click, select **Workspace**, and click the subworkspace that you want to view. An additional subworkspace is Drives drill down.

- ▶ **Processor Value Unit (PVU) details:** This workspace provides PVU details by product, and PVU details by node. It includes information such as node name, product, license name, last used date, try buy, release, and level. If the Tivoli Storage Manager server is not a Version 6.3 server or later, the workspace is blank.
- ▶ **Replication details:** This workspace provides byte by byte replication details. It describes all of the replication details such as node name, file space ID, version, start and end times, status, complete status, incomplete reason, estimated percentage of completion, estimated time remaining, and estimated time to completion.
- ▶ **Replication status:** This workspace provides the replication status for a node without all of the details that the replication details workspace provides. It displays node name, server, file space type, name and ID, target server, and the number of files on the source and target server.
- ▶ **Schedule:** This workspace provides details about client and server schedules. You can group the data by node name, schedule name, or status to identify any potential problems. It displays information such as schedule name, node name, server name, scheduled start, actual start, and the status of the schedule, which can be success, missed, or failed, including any error or warning text.
- ▶ **Sessions:** This workspace provides a view of all the client sessions that are running on the specified server. This workspace is useful for determining which clients are connected to the Tivoli Storage Manager server and how much data was sent or received. The workspace also shows tape mount information that indicates library and tape usage.
- ▶ **Storage pool:** This workspace provides you with detailed information about your storage pools. Tivoli Storage Manager can contain multiple storage pools. These storage pools define the methods and resources that are used to store the data that is backed up or archived to the Tivoli Storage Manager server. The data that is displayed in this workspace includes storage pool names, server name, device classes, total space, utilized space, total volumes used, percentage of space used, disk space used, and data deduplication savings. It also displays a graph with the total space, total usage, and total volumes used.
- ▶ **Server:** This workspace provides the operational status of the Tivoli Storage Manager server. These operations are measured in megabytes per operation. After the operational status is reported, the values are reset to zero. The numbers that are reported for each operation are not cumulative over time.

You can view the following activities or status:

- Length of time it takes activities to complete.
- Any problems that occur after activities complete.
- The status of server-only activities.

The data that is displayed includes information such as server name, disk storage pool space, tape usage count, current database size, information for client operations from a previous day, object count reclamation by byte and duration, migration by byte and duration, and backup by byte and duration.

Bar graphs are also provided to display server operation duration and server operation byte counts.

Figure 13-40 shows the Server Operation Duration and Byte Count as seen in the Tivoli Enterprise Portal.

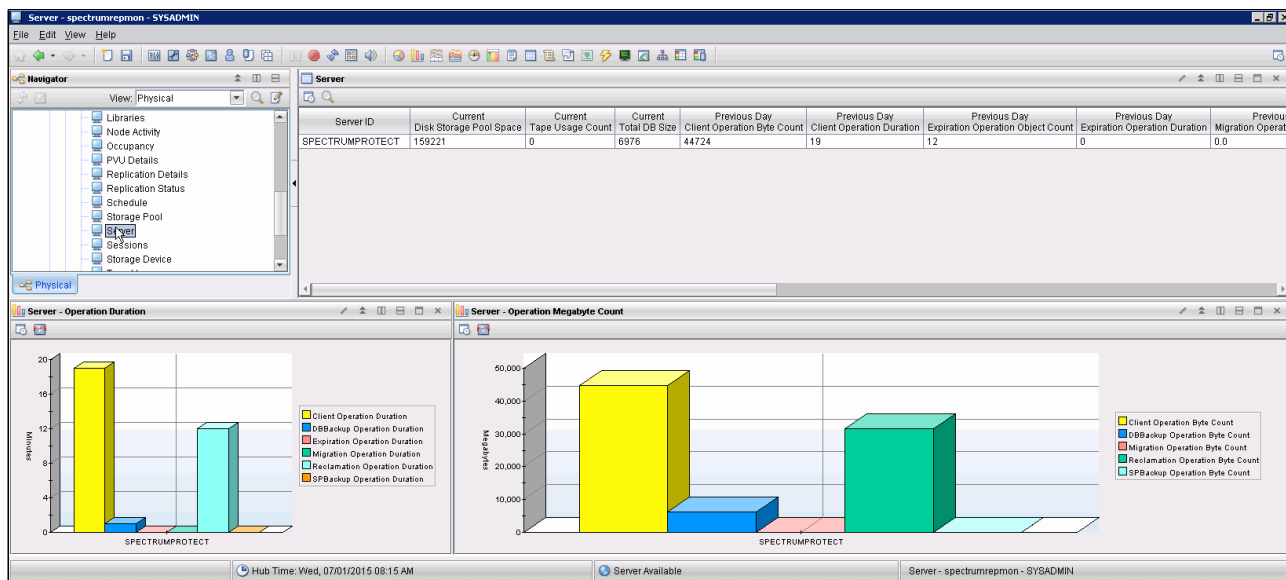


Figure 13-40 Tivoli Enterprise Portal Servers Overview

- ▶ **Storage device:** This workspace provides you with the read and write error status of the storage devices. This status helps you identify possible problems with any of your storage devices. Bar chart graphs also display read and write error count.
- ▶ **Tape usage:** This workspace provides you with tape usage data for each client.
- ▶ **Tape volume:** This workspace provides the status of all tape storage devices. This information can help you identify any storage devices that are near full capacity.

Daily monitoring with Tivoli Enterprise Portal

You can use Tivoli Monitoring for IBM Spectrum Protect to monitor many daily operations to ensure that your system is running in good condition. To do this task, complete the following steps:

1. Start Tivoli Enterprise Portal, log on with your sysadmin ID and password, and go to Tivoli Storage Manager.

Many of the items that you can check daily are displayed in the dashboard view when it opens. The dashboard displays a grouping of commonly viewed items in a single view. Examine items and look for any values that might indicate a potential problem:

- **Node storage space used:** Check this graph for disk, storage, and tape space used.
- **Storage Pool:** Click the chain link icon to drill down for additional details.
- **Unsuccessful server schedules:** Review this table for any missed or failed server schedules. Click the chain link icon for additional details.
- **Unsuccessful client schedules:** Review this table for any missed or failed client schedules. Click the chain link icon for additional details.
- **Drive Status:** Review this table to ensure that all drives are online.
- **Activity log errors:** Review this table to check for error, warning, and severe messages. Click the chain link icon for additional details.

2. In the navigator pane, select the Database workspace. Examine the Percent Space Used value to ensure that the database file system has enough available space. You can also check the Backup Status field to ensure that the database backups completed successfully.
3. Go to the Storage Pool workspace and review the total space that is used to ensure that there is enough space available to manage the anticipated workload.
4. Go to the Activity Log workspace and review the information in the table for any error messages that might indicate a problem that must be resolved.
5. Go to the Drives workspace and check to ensure that all drives are online.
6. Go to the Libraries workspace and check to ensure that the path to the library is online. Click the chain-link icon for additional details.
7. Go to the Tape Volume workspace to view the status and identify devices that are near full.
8. Go to the Server or Activity Log workspace to review the operational status, such as what activities are taking too much time to complete, the status of activities, messages about the activities, and other details that help to identify potential problems.

Cognos reports

IBM Cognos 8 Business Intelligence is an integrated business intelligence suite that is provided as part of Tivoli Common Reporting. You can use Cognos to view and create business reports, analyze data, and monitor events and metrics.

Cognos reports include status and trending data about your Tivoli Storage Manager server and clients.

Cognos reports are available in HTML, PDF, Microsoft Excel, XML, and CSV (delimited text) formats.

Figure 13-41 shows an overview of the default available Status and Trending reports for IBM Spectrum Protect.

Status reports	Trending reports
Client activity status	Client activity success rate
Client backup currency	Client schedule success rate
Client backup status	Client storage usage trends
Client schedule status	Disk utilization trends
Client storage pool usage summary	Node replication growth
Client storage summary and details	Server database growth trend
Current client occupancy summary	Server storage growth trends
Current storage pool summary	Server throughput trends
Highest storage space usage	
Node replication details	
Node replication summary	
Server activity log details	
Server schedule status	
Storage pool deduplication savings	
VE activity status	
VE backup type summary	
VE current occupancy summary	

Figure 13-41 IBM Spectrum Protect Cognos Reports

A detailed description for these reports can be found at the following website:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itm.srv.doc/r_rpt_cognos_rpts.html?lang=en-us#r_rpt_cognos_rpts_crpts

You can use Report Studio to create your own customized Cognos reports.

Report Studio is a product for creating Cognos reports that analyze corporate data according to specific information needs. In Report Studio, you can accomplish the following tasks:

- ▶ Create a Cognos report by developing a query to fetch data from the WAREHOUSE database.
- ▶ Modify an existing Cognos report to change the appearance.
- ▶ View data from a Cognos report to test your new query.

For more information about creating customized reports, see the following website:

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Creating%20Customized%20Reports>



General performance

This chapter describes some of the tools that are available to analyze VersaStack performance.

14.1 IBM Easy Tier

Easy Tier is a performance function that automatically and nondisruptively migrates frequently accessed data from magnetic media to solid-state drives (SSDs). The most frequently accessed data is stored on the fastest storage tier, and the overall performance is improved.

The first generation of Easy Tier introduced automated storage performance management by efficiently boosting enterprise-class performance with SSDs, and automating storage tiering from enterprise-class drives to SSDs. These changes optimized flash deployments with minimal costs. Easy Tier also introduced dynamic volume relocation and dynamic extent pool merge.

The third generation of Easy Tier introduces further enhancements that provide automated storage performance and storage economics management across all three drive tiers (flash, enterprise, and nearline storage tiers). You can use it to consolidate and manage efficiently more workloads on a single IBM Storwize V7000 Gen2 storage system. It also introduces support for storage pool balancing in homogeneous pools. It is based on performance, not capacity.

IBM has a tool to analyze the movements of extents by EasyTier that is called the IBM Storage Tier Advisor Tool, which can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S4000935>

Using the tool with the dpa_heat file that is generated on our example Storwize V7000 storage system shows which volumes have hot data. VDisk 4 was receiving much I/O, so EasyTier has migrated extents onto the SSD tier. The distribution of extents can be shown by running **1svdiskextent**.

Figure 14-1 shows the Volume Heat Distribution by using the STAT tool. The hot data is in red, warm data in orange, and cold in blue.



Figure 14-1 The Volume Heat Distribution that is found by using the STAT tool

14.2 Autotier

Spectrum Control features the Analyze Tiering wizard that can tier volumes automatically or based on the criteria that you set in your tiering policies. For example, you can tier volumes based on the volume workload, on file usage, or both. Depending on the conditions that are set in the tiering policy, recommendations are generated. For example, you can reduce storage costs by moving volumes with low workloads to lower or less expensive tiers. You can also improve performance and use storage more efficiently by moving volumes with heavy workloads to the tiers that best meet their workload requirements.

Volumes can be moved to tiered storage pools on the same storage virtualizer, but volumes cannot be moved from one storage virtualizer pool to another storage virtualizer pool.

A customer can select the resources that they want to analyze. The source storage pools that are related to the resources that you selected are analyzed to determine whether they meet the workload requirements of the volumes. If the workload requirements of the volume in its current tier are not met, the volume is a candidate for relocation. You can perform the following actions:

- ▶ Specify the target storage pools for the volumes.
- ▶ Include or exclude volumes in mirrored volume relationships from the analysis.
- ▶ Optionally, provide more information about storage pools on back-end storage systems. Tivoli Storage Productivity Center might require more information to estimate the workload capability of the source and target storage pools.

To ensure that the performance of the target pools is not degraded when volumes are added, you specify a maximum utilization percentage for the pools. The performance data that is collected on the previous day is used to estimate the average daily utilization of the physical resources, such as controllers, nodes, and disks, that are associated with a pool. The physical resources that are associated with a pool vary depending on the type of storage system.

Note: In our example setup, we put our SSDs in the control enclosure. This is a preferred practice because of our SAS topology; the Storwize V7000 SPCve chip has 16 PHYs, eight of which go to the internal SAS expanders, and the expansion chains receive four each. This means that placing our SSDs on the control enclosure allows us to receive the maximum bandwidth benefit.

To demonstrate the increased performance on SSDs and autotiering, use the HammerDB tool and Spectrum Control to measure the performance and complete the following steps:

1. Place the sql_rdm_data VDisk on Enterprise SAS-only mdiskgrps. Then, create a 250 GB database on this volume.
2. After the database creation finishes, run the HammerDB I/O tool. Create 21 virtual users, with a user delay and repeat delay of 500 ms.

Figure 14-2 shows the transaction counter of HammerDB.




Figure 14-2 The transaction counter of HammerDB while creating 21 virtual users on 10 K SAS drives

3. The transactions max out at 594570. Then, migrate the sql_rdm_data volume to SSD storage by using Spectrum Control. To do so, go to the sql_rdm_data volume in Spectrum Control, right-click it, and select **Transform Storage**, which starts a wizard. In the wizard, select the mdiskgrp with only SSDs (in our example, mdiskgrp2), as shown in Figure 14-3 and Figure 14-4 on page 417.

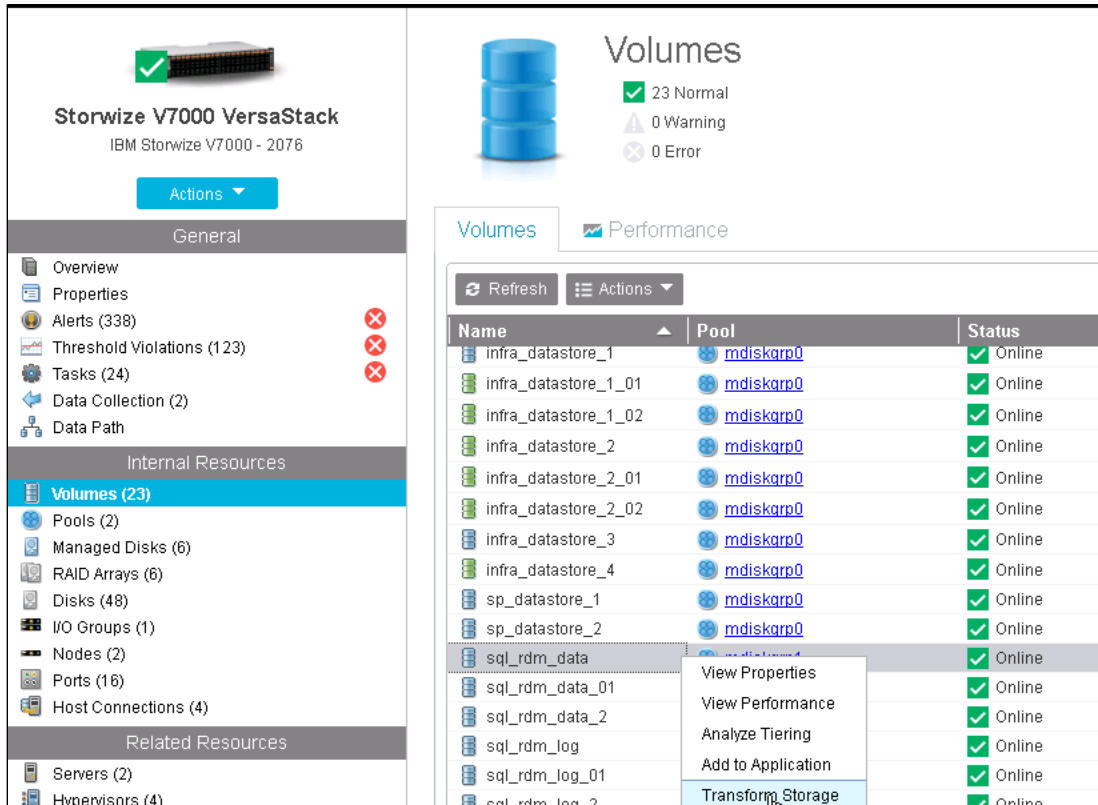


Figure 14-3 The Storwize V7000 volumes window on Spectrum Control

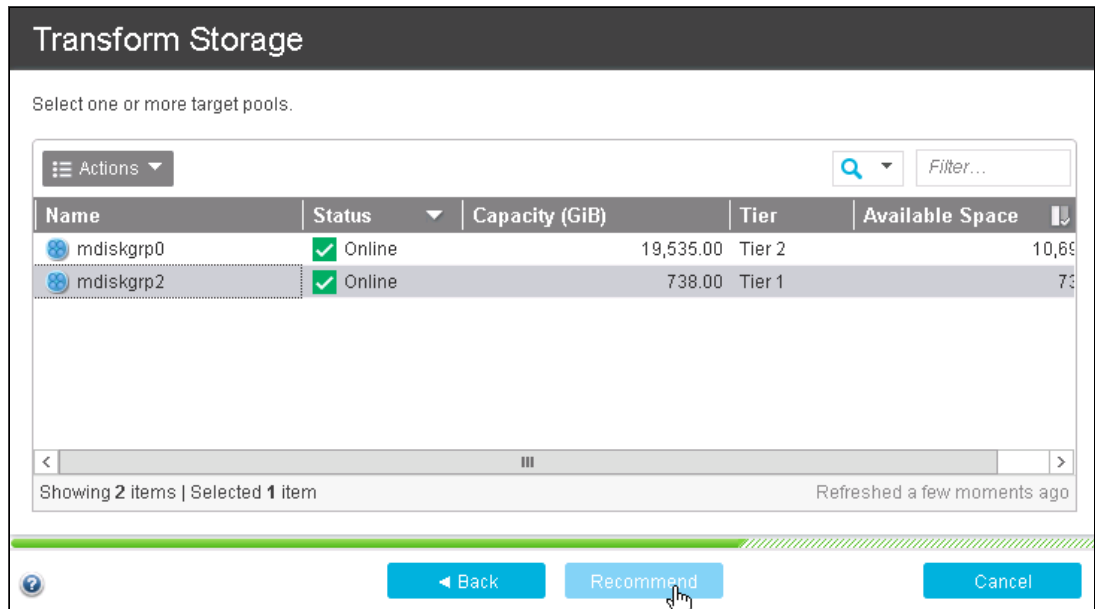


Figure 14-4 Using Spectrum Control to move the volume to a different tier

4. Rerun HammerDB by using the same input as before, as shown in Figure 14-5.




Figure 14-5 The transaction counter of HammerDB while creating 21 virtual users on SSDs

You can see that there is an increase to a maximum transactions per minute (TPM) of 675192.

You can use the automation layer in Spectrum Control to perform the autotiering analysis at scheduled intervals and have the volumes uptiered or downtiered automatically depending on the actual, historical, or expected workload.

14.3 General performance metrics

The HammerDB tool shows an increase in performance when migrating the SQL data volumes from SAS-based disks to SSD-based disks.

Gauging system performance by using a tool such as HammerDB is an intensive process that requires multiple reruns. Moreover, it is difficult to evaluate the results in cases where the general performance and capabilities of the environment supersede the load being put on the system by the benchmarking tool.

For the VersaStack environment, three components determine the general performance (abstracting the impact that is introduced by the OS and hypervisors):

- ▶ Computing blade (B200 M4)
- ▶ I/O backplane (VIC 1340)
- ▶ Storage system (Storwize V7000 storage system)

14.3.1 B200 M4

The CPU performance of the host on which the SQL virtual machines (VMs) are running also determines the processing performance of the database. The results are shown in Figure 14-6.

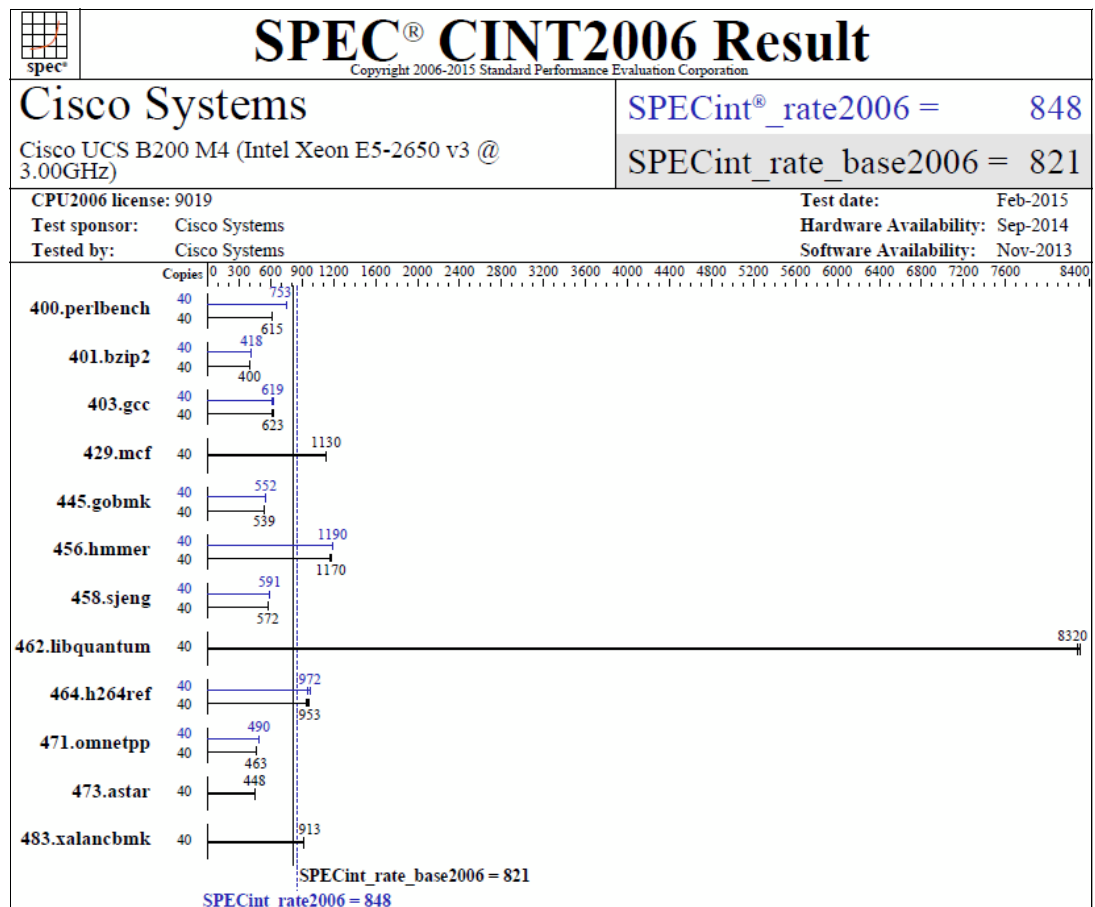


Figure 14-6 Cisco UCS B200 SPEC CINT2006 results

The B200 M4 that we use in our example SQL on VersaStack environment has a SPECINT rate of 848, as shown in Figure 14-6 on page 418.

14.3.2 VIC 1340

The main component determining the I/O backplane performance capability is the Cisco UCS Virtual Interface Card 1340. The VIC 1340 has the following features:

- ▶ Sixteen PCIe Gen3 interfaces.
- ▶ Two 40-Gbps Unified I/O ports or two sets of four 10-Gbps Unified I/O ports.
- ▶ It delivers 80 Gbps to the server.
- ▶ It helps reduce the total cost of operation (TCO) by consolidating the overall number of NICs, HBAs, cables, and switches. LAN and SAN traffic runs over the same mezzanine card and fabric.
- ▶ It adapts to either 10-Gbps or 40-Gbps fabric connections.
- ▶ It has more than 900,000 I/O operations per second (IOPS).

For more information about the VIC 1340, see the following website:

<http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1340/index.html>

14.3.3 Storwize V7000 storage system

Workload simulation and performance metrics for the Storwize V7000 Gen 2 storage system can be modeled by using the IntelliMagic Disk Magic tool.

The IBM Disk Magic for Windows modeling tool helps estimate IBM disk subsystem performance. The IBM disk controllers that are supported are IBM XIV, IBM DS8000, IBM DS6000™, IBM DS5000, IBM DS4000, IBM SAN Volume Controller, IBM Storwize V3500, IBM Storwize V3700, IBM Storwize V5000, IBM Storwize V7000, and IBM Storwize V7000U.

It is beyond the intended scope of this book to go into details of Disk Magic. There is a comprehensive amount of information that is available at the following websites:

- ▶ For IBM employees:
<https://ibm.biz/BdX7ca>
- ▶ For IBM Business Partners (you need your IBM ID to sign in):
<https://ibm.biz/BdX7cb>



General validation

Performing validation testing is important for quality control and to demonstrate that the product performs as expected. This chapter describes the validation testing that we performed for our example VersaStack solution.

15.1 Validation scenarios

These are the scenarios that we validated on our example VersaStack solution:

- ▶ Storwize V7000 storage system:
 - Unexpected Fibre Channel cable failure
 - Unexpected node failure
- ▶ Cisco Nexus Switches: vPC peer switch failure
- ▶ Cisco UCS Service Profile: Service profile migration

15.2 Storwize V7000 failover validation

The pair of nodes within a single Storwize V7000 enclosure is known as an *I/O group*.

When an application server processes I/O to a volume, it can access the volume with either of the nodes in the I/O group. When you create a volume, you can specify a preferred node. Many of the multipathing driver implementations that the system supports use this information to direct I/O to the preferred node. The other node in the I/O group is used only if the preferred node is not accessible.

If you do not specify a preferred node for a volume, the system selects the node in the I/O group that has the fewest volumes to be the preferred node.

An I/O group consists of two nodes. When a write operation is performed to a volume, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group. After the data is protected on the partner node, the write operation to the host application is completed. The data is physically written to disk later.

Read I/O is processed by referencing the cache in the node that receives the I/O. If the data is not found, it is read from the disk into the cache. The read cache can provide better performance if the same node is chosen to service I/O for a particular volume.

I/O traffic for a particular volume is, at any one time, managed exclusively by the nodes in a single I/O group. Thus, although a clustered system can have multiple nodes within it, the nodes manage I/O in independent pairs, which means that the I/O capability of the Storwize V7000 storage system scales well because additional throughput can be obtained by adding additional I/O groups.

When a node fails within an I/O group, the other node in the I/O group assumes the I/O responsibilities of the failed node. Data loss during a node failure is prevented by mirroring the I/O read and write data cache between the two nodes in an I/O group.

If only one node is assigned to an I/O group or if a node fails in an I/O group, the cache is flushed to the disk and then goes into write-through mode. Therefore, any writes for the volumes that are assigned to this I/O group are not cached; they are sent directly to the storage device. If both nodes in an I/O group go offline, the volumes that are assigned to the I/O group cannot be accessed.

15.2.1 Unexpected Fibre Channel cable failure

Removing the Fibre Channel (FC) cables from one node in the Storwize V7000 storage system causes all the I/O traffic to go through the Host Interface Card (HIC) on the other node, but I/O continues and both nodes are still used.

This scenario can be used as a good example of the redundancy of the Storwize V7000 storage system and to show how the Storwize V7000 storage system and Spectrum Control handle errors.

Example 15-1 shows the output of the `lsportfc` command, where you can see that all eight FC ports are active.

Example 15-1 List the FC ports by running lsportfc

```
[09:30:26] mcr-v7000-canister-02:~ # lsportfc
id fc_io_port_id port_id type      port_speed node_id node_name WWPN
nportid status      attachment cluster_use  adapter_location
adapter_port_id
0 1                1      fc      8Gb      4      node1    500507680B214FF4
0C0160 active          switch  local_partner 2      1
1 2                2      fc      8Gb      4      node1    500507680B224FF4
0C0020 active          switch  local_partner 2      2
2 3                3      fc      8Gb      4      node1    500507680B234FF4
A90160 active          switch  local_partner 2      3
3 4                4      fc      8Gb      4      node1    500507680B244FF4
A90020 active          switch  local_partner 2      4
4 5                4      ethernet N/A      4      node1    500507680B314FF4
000000 inactive_unconfigured none  local_partner 3      1
5 6                5      ethernet N/A      4      node1    500507680B324FF4
000000 inactive_unconfigured none  local_partner 3      2
6 7                6      ethernet N/A      4      node1    500507680B334FF4
000000 inactive_unconfigured none  local_partner 3      3
7 8                7      ethernet N/A      4      node1    500507680B344FF4
000000 inactive_unconfigured none  local_partner 3      4
14 1               1      fc      8Gb      2      node2    500507680B214FF5
0C0000 active          switch  local_partner 2      1
15 2               2      fc      8Gb      2      node2    500507680B224FF5
0C0040 active          switch  local_partner 2      2
16 3               3      fc      8Gb      2      node2    500507680B234FF5
A90000 active          switch  local_partner 2      3
17 4               4      fc      8Gb      2      node2    500507680B244FF5
A90040 active          switch  local_partner 2      4
18 5               4      ethernet N/A      2      node2    500507680B314FF5
000000 inactive_unconfigured none  local_partner 3      1
19 6               5      ethernet N/A      2      node2    500507680B324FF5
000000 inactive_unconfigured none  local_partner 3      2
20 7               6      ethernet N/A      2      node2    500507680B334FF5
000000 inactive_unconfigured none  local_partner 3      3
21 8               7      ethernet N/A      2      node2    500507680B344FF5
000000 inactive_unconfigured none  local_partner 3      4
```

To simulate this validation scenario, complete the following steps:

1. Remove the four FC cables from node 2 (the control node at this time). This action creates an error message on the Storwize V7000 CLI and GUI. Accessing the event menu by using the GUI shows more information.
2. To access the event log, click the **Events** tab, as shown in Figure 15-1.



Figure 15-1 The Events tab

3. Figure 15-2 shows two errors inside the event log. A Directed Maintenance Procedure (DMP) can be run by clicking the event in question and then clicking **Run Fix**. Click **Run Fix** for the top error to start a DMP for that error.

Error Code	Last Time Stamp	Status	Description	Object Type	Object ID	Object Name
1061	6/26/15 10:17:17 AM	Alert	Fibre Channel ports not operational	node	2	node2
1450	6/26/15 10:17:17 AM	Alert	Fibre Channel I/O ports not operational	node	2	node2

Figure 15-2 Event log in the GUI

You are asked if the change is on purpose and, if not, what you want like to do to fix it.

Figure 15-3 on page 425 shows the window that explains the error. In this case, four FC ports are inactive.

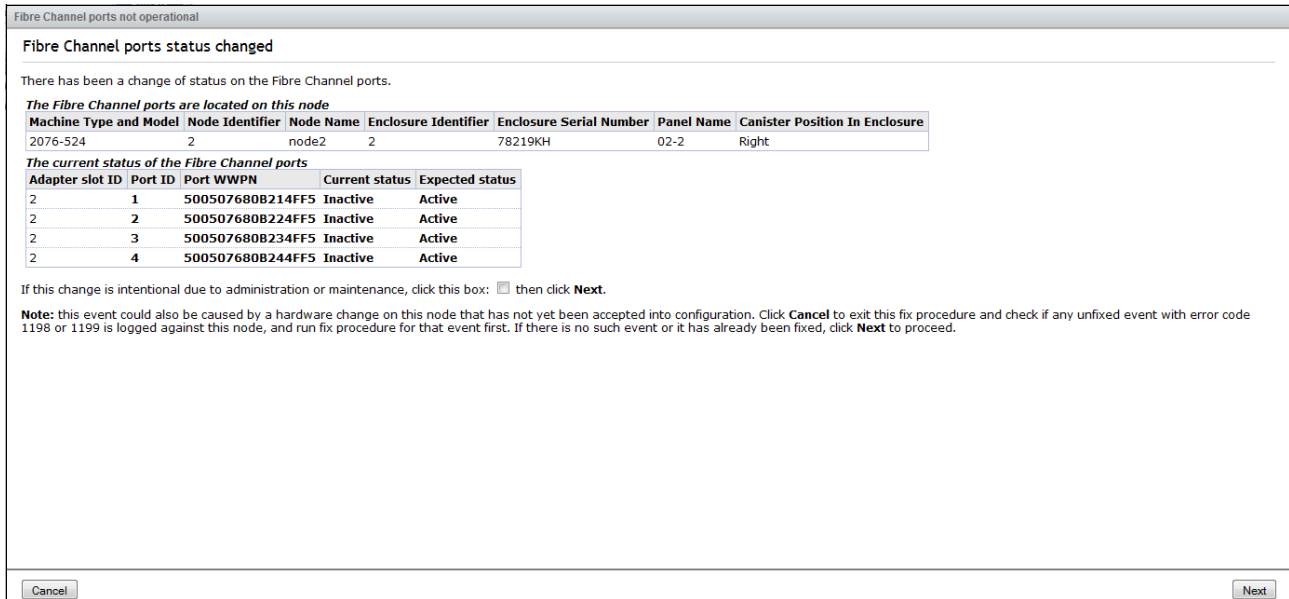


Figure 15-3 DMP showing the four inactive FC ports

4. Click **Next**. The DMP shows you possible ways to fix the issue.

Figure 15-4 shows how the DMP directs you to fix the error by checking the FC connections.

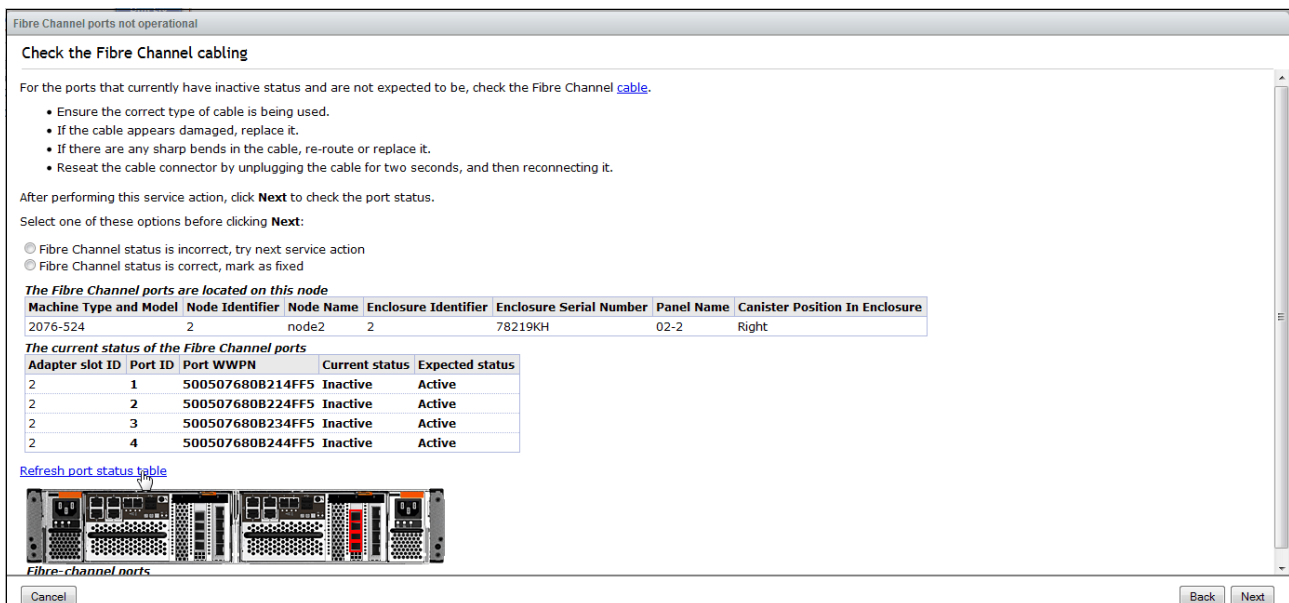


Figure 15-4 DMP prompt to check the cables

- If you plug in the FC cables that were removed from node 2 and refresh this window, you see the status of the ports go to Active and the event is marked as fixed.

Figure 15-5 shows that the problem is solved and the event is marked as fixed.



Figure 15-5 Final window of the DMP

- Go to Spectrum Control and click **Storage** → **V7000 VersaStack** → **Nodes**. When you look at this period, you see that I/O continued, even when the cables were unplugged.

Figure 15-6 shows Spectrum Control showing the total I/O rate and overall response time. The FC cables were removed from 10:30 - 10:45, and you can see that I/O continued throughout this period and response time stayed constant.




Figure 15-6 Spectrum Control chart showing total I/O and overall response time

15.2.2 Unexpected node failure

A Storwize V7000 storage system is an active/active storage controller that seamlessly allows for the failure of one node.

To simulate this failure, complete the following steps:

1. Physically remove one of the nodes from the enclosure. This is not a recommended action in an actual production environment, but is done to demonstrate various features only.

Figure 15-7 shows the performance window on the Storwize V7000 GUI. I/O is running and the health status is green. The Storwize V7000 performance window shows only 5 minutes of data.

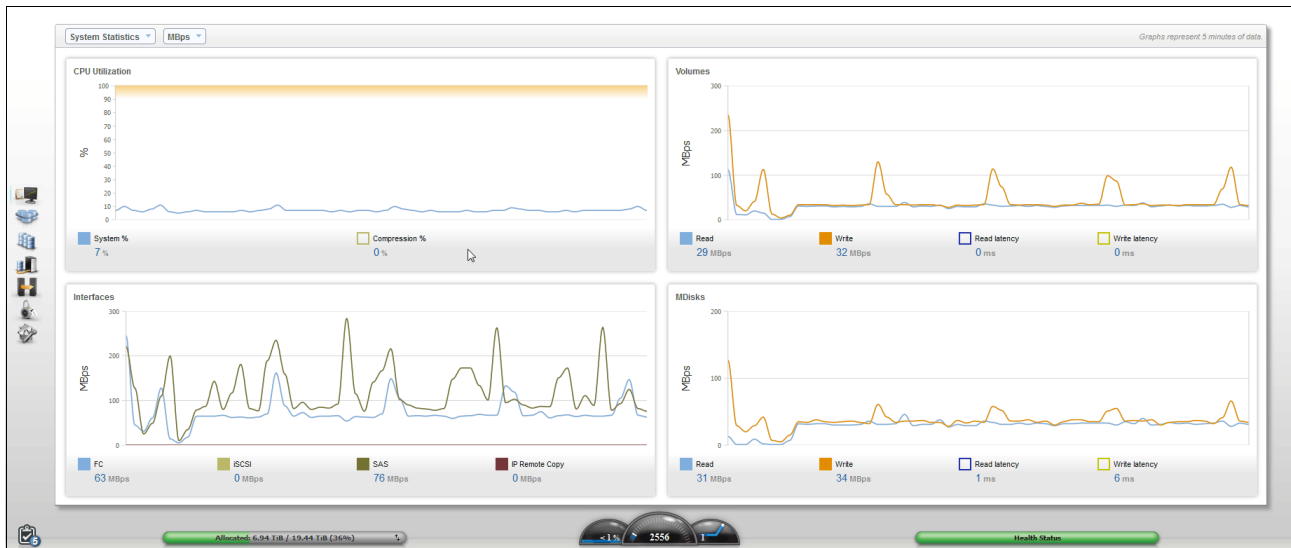


Figure 15-7 The performance window of the Storwize V7000 GUI

2. Remove the control node (node 2 in this case), which causes the cluster IP to fail over from node 2 to node 1. You briefly lose access to the GUI.

Figure 15-8 shows that removing the control node takes the GUI offline, as shown in the upper right of the window.




Figure 15-8 The GUI is offline

3. Access the GUI again by refreshing the GUI after a few minutes. There are errors in the event log. For more information, go to the System tab in Monitoring.

Figure 15-9 shows the System window, which shows both enclosures with an error within the control enclosure, that is, enclosure 1.




Figure 15-9 System tab showing an error in the control enclosure

4. Rotate the enclosure by using the red arrow, and hover your cursor over the canister to see more information.

Figure 15-10 shows you hovering the cursor over the canister, which shows its ID, state, configuration node, WWNN, and service IP.

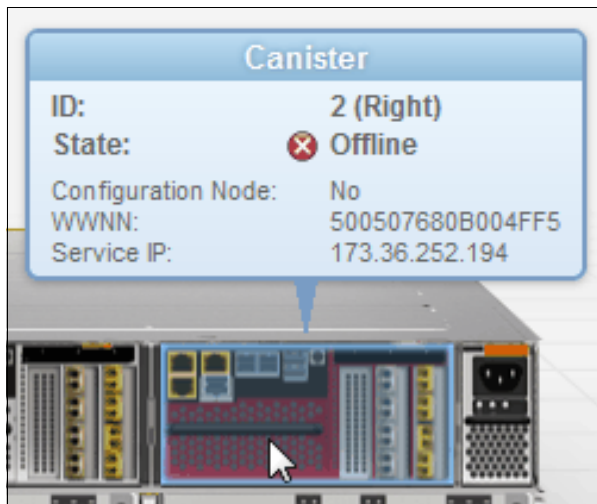


Figure 15-10 The canister is offline in the GUI

5. Reinsert node 2. When it starts, it seamlessly joins the cluster, and the systems window updates to show that it joined the cluster.

Figure 15-11 on page 429 shows the fully recovered cluster, which shows no errors.




Figure 15-11 The systems window after reinserting the node

- We can confirm on Spectrum Control that I/O continued throughout by clicking **Storage** → **V7000 VersaStack** → **Nodes** and selecting the period that the node was removed. The node was removed between 15:15 and 15:30.

Figure 15-12 shows Spectrum Control displaying the read and write I/O during the time when a node was removed from the cluster.




Figure 15-12 Show read and write I/O on spectrum control.

As previously explained, with only one node active, the cache is immediately flushed to disk, so the host does not write over data on cache that has yet to be destaged. This means that you have a write cache hit of 0% when a node is removed, which can be shown with the analytics available to you through Spectrum Control.

To view the write cache hit percent, press the + next to Metrics.

Figure 15-13 shows the Spectrum Control window for the Storwize V7000 storage system; the + expands the metrics that are available to view.

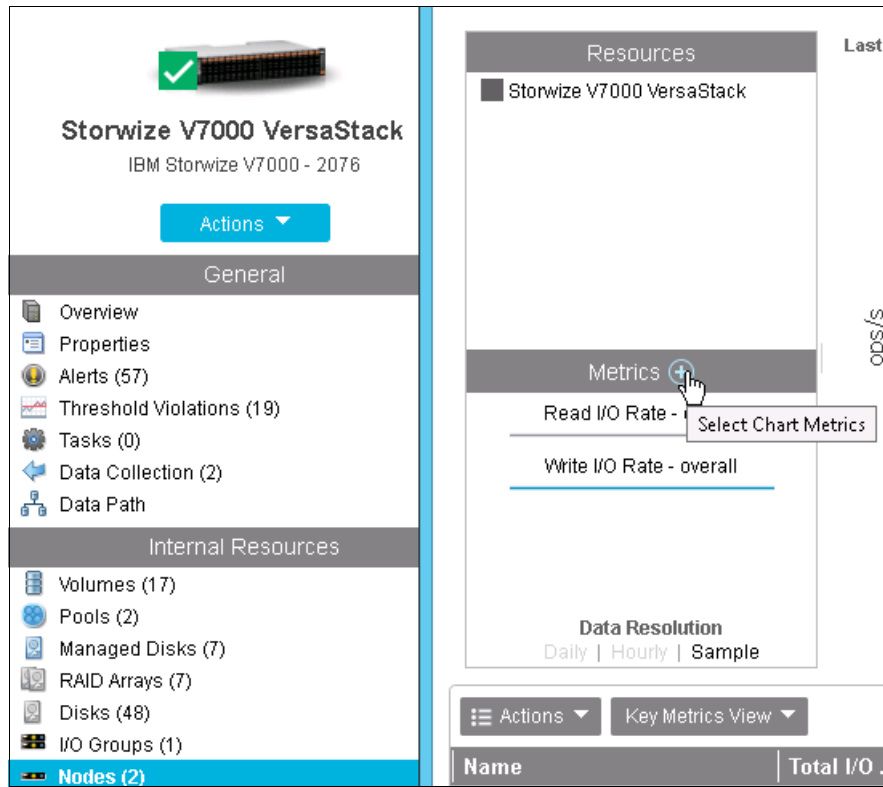


Figure 15-13 The Spectrum Control view of the Storwize V7000 storage system

You can change the metrics that you want to display. In our example, we display Cache Write Delay and Cache Hit Percent.

Figure 15-14 on page 431 shows the different metrics that are available, Cache Write Delay and Cache Hit Percent are selected and everything else is clear.

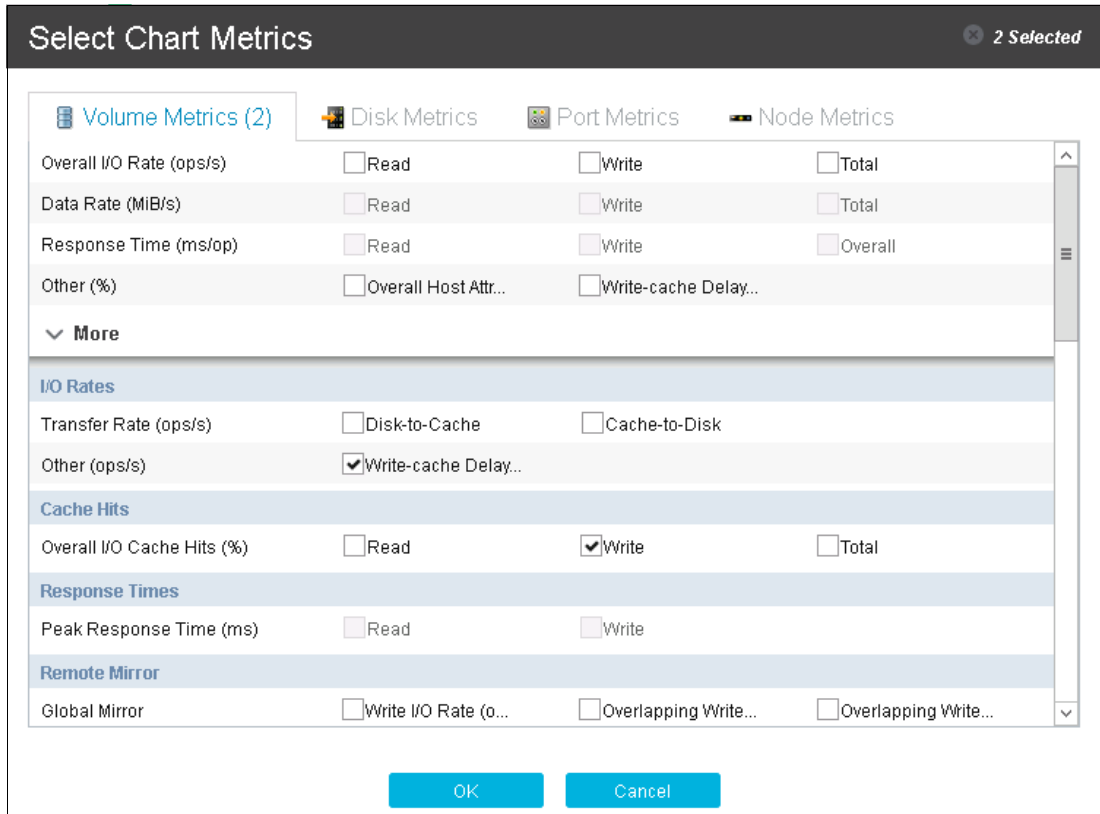


Figure 15-14 The different Spectrum Control metrics that are available

Selecting the period that we are interested in shows the cache hits dropping from 100% to 0% when one canister is removed and then returning to 100% when the canister is returned.

Figure 15-15 shows the Write Delay and Cache Hit Percent, and the node was removed between 10:15 and 10:30.

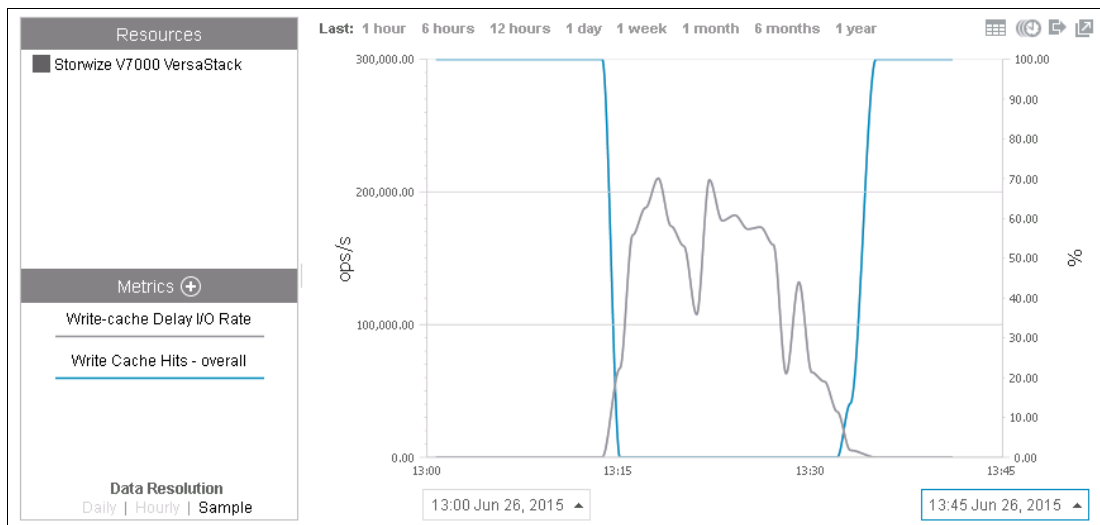


Figure 15-15 Use Spectrum Control to show the behavior of a cache with one node and with two nodes

15.3 Cisco Nexus devices

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single PortChannel to a third device. The third device can be a Cisco Nexus 2000 Series Fabric Extender or a switch, server, or any other networking device. A vPC can provide Layer 2 multipathing, which you can use to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

Figure 15-16 shows the Cisco Nexus vPC physical and logical topology.




Figure 15-16 Cisco Nexus vPC topology

A vPC provides the following benefits:

- ▶ Allows a single device to use a PortChannel across two upstream devices
- ▶ Eliminates Spanning Tree Protocol blocked ports
- ▶ Provides a loop-free topology
- ▶ Uses all available uplink bandwidth
- ▶ Provides fast convergence if either the link or a device fails
- ▶ Provides link-level resiliency
- ▶ Helps ensure high availability

15.3.1 vPC peer switch failure validation

This validation scenario describes a vPC peer switch failure by bringing down one of the Nexus 9372 PX switches. This scenario highlights the high availability and redundancy of Nexus switches in the VersaStack environment.

Test procedure

Figure 15-17 on page 433 shows the status of vPC configuration when both Nexus 9372 peer switches are up and running.

Figure 15-17 on page 433 showing the Cisco Nexus vPC status.

```

N9K-A# sh vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id      : 101
Peer status        : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 inconsistency reason : Consistency Check Not Performed
vPC role           : primary
Number of vPCs configured : 2
Peer Gateway       : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status : Enabled (timeout = 240 seconds)

vPC Peer-link status
-----
id  Port  Status Active vlans
-----
1   Po10  up    1,30,40,50,60

vPC status
-----
id  Port  Status Consistency Reason      Active vlans
-----
13  Po13  up    success  success                    1,30,40,50,60
14  Po14  up    success  success                    1,30,40,50,60

N9K-A#

N9K-B# sh vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id      : 101
Peer status        : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 inconsistency reason : Consistency Check Not Performed
vPC role           : secondary
Number of vPCs configured : 2
Peer Gateway       : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status : Enabled (timeout = 240 seconds)

vPC Peer-link status
-----
id  Port  Status Active vlans
-----
1   Po10  up    1,30,40,50,60

vPC status
-----
id  Port  Status Consistency Reason      Active vlans
-----
13  Po13  up    success  success                    1,30,40,50,60
14  Po14  up    success  success                    1,30,40,50,60

N9K-B#

```

Figure 15-17 Cisco Nexus vPC status

Complete the following steps:

1. Before reloading the Nexus switch with the primary role, initiate an OLTP workload on the SQL clustered instance from outside the VersaStack environment. The tool that is used in this example for generating a workload is HammerDB.

Figure 15-18 shows the HammerDB OLTP workload running.

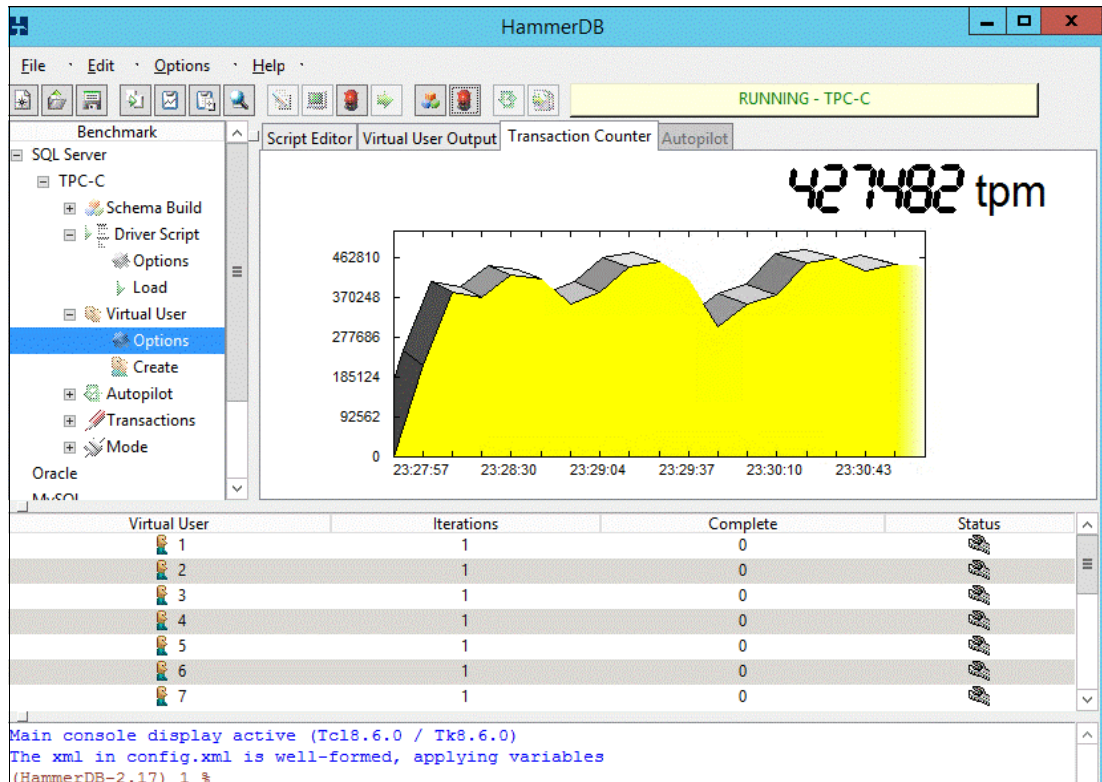


Figure 15-18 HammerDB OLTP workload tool

2. Connect to the Nexus 9372 switch with the vPC role as the primary through Secure Shell and run **reload**.

Figure 15-19 shows the **reload** command that is run on the primary switch.

```
N9K-A# reload
This command will reboot the system. (y/n)? [n] y
```

Figure 15-19 Cisco Nexus command line

Test observation

When the primary Nexus peer switch was reloading, the secondary peer switch that is running assumes the vPC role of operational primary.

The peer status and vPC keep-alive status are seen as Down and in a suspended state, as shown in Figure 15-20.

The screenshot shows the HammerDB interface on the left and a PuTTY terminal window on the right. The HammerDB interface displays a benchmark running at 440256 tpm. The PuTTY terminal window shows the output of the 'sh vpc brief' command on a Cisco Nexus switch, indicating that the vPC peer link is down and the vPC keep-alive status is suspended.

```

Legend:
    (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 101
Peer status            : peer link is down
vPC keep-alive status  : Suspended (Destination IP not reachable)
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 inconsistency reason : Consistency Check Not Performed
vPC role                : secondary, operational primary
Number of vPCs configured : 2
Peer Gateway           : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status   : Enabled (timeout = 240 seconds)

vPC Peer-link status
-----
id  Port  Status Active vlans
-----
1   Po10  down  -

vPC status
-----
id  Port  Status Consistency Reason Active vlans
-----
13  Po13  up     success  success  1,30,40,50,60
14  Po14  up     success  success  1,30,40,50,60

```

Figure 15-20 Cisco Nexus vPC Peer Status down

During the reload of the primary switch, half of the network bandwidth is lost and the remaining vPC switch maintains the network connectivity. There is no impact to the vPC operation or data forwarding.

Figure 15-20 on page 434 shows the vPC peer status as Down, but the data and control planes are still operational with the OLTP workload also running in the background.

After the reloaded switch comes back up, the vPC status is back to normal, as shown in Figure 15-21. Network bandwidth is restored to full capacity.

```

192.168.10.31 - PuTTY
N9K-A# sh vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id       : 101
Peer status         : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 inconsistency reason : Consistency Check Not Performed
vPC role            : primary, operational secondary
Number of vPCs configured : 2
Peer Gateway        : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status : Enabled (timeout = 240 seconds)

vPC Peer-link status
-----
id  Port  Status Active vlans
-----
1   Po10  up    1,30,40,50,60

vPC status
-----
id  Port  Status Consistency Reason      Active vlans
-----
13  Po13  down*  Not      Consistency Check Not      -
      Applicable Performed
14  Po14  down*  Not      Consistency Check Not      -
      Applicable Performed

N9K-A#

192.168.10.32 - PuTTY
-----
13  Po13  up    success  success      1,30,40,50,60
14  Po14  up    success  success      1,30,40,50,60

N9K-B# sh vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id       : 101
Peer status         : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : success
Type-2 inconsistency reason : Consistency Check Not Performed
vPC role            : secondary, operational primary
Number of vPCs configured : 2
Peer Gateway        : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status : Enabled (timeout = 240 seconds)

vPC Peer-link status
-----
id  Port  Status Active vlans
-----
1   Po10  up    1,30,40,50,60

vPC status
-----
id  Port  Status Consistency Reason      Active vlans
-----
13  Po13  up    success  success      1,30,40,50,60
14  Po14  up    success  success      1,30,40,50,60

N9K-B#
  
```

Figure 15-21 Cisco Nexus Switch vPC Status as Normal

Figure 15-21 shows the Cisco Nexus vPC status restored to normal state after the failed switch successfully came back up.

15.4 Cisco UCS service profile

Conceptually, a *service profile* is an extension of the VM abstraction that is applied to physical servers. The definition is expanded to include elements of the environment that span the entire datacenter, encapsulating the server identity (LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and quality of service (QoS) policies) in logical “service profiles” that can be dynamically created and associated with any physical server in the system within minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation. It enables migration of identities between servers in the environment without requiring any physical configuration changes, and facilitates rapid bare-metal provisioning of replacements for failed servers.

Service profiles also include operational policy information, such as information about firmware versions.

This highly dynamic environment can be adapted to meet rapidly changing needs in today's datacenters with just-in-time deployment of new computing resources and reliable movement of traditional and virtual workloads. Datacenter administrators can now focus on addressing business policies and data access on the basis of application and service requirements, rather than physical server connectivity and configurations.

Service profiles can be abstracted from the specifics of a given server to create a service profile template, which defines policies that can be applied any number of times to provision any number of servers. Service profile templates help enable large-scale operations in which many servers are provisioned as easily as a single server.

In addition, by using service profiles, Cisco UCS Manager provides logical grouping capabilities for both physical servers and service profiles and their associated templates. This pooling or grouping, combined with fine-grained role-based access, allows businesses to treat a farm of compute blades as a flexible resource pool that can be reallocated in real time to meet their changing needs, while maintaining any organizational overlay on the environment that they want.

Figure 15-22 shows the Cisco UCS service profile incorporating a complete metadata description of the information that is required to provision a server in a datacenter, including storage, network, and operational policies.

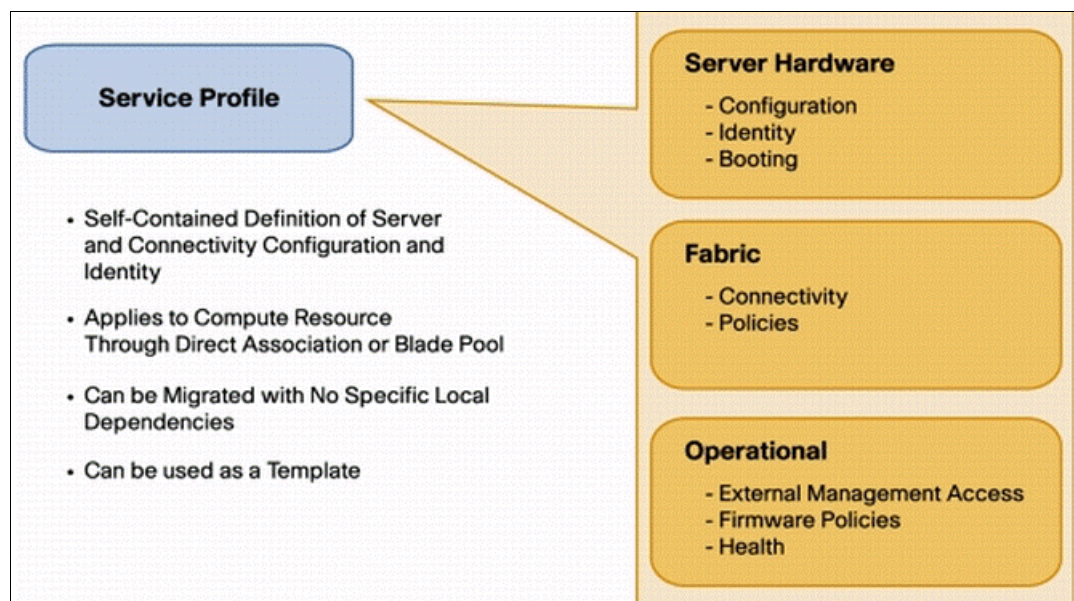


Figure 15-22 Cisco UCS service profile

15.4.1 Service profile migration validation

This validation scenario describes a use case of a Cisco UCS service profile migration in case there is an unplanned Cisco UCS B200 M4 hardware failure. This scenario is tested on a server that boots from SAN and needs spare hardware to replace the failed one.

Test procedure

Complete the following steps:

1. Power off the Cisco UCS B200 M4 server in slot 1 to simulate the hardware failure scenario.

Figure 15-23 shows a decommissioned server in Cisco UCS Manager.

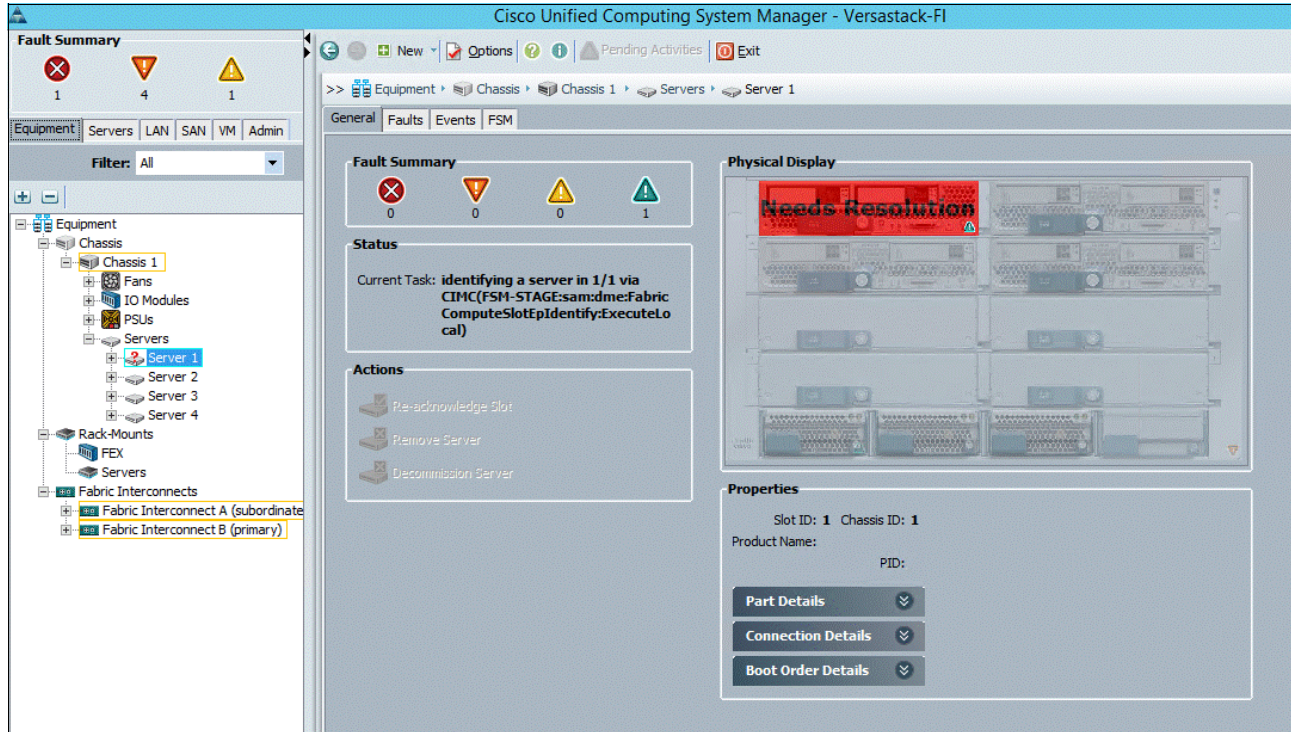


Figure 15-23 Cisco UCS Manager showing a decommissioned blade

2. Disassociate the service profile from the failed blade server.
3. Decommission the blade and swap it with a new blade with an equal configuration.

4. Reacknowledge the slot. The new blade is discovered by the UCSM.

Figure 15-24 shows a new blade being discovered in Cisco UCS Manager.



Figure 15-24 Cisco UCS Manager new blade discovery

5. Power off the B200 M4 server in slot 1 to simulate the hardware failure scenario.

6. Reassociate the service profile to the new hardware. Figure 15-25 shows the service profile association in Cisco UCS Manager.

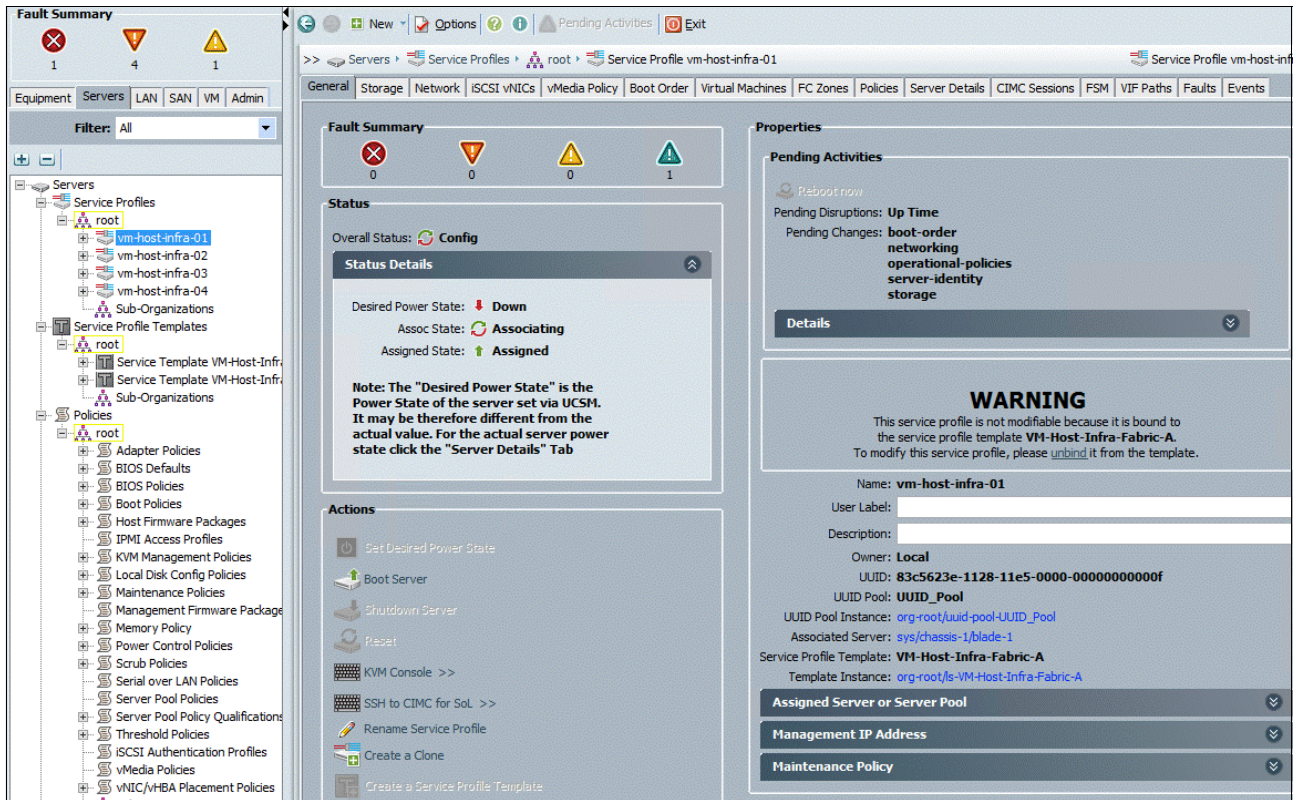


Figure 15-25 Service profile association in Cisco UCS Manager

Test observations

The service profile migration from the failed hardware to the new hardware was successful and the new server booted from SAN successfully.

Figure 15-26 shows the vSphere ESXi starting after the successful migration of the Cisco UCS service profile.

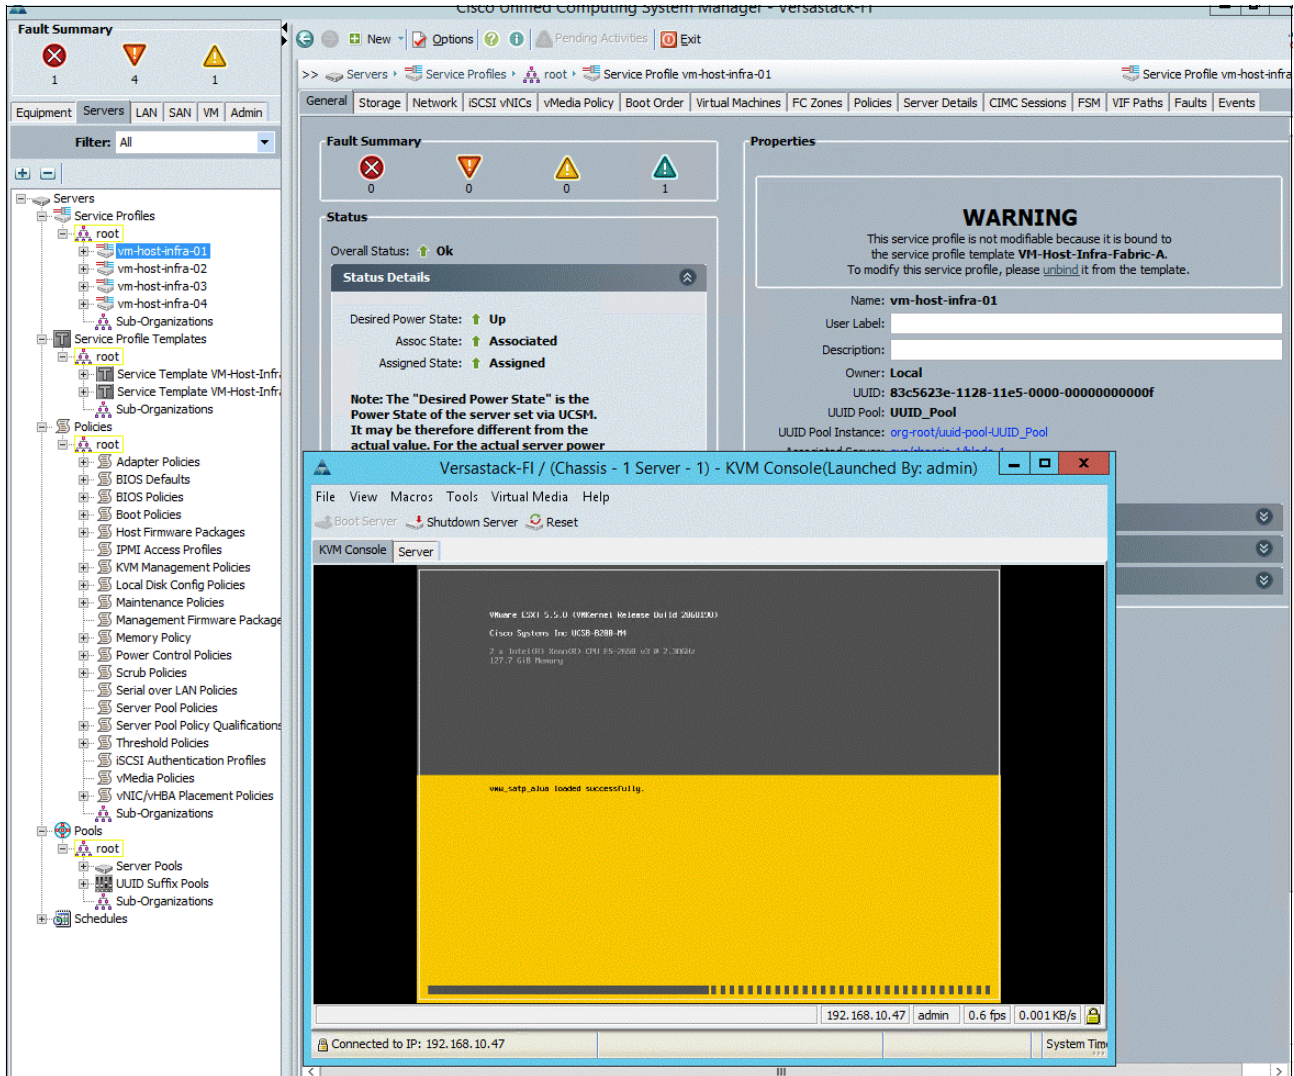


Figure 15-26 ESXi starting

The WSFC active node running on the failed ESXi host did not migrate to the second ESXi host because the vSphere HA/Distributed Resource Scheduler (DRS) anti-affinity rule is configured.

The WSFC and SQL Server FCI active node failed over successfully to the second virtual machine node running on the other ESXi host.

A couple of other VMs with a default vSphere HA/DRS configuration migrated and restarted on the second ESXi host successfully.

All the above outcomes were the expected behavior, and the services recovered quickly with a minimum of downtime.



A

Windows Active Directory and running configurations

This appendix shows how to build Windows Active Directory Server virtual machines (VMs), and the running configurations for the Nexus-A and Nexus-B switches.

Building Windows Active Directory Server virtual machines

To build an Active Directory Server virtual machine (VM) for the vm-host-infra-01 ESXi host, complete the following steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the vSphere Client, select the host in the inventory pane.
3. Right-click the host and select **New Virtual Machine**.
4. Select **Custom** and click **Next**.
5. Enter a name for the VM. Click **Next**.
6. Select infra_datastore_1. Click **Next**.
7. Select **Virtual Machine Version: 10**. Click **Next**.
8. Verify that the **Windows** option and the **Microsoft Windows Server 2012 R2 (64-bit) version** are selected. Click **Next**.
9. Select two virtual sockets and one core per virtual socket. Click **Next**.
10. Select 4 GB of memory. Click **Next**.
11. Select one network interface card (NIC).
12. For NIC 1, select the **MGMT Network** option and the VMXNET 3 adapter. Click **Next**.
13. Keep the **LSI Logic SAS** option for the SCSI controller selected. Click **Next**.
14. Keep the **Create a New Virtual Disk** option selected. Click **Next**.
15. Make the disk size at least 60 GB. Click **Next**.
16. Click **Next**.
17. Select the **Edit the Virtual Machine Settings Before Completion** check box. Click **Continue**.
18. Click the **Options** tab.
19. Select **Boot Options**.
20. Select the **Force BIOS Setup** check box.
21. Click **Finish**.
22. From the left pane, expand the host field by clicking the plus sign (+).
23. Right-click the newly created AD Server VM and click **Open Console**.
24. Click the third button (green right arrow) to power on the VM.
25. Click the ninth button (CD with a wrench) to map the Windows Server 2012 R2 ISO, and then select **Connect to ISO Image on Local Disk**.
26. Go to the Windows Server 2008 R2 SP1 ISO, select it, and click **Open**.
27. Click in the BIOS Setup Utility window and use the right arrow key to go to the Boot menu. Use the down arrow key to select CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.
28. The Windows Installer boots. Select the appropriate language, time and currency format, and keyboard. Click **Next**.
29. Click **Install now**.
30. Make sure that the **Windows Server 2012 R2 Standard (Full Installation)** option is selected. Click **Next**.

31. Read and accept the license terms and click **Next**.
32. Select **Custom (Advanced)**. Make sure that **Disk 0 Unallocated Space** is selected. Click **Next** to allow the Windows installation to complete.
33. After the Windows installation is complete and the VM restarts, click **OK** to set the Administrator password.
34. Enter and confirm the Administrator password and click the blue arrow to log in. Click **OK** to confirm the password change.
35. After logging in to the VM desktop, from the VM console window, select the **VM** menu. Under Guest, select **Install/Upgrade VMware Tools**. Click **OK**.
36. If prompted to eject the Windows installation media before running the setup for the VMware tools, click **OK**, then click **OK** again.
37. In the dialog box, select **Run setup64.exe**.
38. In the VMware Tools installer window, click **Next**.
39. Make sure that **Typical** is selected and click **Next**.
40. Click **Install**.
41. Click **Finish**.
42. Click **Yes** to restart the VM.
43. After the restart is complete, select the **VM** menu. Under Guest, select **Send Ctrl+Alt+Del**. Then, enter the password to log in to the VM.
44. Set the time zone for the VM, IP address, gateway, and host name.

Note: A restart is required.

45. If necessary, activate Windows.
46. Download and install all the required Windows updates.

Note: This process requires several restarts.

47. Open Server Manager.
48. On the left pane, click **Roles**, then select **Add Roles** on the right.
49. Click **Next**.
50. In the list, select the **Active Directory Domain Services** check box.
51. In the dialog box that opens, click **Add Required Features** to add .NET Framework 3.5.1.
52. Click **Next**.
53. Click **Next**.
54. Click **Install**.
55. In the middle of the window, click **Close this wizard and launch the Active Directory Domain Services Installation Wizard (dcpromo.exe)**.
56. In the Active Directory Domain Services Installation wizard, click **Next**.
57. Click **Next**.
58. Select **Create a new domain in a new forest** and click **Next**.
59. Type the FQDN of the Windows domain for this VersaStack environment and click **Next**.

60. Select the appropriate forest functional level and click **Next**
61. Keep DNS server selected and click **Next**.
62. If one or more DNS servers exist that this domain can resolve from, select **Yes** to create a DNS delegation. If this AD server is being created on an isolated network, select **No** to not create a DNS delegation. The remaining steps in this procedure assume that a DNS delegation is not created. Click **Next**.
63. Click **Next** to accept the default locations for database and log files.
64. Enter and confirm <<var_password>> for the Directory Services Restore Mode Administrator Password. Click **Next**.
65. Review the Summary information and click **Next**. Active Directory Domain Services installs.
66. Click **Finish**.
67. Click **Restart Now** to restart the AD Server.
68. After the machine restarts, log in as the domain administrator.
69. Open the DNS Manager by clicking **Start** → **Administrative Tools** → **DNS**.
70. Optional: Add Reverse Lookup Zones for your IP address ranges.
71. Expand the Server and Forward Lookup Zones. Select the zone for the domain. Right-click and select **New Host (A or AAAA)**. Populate the DNS Server with Host Records for all components in the VersaStack environment.
72. Optional: Build a second AD server VM. Add this server to the newly created Windows Domain and activate Windows. Install Active Directory Domain Services on this machine. Start `dcpromo.exe` at the end of this installation. Choose to add a domain controller to a domain in an existing forest. Add this domain controller to the domain created earlier. Complete the installation of this second domain controller. After vCenter Server is installed, affinity rules can be created to keep the two AD servers running on different hosts.

Nexus 9000 running configuration

This section shows the **running config** information for Nexus-A and Nexus-B.

These configurations are generated by running **running-config**, as shown in Example A-1.

Example A-1 The running-config command

```
N9K-A# sh running-config
```

Nexus 9000 A running configuration

Here is the content of the Nexus 9000 A running configuration:

```
!version 6.1(2)I3(3a)
switchname N9K-A
vdc N9K-A id 1
  allocate interface Ethernet1/1-54
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 512
  limit-resource u4route-mem minimum 248 maximum 248
```

```

    limit-resource u6route-mem minimum 96 maximum 96
    limit-resource m4route-mem minimum 58 maximum 58
    limit-resource m6route-mem minimum 8 maximum 8
cfs eth distribute
feature udl
feature lacp
feature vpc
username admin password 5 $1$vFdUE8vJ$CDbxkfFaGGQjCaxM6JKsz. role network-admin
ip domain-lookup
copp profile strict
snmp-server user admin network-admin auth md5 0x546a7b8b3b91374ff18cdc3997e0d17
2 priv 0x546a7b8b3b91374ff18cdc3997e0d172 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vlan 1,30,40,50,60
vlan 30
    name vMotion
vlan 40
    name WinClus
vlan 50
    name WinCSV
vlan 60
    name Backup
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
vrf context management
    ip route 0.0.0.0/0 192.168.10.1
vpc domain 101
    peer-switch
    role priority 10
    peer-keepalive destination 192.168.10.32 source 192.168.10.31
    delay restore 150
    peer-gateway
    auto-recovery
    ip arp synchronize
interface port-channel10
    description vPC peer-link
    switchport mode trunk
    spanning-tree port type network
    vpc peer-link
interface port-channel13
    description to FI-A
    switchport mode trunk
    spanning-tree port type edge trunk
    mtu 9216
    vpc 13
interface port-channel14
    description to FI-B
    switchport mode trunk
    spanning-tree port type edge trunk
    mtu 9216

```

```
vpc 14
interface Ethernet1/1
interface Ethernet1/2
interface Ethernet1/3
interface Ethernet1/4
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
  description FI-A:1/25
  switchport mode trunk
  mtu 9216
  channel-group 13 mode active
interface Ethernet1/26
  description FI-B:1/26
  switchport mode trunk
  mtu 9216
  channel-group 14 mode active
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
```



```

interface Ethernet1/47
  description vPC Peer N9K-B:1/47
  switchport mode trunk
  channel-group 10 mode active
interface Ethernet1/48
  description vPC Peer N9K-B:1/48
  switchport mode trunk
  channel-group 10 mode active
interface Ethernet1/49
interface Ethernet1/50
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
  vrf member management
  ip address 192.168.10.31/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I3.3a.bin
N9K-A#

```

Nexus 9000 B running configuration

Here is the content of the Nexus 9000 B running configuration:

```

version 7.0(3)I1(1a)
switchname N9K-B
vdc N9K-B id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
cfs eth distribute
feature udd
feature lacp
feature vpc
username admin password 5 $1$h0zBLP15$ZFoDlelseUIJ3gX6ugx54. role network-admin
ip domain-lookup
copp profile strict
snmp-server user admin network-admin auth md5 0x672fc1ebf92b0e84c5443ce2f1c34b69
priv
  0x672fc1ebf92b0e84c5443ce2f1c34b69 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
vlan 1,30,40,50,60
vlan 30
  name vMotion
vlan 40

```

```

    name WinClus
vlan 50
    name WinCSV
vlan 60
    name Backup
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
spanning-tree port type network default
vrf context management
    ip route 0.0.0.0/0 192.168.10.1
vpc domain 101
    peer-switch
    role priority 20
    peer-keepalive destination 192.168.10.31 source 192.168.10.32
    delay restore 150
    peer-gateway
    auto-recovery
    ip arp synchronize
interface port-channel10
    description vPC peer-link
    switchport mode trunk
    spanning-tree port type network
    vpc peer-link
interface port-channel13
    description FI-A
    switchport mode trunk
    spanning-tree port type edge trunk
    mtu 9216
    vpc 13
interface port-channel14
    description to FI-B
    switchport mode trunk
    spanning-tree port type edge trunk
    mtu 9216
    vpc 14
interface Ethernet1/1
interface Ethernet1/2
interface Ethernet1/3
interface Ethernet1/4
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20

```

```
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
  description FI-B:1/25
  switchport mode trunk
  mtu 9216
  channel-group 14 mode active
interface Ethernet1/26
  description FI-A:1/26
  switchport mode trunk
  mtu 9216
  channel-group 13 mode active
interface Ethernet1/27
interface Ethernet1/28
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
interface Ethernet1/47
  description vPC Peer N9K-A:1/47
  switchport mode trunk
  channel-group 10 mode active
interface Ethernet1/48
  description vPC Peer N9K-A:1/48
  switchport mode trunk
  channel-group 10 mode active
interface Ethernet1/49
interface Ethernet1/50
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
  vrf member management
  ip address 192.168.10.32/24
line console
line vty
boot nxos bootflash:/n9000-dk9.7.0.3.I1.1a.bin
```


Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are 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.

- ▶ *Implementing the IBM System Storage SAN Volume Controller V7.4*, SG24-7933
- ▶ *Implementing the IBM Storwize V7000 V7.4*, SG24-7938
- ▶ *Introducing and Implementing IBM FlashSystem V9000*, SG24-8273
- ▶ *IBM Real-time Compression in IBM SAN Volume Controller and IBM Storwize V7000*, REDP-4859

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=E0W4gggN99o>
- ▶ 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>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Redbooks

VersaStack Solution by Cisco and IBM with IBM DB2, IBM Spectrum Control, and IBM Spectrum Protect

SG24-8302-00
ISBN 073844135X



(1.0" spine)
0.875" x 1.498"
460 <-> 788 pages



SG24-8302-00

ISBN 073844135X

Printed in U.S.A.

Get connected

